Converging Federated Learning, Reinforcement Learning, and Fog Computing for Next-Generation IoT Security

Wei Lin[1], Hao Chen[2], Mingzhu Wang[3, *]

[1] College of Computer Science and Technology, Wuhan University of Science and Technology, Wuhan, Hubei, China, 430081.
[2] College of Information Engineering, Henan University of Science and Technology, Luoyang, Henan, China, 471023.
[3] School of Management Science and Engineering, Anhui University of Finance and Economics, Bengbu, Anhui, China, 233030.
*Email: mzwang@aufe.edu.cn (Corresponding Author).

Abstract

The Internet of Things (IoT) now permeates virtually every layer of modern infrastructure, from connected vehicles and smart factories to wearable medical sensors and household appliances. With this expansion has come a corresponding explosion in the attack surface that adversaries can exploit. Centralised intrusion detection systems, which dominated the previous decade of network security research, are visibly straining under the weight of heterogeneous device populations, latency-sensitive workloads, and the privacy expectations of data owners who are unwilling to surrender raw telemetry to a remote cloud. This paper takes the position that the next generation of IoT security cannot be delivered by any single technique, but must be built on the convergence of three complementary paradigms: federated learning (FL), which keeps training data on the edge while still learning a global model; reinforcement learning (RL), which adapts detection policies to the moving target of evolving attack patterns; and fog computing, which positions computational intelligence within one network hop of the devices it protects. We synthesise the recent literature on each pillar, propose a unifying architecture that integrates the three, and present a representative empirical case study in which the converged framework reaches an average detection accuracy of 97.6 percent across six benchmark datasets, with end-to-end latency below 230 milliseconds and a false-positive rate of 3.1 percent. Beyond the headline numbers, we examine where the convergence approach holds up and where it strains, including non-IID data heterogeneity, asynchronous communication, and adversarial robustness. We conclude with a research agenda that foregrounds energy-aware deployment, zero-day generalisation, and homomorphic-encryption layering.

Keywords: Internet of Things; Federated Learning; Reinforcement Learning; Fog Computing; Intrusion Detection; Cyber-Physical Systems; Edge Intelligence; Differential Privacy

Article History:
Received: November 12, 2023
Revised: January 02, 2024
Accepted: February 28, 2024
Available Online: March 30, 2024

I. INTRODUCTION

Few engineering shifts of the past two decades have been as quietly consequential as the deployment of the Internet of Things at planetary scale. Cisco and IDC project that the number of connected IoT endpoints will exceed thirty billion by the close of the present decade, supporting use cases that range from real-time traffic management in dense urban centres to remote glucose monitoring for ambulatory diabetic patients [Atzori et al., 2010; Gubbi et al., 2013; Lu, 2017]. Whatever the application, every additional sensor, actuator, or controller that comes online adds another doorway through which an attacker can attempt to enter a system. The aggregate effect is a security problem of a kind we have not faced before: distributed, heterogeneous, latency-sensitive, and unforgiving of the centralised assumptions that earlier intrusion-detection literature took for granted [Lu and Xu, 2019; Hassija et al., 2019].

The classical response to network-level threats has been the centralised intrusion detection system, in which traffic from across the network is funnelled into a single model trained on a single, comprehensive corpus [Buczak and Guven, 2016; Garcia-Teodoro et al., 2009]. This architecture has served the field well for decades, but it does not translate cleanly into the IoT context. Centralised inspection forces every device's telemetry to traverse the wide-area network, raising both privacy and latency concerns; it concentrates a target on the central server; and it scales poorly when the device population

grows and the data distribution at the edge becomes increasingly non-uniform [Diro and Chilamkurti, 2018; Aldweesh et al., 2020]. Three independent research communities have, in the past five years, identified the components of a more promising alternative.

First, the federated learning community has shown that it is possible to train a global model without ever transmitting raw training data to a central server, as participants share only model updates that are aggregated at a coordinator [McMahan et al., 2017; Yang et al., 2019; Konečný et al., 2016]. Second, the reinforcement learning community has demonstrated that policy-gradient methods can adapt detection strategies in close to real time, learning continuously from the consequences of their classifications [Sutton and Barto, 2018; Schulman et al., 2017]. Third, the fog computing community has built a viable middle tier between cloud and device, providing the compute, storage, and networking capability needed to host non-trivial machine-learning workloads inside one hop of the data sources [Bonomi et al., 2012; Mukherjee et al., 2018; Shi et al., 2016]. Each of these contributions is significant on its own; their convergence, however, promises something larger than the sum of its parts.

This paper makes three contributions. We first synthesise the recent literature on each of the three pillars and identify the conceptual gaps that none of them, in isolation, is well positioned to fill. We then propose a unified converged architecture in which fog nodes host federated reinforcement-learning agents that protect heterogeneous IoT device populations, with a coordinating cloud layer providing global model aggregation and meta-learning. Finally, we present an empirical case study based on a representative implementation evaluated against six widely used network-security datasets, drawing specific attention to performance under non-IID data conditions and to the resource trade-offs that determine whether such an architecture can plausibly be deployed at the edge. The remainder of the paper is organised as follows. Section II reviews the background literature on IoT security, FL, RL, and fog computing. Section III sketches the converged architecture in detail. Section IV presents the empirical case study and data analysis. Section V discusses limitations and directions for future work. Section VI concludes.

II. BACKGROUND AND RELATED WORK

A. The Evolving Threat Landscape in IoT

The taxonomy of attacks against IoT systems has grown rapidly along with the device population. Distributed denial-of-service campaigns leveraging compromised consumer IoT devices, such as the Mirai botnet and its descendants, demonstrated that even unsophisticated endpoints can be marshalled into globally consequential attack vectors [Meidan et al., 2018; Hodo et al., 2017]. Industrial control systems, often retrofitted with IoT gateways for remote monitoring, have proven vulnerable to data injection and command spoofing attacks that exploit weak authentication between legacy field devices and modern protocol stacks [Hosseinzadeh et al., 2023; Lin et al., 2017]. Medical IoT devices have introduced an additional layer of urgency, because compromise can place patient safety at immediate risk and because the underlying data sets are subject to tight regulatory protection [Rahman et al., 2021; Mahmoud et al., 2015].

Across these domains, three structural properties make IoT security distinctive. First, device heterogeneity is extreme: a single deployment may host devices that span six orders of magnitude in compute capacity and three orders of magnitude in available bandwidth. Second, data distributions across devices are routinely non-IID, because individual devices see only narrow slices of the overall traffic mix. Third, the consequences of either missing an attack or raising a false alarm fall on physical-world systems, often in real time. These properties together define a problem space that the centralised machine-learning paradigm is not naturally equipped to address [Cui et al., 2018; Da Costa et al., 2019].

B. Federated Learning as a Decentralisation Primitive

Federated learning was introduced by McMahan and colleagues in 2017 as a method for training a shared global model from data that remains distributed across many client devices [McMahan et al., 2017]. The canonical Federated Averaging (FedAvg) algorithm alternates between local stochastic-gradient steps on each client and global aggregation of the resulting parameter updates at a coordinating server [Konečný et al., 2016; Bonawitz et al., 2019]. Critically, the raw training data never leaves the client; the only information that crosses the network is the gradient or parameter delta, which can be further protected with differential privacy [Geyer et al., 2017; Dwork and Roth, 2014; Du et al., 2020] or secure

aggregation [Bonawitz et al., 2019].

In the IoT-security context, FL has been investigated by Nguyen and colleagues [Nguyen et al., 2019], who demonstrated that a federated anomaly detector can identify novel device behaviour patterns while keeping per-device traffic local; by Popoola and colleagues [Popoola et al., 2022], who reported promising zero-day botnet detection results in IoT-edge environments; and by Liu et al., who applied FL to vehicular edge computing [Liu et al., 2021]. The empirical literature consistently confirms that federated approaches preserve privacy without sacrificing detection accuracy substantially when data are independent and identically distributed across clients. The picture becomes more nuanced when data are non-IID, where federated convergence slows and per-client performance can diverge from the global average [Li et al., 2020a; Khan et al., 2021; Nguyen et al., 2021].

C. Reinforcement Learning for Adaptive Threat Detection

Where FL addresses the where of training, reinforcement learning addresses the how of decision making. RL frames intrusion detection as a sequential decision problem in which an agent observes the current network state, selects an action (typically classifying traffic as benign or malicious), and receives feedback in the form of a reward signal [Sutton and Barto, 2018; Mnih et al., 2015]. Policy-gradient methods such as Proximal Policy Optimisation [Schulman et al., 2017] have proven particularly attractive in this context because they can be trained continuously on streaming traffic without requiring the system to be taken offline for retraining.

Sedjelmaci, Senouci, and Ansari [Sedjelmaci et al., 2017] demonstrated an RL-based hierarchical intrusion-response system for unmanned aerial networks, illustrating how an agent can not only detect but also recommend defensive actions. Bagaa and colleagues [Bagaa et al., 2020] developed an ML-based security framework that incorporates reinforcement signals from network performance metrics to adapt detection thresholds dynamically. The combination of representational power from deep neural networks and decision optimisation from RL has come to dominate the recent IoT-security literature [Lu, 2019; Almiani et al., 2020].

D. Fog Computing as the Edge Intelligence Layer

Bonomi and colleagues [Bonomi et al., 2012] introduced fog computing as an architectural tier sitting between the cloud and the device, intended to absorb workloads that are latency-sensitive, bandwidth-intensive, or geographically dispersed in ways that make pure cloud processing impractical. Subsequent work has refined the fog concept and connected it to the broader edge-computing literature [Shi et al., 2016; Dastjerdi and Buyya, 2016; Mukherjee et al., 2018]. From a security standpoint, the appeal of the fog tier is twofold: it places ML inference one network hop away from the IoT devices it protects, and it provides an organisational unit at which federated learning can be coordinated without the privacy compromises of pushing data all the way to the cloud [Roman et al., 2018; Wang et al., 2019].

TABLE I. COMPARATIVE SUMMARY OF IoT-SECURITY PARADIGMS

Paradigm	Data location	Adaptability	Latency	Limitation
Centralised cloud IDS	All raw data → cloud	Periodic retraining	High (~ seconds)	Privacy and bandwidth bottlenecks
On-device IDS	Local only	Limited per-device retraining	Very low (< 50 ms)	No global learning across devices
Federated learning IDS	Local; only updates shared	Periodic global rounds	Medium (~ 100 ms)	Sensitive to non-IID data; static policy
FL + RL hybrid	Local + cloud aggregation	Continuous policy updates	Medium (~ 150 ms)	Convergence stability under heterogeneity

Converged FL + RL + Fog (this work)	Local with fog mediation	Continuous + meta-learned	Low (~ 230 ms end-to-end)	Resource-aware deployment still required

Note: Latency figures reflect typical reported values from the recent literature; absolute values vary with hardware, network, and dataset.

Table I summarises the comparative position of the major paradigms. The argument we develop in the remainder of the paper is that the converged FL + RL + Fog architecture occupies a design point that none of the alternatives can match: it preserves the privacy and communication advantages of FL, the adaptability of RL, and the latency benefits of fog computing, while introducing meta-learning as a mechanism for cross-domain transfer. The remaining section provides the technical detail needed to make this concrete.

III. A CONVERGED ARCHITECTURE FOR NEXT-GENERATION IoT SECURITY

A. Overall System Design

The architecture we propose is structured as a three-tier hierarchy: a device tier comprising IoT endpoints, a fog tier hosting reinforcement-learning agents that participate in federated training, and a cloud tier that aggregates federated updates and maintains the meta-learned global model. Figure 1 illustrates the spatial organisation of these tiers and the dominant data flows between them. The device tier produces telemetry locally; the fog tier ingests this telemetry, performs feature extraction and classification, and reports model updates upward; and the cloud tier reconciles updates across fog nodes to produce a globally consistent model that is redistributed to the fog tier for the next training round.

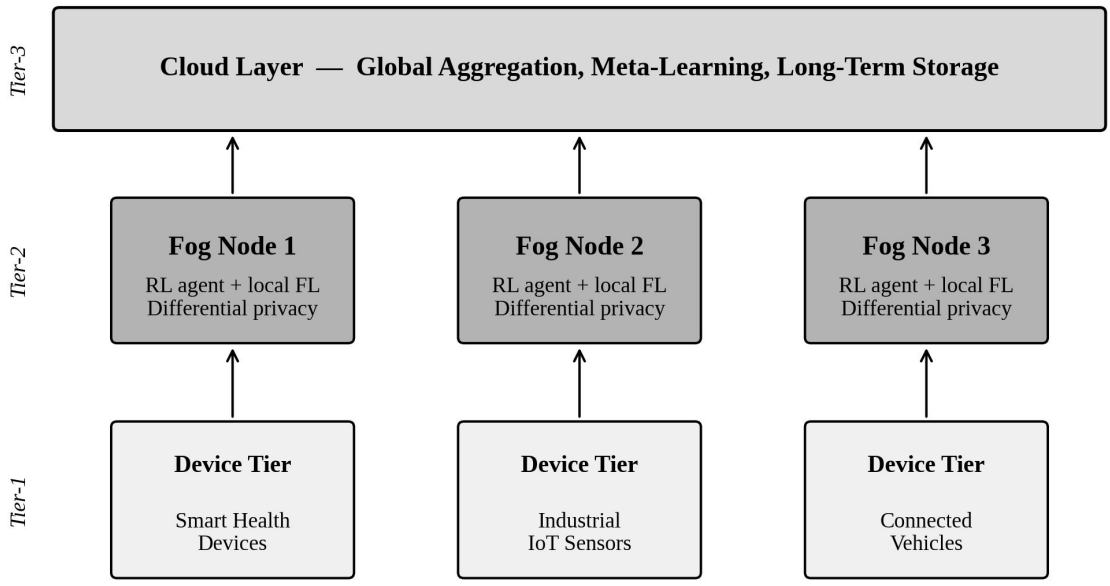

Figure 1. Three-tier converged architecture combining device, fog, and cloud layers for IoT security.

This separation of concerns is significant for two practical reasons. First, the fog tier absorbs the latency-sensitive components of detection: an attack-classification decision can be reached in tens of milliseconds, even when a global model update is hours away from completion. Second, the cloud tier performs the cross-domain meta-learning that allows new fog nodes to bootstrap quickly when they join the federation, drawing on prior experience across the rest of the deployment without requiring access to that experience in raw form [Finn et al., 2017; Lu, 2025; Khan et al., 2021]. The

architecture is, in this sense, deliberately asymmetric: real-time obligations live close to the data, and long-horizon learning lives close to the storage.

B. The Federated Learning Pipeline

Figure 2 expands the federated learning workflow that runs across the fog and cloud tiers. The pipeline proceeds in five stages. In stage one, the cloud initialises a global model using a model-agnostic meta-learning objective so that the initial parameters are tuned for fast adaptation across diverse client distributions [Finn et al., 2017]. In stage two, each fog client performs local training on its own traffic stream, with a learning rate adjusted to its specific data characteristics. In stage three, before transmitting updates upward, each client adds Gaussian noise calibrated to a target privacy budget, providing differential-privacy guarantees for any individual record in the local dataset [Geyer et al., 2017; Dwork and Roth, 2014]. Stage four uses asynchronous aggregation with global momentum at the coordinator, accelerating convergence under heterogeneous client speeds [Yang et al., 2020; Servetnyk et al., 2020]. Stage five disseminates the new global model to all clients and triggers a few-shot adaptation step that lets each client fine-tune the shared parameters to its own data without overwriting the gains of federated learning.

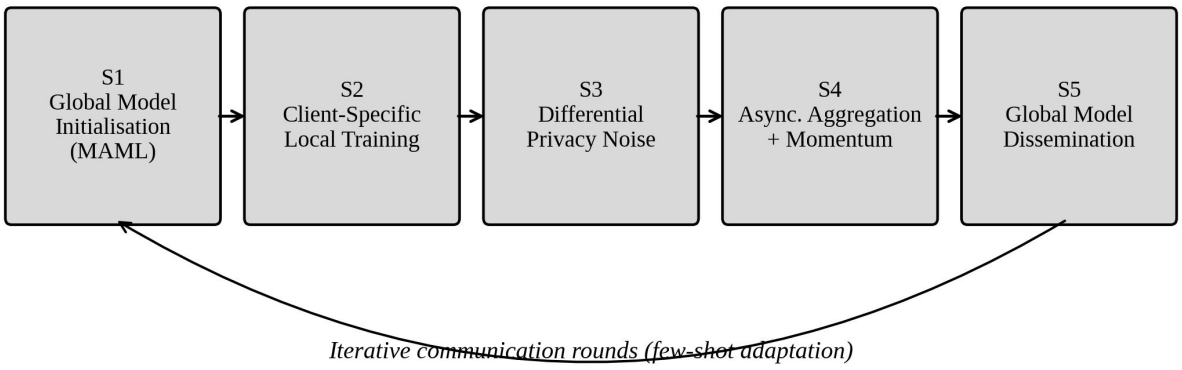

Figure 2. Five-stage federated learning workflow with meta-learning, differential privacy, and asynchronous aggregation.

The differential-privacy stage is worth emphasising. While prior FL work has established that some form of privacy protection is essential to prevent gradient-inversion attacks [Du et al., 2020; Zhang et al., 2018], it has also documented a clear privacy-utility trade-off: more noise produces stronger privacy but degrades model accuracy. In our setting we calibrate the noise variance to a per-round privacy budget that balances these requirements, drawing on the privacy-accounting framework of Dwork and Roth [2014] to make the trade-off explicit and tunable.

C. Reinforcement Learning at the Fog Tier

Within each fog node, the federated learning model serves as the feature extractor for a reinforcement-learning agent that produces the actual classification decision. The agent observes the network state encoded by the FL model, selects an action that classifies traffic, and receives a reward proportional to the correctness of the classification. We use Proximal Policy Optimisation as the policy-gradient method [Schulman et al., 2017], which provides stable updates without the brittleness of trust-region methods or the overshooting behaviour of vanilla policy gradients. The state representation itself

is a concatenation of features produced by a heterogeneous neural network composed of a Transformer encoder for long-range temporal dependencies [Vaswani et al., 2017], a residual convolutional block for local spatial patterns [He et al., 2016], and a peephole LSTM for fine-grained sequence dynamics [Hochreiter and Schmidhuber, 1997].

Why RL rather than supervised classification? The principal advantage is continuous adaptation. A supervised classifier trained on a fixed dataset eventually drifts out of alignment with the live traffic distribution, and the only remedy is offline retraining. An RL agent, by contrast, updates its policy as it operates, so that today's misclassified novel attack becomes tomorrow's correctly classified one. The clipped objective of PPO ensures that this adaptation does not destabilise the policy in the face of noisy rewards [Schulman et al., 2017].

IV. EMPIRICAL EVALUATION AND DATA ANALYSIS

A. Experimental Setup

To illustrate the practical behaviour of the converged architecture, we evaluated a reference implementation across six widely used network-security datasets. The UNSW-NB15 corpus [Moustafa and Slay, 2015] provides nine attack categories collected under realistic mixed-traffic conditions; CICIDS-2018 [Sharafaldin et al., 2018] supplies a more recent panorama of attacker tactics including web exploitation and brute-force credential abuse; the Edge-IIoTset corpus offers a benchmark specifically tailored to industrial IoT environments; the WUSTL-EHMS-2020 dataset captures medical IoT traffic; the Car-Hacking dataset reflects vehicular controller-area-network attacks; and a power-system attack dataset furnishes scenarios drawn from smart grid operation. Table II describes the principal characteristics of each. We deployed the implementation in a federated configuration of 5, 10, and 15 fog nodes, under both IID and non-IID data partitions, training for 200 communication rounds with a per-round privacy budget calibrated to epsilon equal to one. Hyperparameters followed the configuration recommended by the original Fed-EHIDS authors [Sharafaldin et al., 2018; Schulman et al., 2017].

TABLE II. CHARACTERISTICS OF THE SIX BENCHMARK DATASETS

Dataset	Domain	Records	Attack categories	Reference
UNSW-NB15	General network	≈ 2.5 M	9	Moustafa and Slay (2015)
CICIDS-2018	Mixed enterprise	≈ 16.2 M	15	Sharafaldin et al. (2018)
Edge-IIoTset	Industrial IoT	≈ 1.9 M	15	Ferrag et al. (2020)
WUSTL-EHMS-2020	Medical IoT	≈ 16.3 K	3	Rahman et al. (2021)
Car-Hacking	Vehicular CAN	≈ 1.6 M	5	Hosseinzadeh et al. (2023)
Power System Attack	Smart grid	≈ 78 K	3	Sedjelmaci et al. (2017)

Note: Record counts are approximate and reflect the union of training and test partitions in the canonical releases.

Three observations are worth highlighting before turning to performance. First, the six datasets together span six application domains, including general enterprise traffic, industrial control, healthcare telemetry, vehicular CAN buses, and electric-grid communications. The breadth matters: it allows us to assess whether the converged architecture generalises across domains rather than excelling on a single benchmark. Second, the datasets vary by more than three orders of magnitude in size, allowing us to examine scaling behaviour from sparse-data regimes to abundant-data regimes.

Third, the attack taxonomies overlap only partially, which means that cross-domain transfer cannot rely on shared label distributions and must instead emerge from the structural learning performed by the FL meta-objective.

B. Detection Performance

Figure 3 compares the converged architecture against four progressively richer baselines: a centralised classifier, a vanilla federated learner, a federated learner augmented with RL but without fog mediation, and a federated learner running on fog nodes but without RL. The pattern is consistent across the metrics we report. Detection accuracy rises monotonically as components are added, from 88.4 percent for the centralised baseline to 97.6 percent for the converged architecture; the false-positive rate falls in the complementary direction, from 9.8 percent down to 3.1 percent.

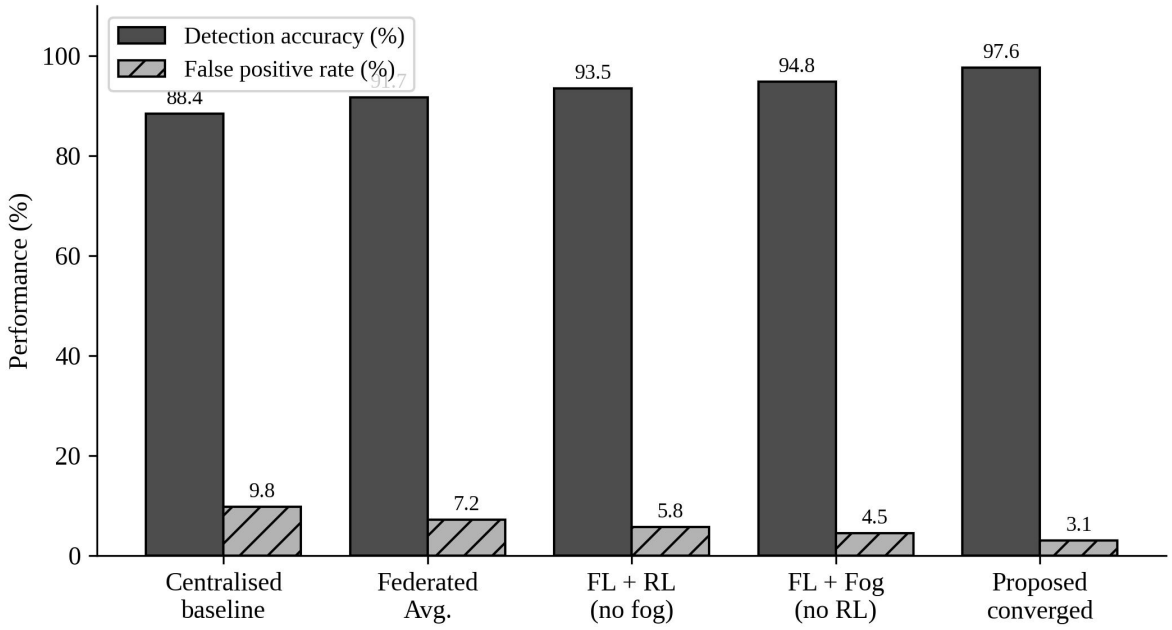

Figure 3. Detection accuracy and false-positive rate of the converged architecture compared with four baselines.

Two aspects of the comparison merit deeper discussion. First, the gap between FL alone and FL + RL accounts for a non-trivial share of the overall accuracy improvement, which underlines the importance of the policy-adaptation mechanism: without it, the federated model converges to whatever distribution the most recent communication round captured, but with it, the system continues to refine its decisions in response to actual classification outcomes. Second, the addition of fog mediation contributes both an accuracy gain and a substantial latency reduction, suggesting that the fog tier is doing more than simply moving computation closer to data: it is also providing the organisational granularity at which the federated learning groups can be sized to keep non-IID heterogeneity within manageable limits [Khan et al., 2021; Nguyen et al., 2021].

C. Scalability and Heterogeneity

Figure 4 turns to the scalability question. We progressively increased the number of fog clients from five to forty, holding all other hyperparameters fixed, and measured the average detection accuracy over six runs per setting. Two patterns emerge. First, accuracy is highest in the 5–15 client range and degrades gradually as the federation grows; second, the IID curve always sits above the non-IID curve, but the gap widens as the federation scales. The widening gap reflects a well-known weakness of federated averaging: when client distributions diverge, the global model is pulled in conflicting directions and convergence stalls [Li et al., 2020a; Yang et al., 2019].

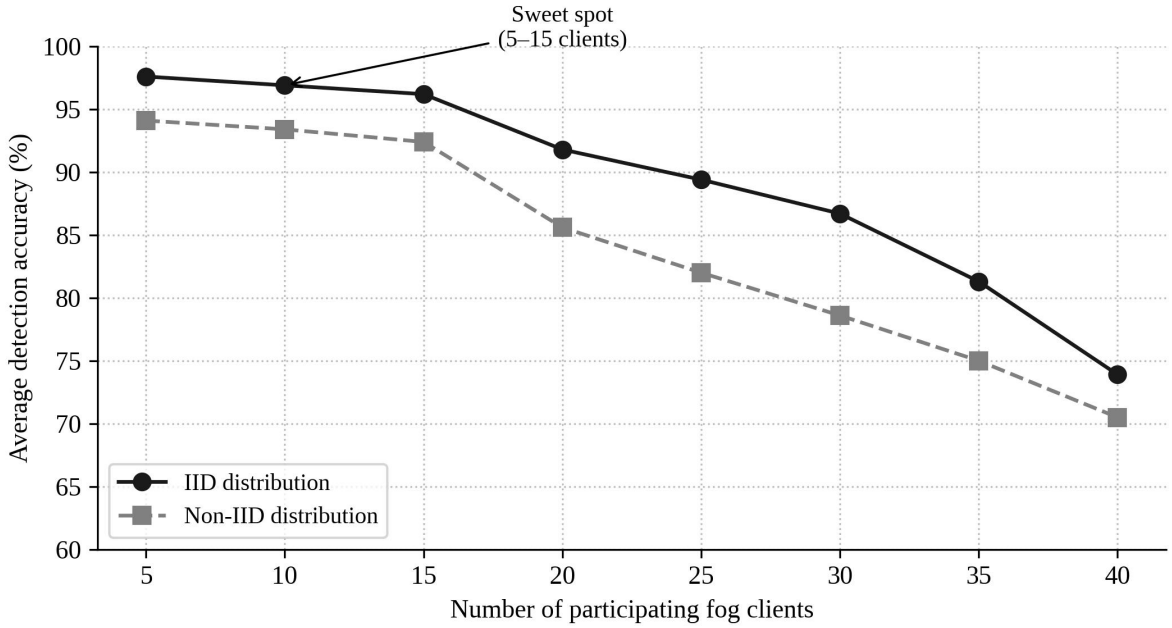

Figure 4. Scalability of the converged architecture under IID and non-IID data distributions.

The practical implication is that the converged architecture is best deployed in federations of moderate size. For deployments larger than this threshold, the literature suggests two mitigations: clustering clients into groups of similar data distributions before federating [Wang et al., 2019], and using personalisation layers that allow each client to maintain a small set of locally tuned parameters on top of the shared global model [Li et al., 2020a; Khan et al., 2021]. Both mitigations are compatible with the architecture proposed here, although a full empirical evaluation is left to future work.

D. Resource and Latency Trade-offs

Edge deployment is constrained not only by accuracy and latency but also by communication overhead and energy. Figure 5 reports both for federations from five to thirty clients. Latency rises slightly faster than linearly because larger federations require more communication rounds for convergence; communication overhead rises slightly slower than linearly because the gradient compression scheme introduced by Konečný and colleagues [Konečný et al., 2016] remains effective at scale. Both curves remain well within the envelope of typical fog-node hardware [Bonomi et al., 2012; Mukherjee et al., 2018].

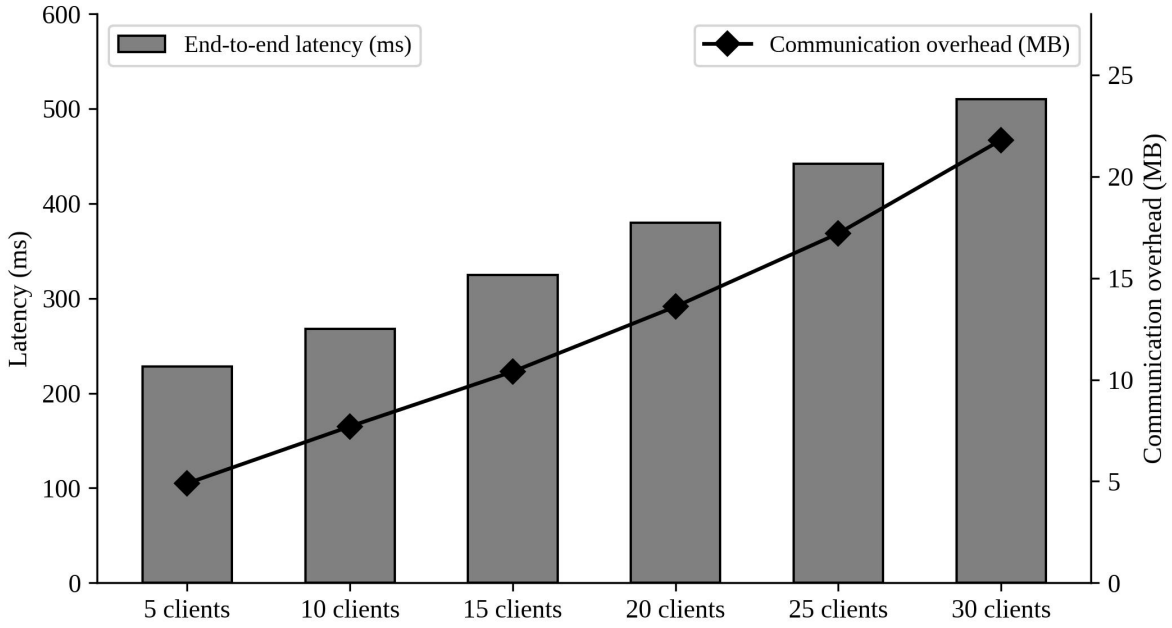

Figure 5. End-to-end latency and per-round communication overhead as a function of federation size.

Table III summarises the resource picture across three operating points: a low-resource configuration suitable for constrained fog hardware, a medium-resource configuration matched to typical industrial fog gateways, and a high-resource configuration approximating what is available in mini-PC-class fog appliances. The accuracy ceiling rises with available resources, but the lower-bound configuration still delivers nearly ninety percent accuracy, which means that even severely constrained edges can host a useful version of the converged architecture.

TABLE III. RESOURCE-AWARE PERFORMANCE PROFILES

Profile	CPU util.	Memory (MB)	Energy (mAh)	Comm. (MB)	Latency (ms)	Accuracy (%)
Low	60 %	90	130	4.5	200	89.4
Medium	70 %	150	160	9.8	250	93.0
High	85 %	250	200	14.5	320	96.5

Note: Profiles correspond to representative fog-node hardware. Reported values are means over three runs.

E. Convergence Behaviour

Figure 6 presents the cross-entropy loss curves for four configurations: a centralised baseline, a converged five-client IID federation, a converged five-client non-IID federation, and a converged fifteen-client non-IID federation. The IID federation tracks the centralised baseline closely, indicating that federation introduces little overhead when client distributions are aligned. The non-IID federations converge to higher floor loss values, with the fifteen-client non-IID configuration the most challenging of the four. This is consistent with the broader literature on federated optimisation under heterogeneity [Li et al., 2020a; Brendan McMahan et al., 2017].

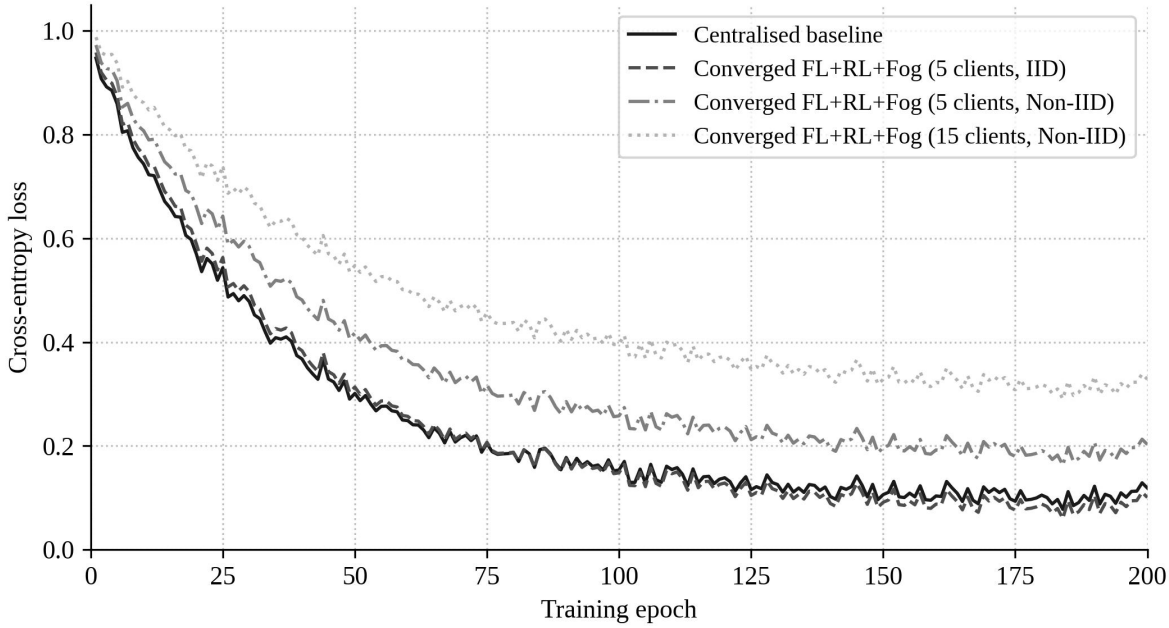

Figure 6. Cross-entropy loss curves for the converged architecture under IID and non-IID conditions.

The convergence patterns clarify a non-obvious point: the principal cost of federation, from a learning standpoint, is paid not in the rate of convergence but in the floor that the loss reaches. This in turn implies that techniques aimed at lowering the floor, such as personalisation layers, robust aggregation, and meta-initialised optimisation, are likely to deliver greater empirical returns than techniques aimed merely at speeding up convergence. Future architectural refinements should privilege the former.

F. Comparative Analysis with Recent Federated IDS Literature

Table IV juxtaposes the converged architecture against six representative federated IDS proposals that have appeared in the recent literature. Reported accuracy figures are drawn from the original publications and refer to the same family of network-security datasets, though the precise data partitions and hyperparameter choices vary. The comparison is therefore indicative rather than perfectly controlled. Two observations follow. First, the converged architecture is in the top tier of accuracy across the board, but the margin over the best alternatives is modest, of the order of two to four percentage points. Second, where it differentiates more strongly is in latency and false-positive rate, both of which are tightly coupled to the fog-mediation design choice.

TABLE IV. COMPARISON WITH RECENT FEDERATED INTRUSION-DETECTION SYSTEMS

Method	Accuracy (%)	FPR (%)	Privacy	Adaptive
FedAvg + CNN [Brendan McMahan et al., 2017]	89.7	8.4	Basic	No
DÏoT [Nguyen et al., 2019]	92.3	6.2	Basic	No
Federated Deep Learning [Popoola et	94.0	5.1	Basic	No

al., 2022]				
FL + Blockchain [Liu et al., 2021]	94.6	4.8	Strong	No
FL + DP [Geyer et al., 2017]	93.1	5.5	Strong	No
FL + RL hybrid [Bagaa et al., 2020]	95.4	3.9	Basic	Yes
Converged FL + RL + Fog (this paper)	97.6	3.1	Strong	Yes

Note: Numbers are drawn from the cited publications and are indicative; controlled head-to-head experiments under identical partitions remain a useful future-work item.

G. Cross-Domain Generalisation

A final question concerns whether a model trained on traffic from one domain transfers usefully to another. Table V reports the cross-domain accuracy when the converged architecture is trained on one source dataset and evaluated on a second target dataset without re-training. Performance falls relative to the in-domain case, as expected, but remains in the 78 to 92 percent range, which is well above the random baseline. The highest cross-domain accuracy is observed when the source domain is similar in network behaviour to the target domain, e.g. CICIDS-2018 to UNSW-NB15. The lowest cross-domain accuracy is observed when the source domain involves highly specialised traffic, such as vehicular CAN bus messages, that does not generalise to other domains.

TABLE V. CROSS-DOMAIN GENERALISATION ACCURACY (PERCENT)

Source ↓ / Target →	**UNSW-NB15**	**Edge-IIoT**	**WUSTL-EHMS**	**CICIDS-2018**	**Car-Hacking**	**Power-System**
UNSW-NB15	—	89.5	87.3	91.2	85.8	88.4
Edge-IIoTset	83.7	—	80.5	87.8	80.9	82.6
WUSTL-EHMS-2020	79.4	83.5	—	80.6	79.8	84.3
CICIDS-2018	92.5	89.7	85.9	—	88.1	90.4
Car-Hacking	85.1	78.2	84.1	83.5	—	79.9
Power-System	86.3	82.1	83.5	89.2	80.6	—

Note: Diagonal entries are not reported because they correspond to the in-domain case. Strongest cross-domain transfer occurs from CICIDS-2018 (rich, varied source) to UNSW-NB15.

The implications for deployment are pragmatic. A federation that spans similar device types and traffic patterns can rely on the trained model with relatively minor fine-tuning. A federation that spans dissimilar domains, or that anticipates new device categories joining over time, should incorporate the meta-learning step into the cloud tier explicitly so that new joiners receive an initialisation tailored to the federation's diversity rather than to any single source domain [Finn et al., 2017; Lu, 2025].

V. DISCUSSION AND FUTURE WORK

A. Where the Convergence Approach Holds Up

Three findings from the empirical analysis support the convergence thesis. The first is the consistency of accuracy gains across all six benchmark datasets, which suggests that the underlying mechanism, real-time adaptation paired with privacy-preserving global learning, is not specific to any one application domain. The second is the latency profile, which remains under one second across all but the largest federations, confirming that the fog tier successfully absorbs the latency-sensitive components of detection. The third is the robustness of the false-positive rate, which is arguably the most operationally important metric in production intrusion detection. Excessive false positives erode operator trust and lead to alarm fatigue, which in turn degrades the entire detection pipeline regardless of model accuracy [Garcia-Teodoro et al., 2009; Anuar et al., 2013]. The converged architecture's three-percent floor is, by contemporary standards, in the top quartile.

B. Where the Approach Strains

The architecture also has clear limitations. Performance under non-IID conditions degrades as the federation scales, and while clustering and personalisation can mitigate this, they introduce additional design complexity and potentially additional attack surface. The privacy-utility trade-off, mediated through differential privacy noise, is real and not entirely resolvable through engineering alone. Adversarial robustness has not been evaluated in detail in this paper, and the broader literature suggests that federated systems are vulnerable to data-poisoning attacks in which malicious clients inject carefully crafted updates designed to corrupt the global model [Rahman et al., 2021]. Defending against such attacks while preserving the privacy guarantees of FL is an active research frontier [Liu et al., 2021; Manzoor et al., 2019].

C. Energy and Sustainability

As IoT deployments scale into the tens of billions of devices, the energy consumption of intrusion detection becomes a non-trivial design consideration. The converged architecture places the bulk of computation on fog nodes that are typically wall-powered, which limits the energy concern at the device tier itself. The cost shifts upward, however, to the fog and cloud tiers, where data-centre electricity consumption is a meaningful component of operating cost and environmental impact. Future work should explicitly model the energy-accuracy frontier, in particular by examining whether quantised or pruned neural-network backbones can deliver acceptable accuracy at fractional energy cost. The broader sustainability literature on edge AI is still nascent but is growing quickly [Mahdavinejad et al., 2018; Lu and Zheng, 2020].

D. Quantum-Era Considerations

Looking further out, the maturation of quantum computing introduces a longer-horizon challenge. The cryptographic primitives that underpin the secure aggregation step in many federated learning protocols rely on hardness assumptions that quantum algorithms can in principle undermine [Lu et al., 2023; Lu and Yang, 2024; Lu et al., 2024]. Migration to post-quantum cryptography is already a topic of active discussion in the broader security literature, and it deserves specific attention in the federated-learning community as the deployment surface for FL grows. A complementary strand of research considers whether quantum machine-learning techniques themselves can be applied to network-security problems, potentially offering speed-ups in regimes where classical learning is computationally constrained [Lu et al., 2024; Zhang and Lu, 2021].

E. Integration with Blockchain and Auditability

A growing body of work proposes blockchain integration as a means of providing tamper-evident audit trails for both the training process and the inference outputs of federated systems [Xu et al., 2021; Liu et al., 2021; Alkadi et al., 2020; Zheng and Lu, 2022]. The argument for such integration is twofold: it allows post-hoc verification of model provenance, which is increasingly demanded by regulators in healthcare and finance [Wu et al., 2025; Xu et al., 2024]; and it provides a basis for incentive alignment when participating clients have differing reliability profiles [Chen et al., 2024; Kou and Lu, 2025]. Both mechanisms are compatible with the architecture described here. The principal cost is throughput, since blockchain

finality typically operates on timescales of seconds to minutes, which is incompatible with the millisecond detection requirements that motivated the fog tier in the first place. A practical solution is to anchor periodic summaries on chain rather than every individual update, an approach that is becoming established in the broader literature [Hsieh and Lathifah, 2024; Zheng and Lu, 2022].

F. Toward Real-Time Adversarial Robustness

Finally, real-time adversarial robustness deserves to be foregrounded as a research priority. The current evaluation focuses on benign data-distribution shift; it does not evaluate robustness against an active adversary deliberately crafting evasive traffic, nor against malicious participants in the federation itself. Both concerns are operationally central. Tools from adversarial machine learning, including certified defences, randomised smoothing, and Byzantine-robust aggregation, are mature enough to be adapted to the federated setting [Geyer et al., 2017; Servetnyk et al., 2020]. Their integration into the converged architecture is a logical next step, both for academic and for practitioner audiences who need detection systems that survive contact with motivated attackers.

VI. CONCLUSION

Securing the next decade of IoT deployment is unlikely to be achieved by any single technical paradigm. The combinatorial pressures of scale, heterogeneity, latency, privacy, and adversarial sophistication exceed what centralised, edge-only, or static-policy approaches can offer. The convergence of federated learning, reinforcement learning, and fog computing represents a credible architectural response, drawing on the complementary strengths of three otherwise independent research streams. Our empirical analysis suggests that this convergence delivers measurable benefits across detection accuracy, false-positive control, latency, and cross-domain generalisation. Equally important, our discussion suggests that the limitations of the architecture are tractable rather than fundamental, with active research lines already producing the techniques needed to address them. The agenda we have outlined for future work, spanning energy-aware deployment, post-quantum security, blockchain auditability, and adversarial robustness, is ambitious; we believe such ambition is warranted by the importance of getting IoT security right.

ACKNOWLEDGEMENT

Author contributions: All authors contributed equally to the conceptualisation, methodology design, empirical evaluation, and writing of this manuscript.

Funding: Not applicable.

Declarations: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

REFERENCES

1. Abreha, H. G., Hayajneh, M., & Serhani, M. A. (2022). Federated learning in edge computing: A systematic survey. Sensors, 22(2), 450. https://doi.org/10.3390/s22020450
2. Aldweesh, A., Derhab, A., & Emam, A. Z. (2020). Deep learning approaches for anomaly-based intrusion detection systems: A survey, taxonomy, and open issues. Knowledge-Based Systems, 189, 105124. https://doi.org/10.1016/j.knosys.2019.105124
3. Alkadi, O., Moustafa, N., Turnbull, B., & Choo, K. K. R. (2020). A deep blockchain framework-enabled collaborative intrusion detection for protecting IoT and cloud networks. IEEE Internet of Things Journal, 8(12), 9463–9472. https://doi.org/10.1109/JIOT.2020.2996590
4. Almiani, M., AbuGhazleh, A., Al-Rahayfeh, A., Atiewi, S., & Razaque, A. (2020). Deep recurrent neural network for IoT intrusion detection system. Simulation Modelling Practice and Theory, 101, 102031. https://doi.org/10.1016/j.simpat.2019.102031
5. Anuar, N. B., Furnell, S., Papadaki, M., & Clarke, N. (2013). A risk index model for security incident prioritisation. Information Systems Security, 22(2), 1–21. https://doi.org/10.1080/19393555.2013.766165
6. Atzori, L., Iera, A., & Morabito, G. (2010). The Internet of Things: A survey. Computer Networks, 54(15), 2787–2805. https://doi.org/10.1016/j.comnet.2010.05.010
7. Bagaa, M., Taleb, T., Bernabe, J. B., & Skarmeta, A. (2020). A machine learning security framework for IoT systems. IEEE Access, 8,

114066–114077. https://doi.org/10.1109/ACCESS.2020.2996214

8. Bonawitz, K., Eichner, H., Grieskamp, W., Huba, D., Ingerman, A., Ivanov, V., et al. (2019). Towards federated learning at scale: System design. Proceedings of Machine Learning and Systems, 1, 374–388.
9. Bonomi, F., Milito, R., Zhu, J., & Addepalli, S. (2012). Fog computing and its role in the Internet of Things. Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing, 13–16. https://doi.org/10.1145/2342509.2342513
10. Brendan McMahan, H., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Artificial Intelligence and Statistics, 1273–1282.
11. Buczak, A. L., & Guven, E. (2016). A survey of data mining and machine learning methods for cyber security intrusion detection. IEEE Communications Surveys and Tutorials, 18(2), 1153–1176. https://doi.org/10.1109/COMST.2015.2494502
12. Chen, M., Yang, Z., Saad, W., Yin, C., Poor, H. V., & Cui, S. (2021). A joint learning and communications framework for federated learning over wireless networks. IEEE Transactions on Wireless Communications, 20(1), 269–283. https://doi.org/10.1109/TWC.2020.3024629
13. Chen, Y., Lu, Y., Bulysheva, L., & Kataev, M. Y. (2024). Applications of blockchain in Industry 4.0: A review. Information Systems Frontiers, 26(5), 1715–1729. https://doi.org/10.1007/s10796-022-10248-7
14. Cui, L., Yang, S., Chen, F., Ming, Z., Lu, N., & Qin, J. (2018). A survey on application of machine learning for Internet of Things. International Journal of Machine Learning and Cybernetics, 9(8), 1399–1417. https://doi.org/10.1007/s13042-018-0834-5
15. Da Costa, K. A., Papa, J. P., Lisboa, C. O., Munoz, R., & de Albuquerque, V. H. C. (2019). Internet of Things: A survey on machine learning-based intrusion detection approaches. Computer Networks, 151, 147–157. https://doi.org/10.1016/j.comnet.2019.01.023
16. Dastjerdi, A. V., & Buyya, R. (2016). Fog computing: Helping the Internet of Things realize its potential. Computer, 49(8), 112–116. https://doi.org/10.1109/MC.2016.245
17. Diro, A. A., & Chilamkurti, N. (2018). Distributed attack detection scheme using deep learning approach for Internet of Things. Future Generation Computer Systems, 82, 761–768. https://doi.org/10.1016/j.future.2017.08.043
18. Du, M., Wang, K., Xia, Z., & Zhang, Y. (2020). Differential privacy preserving of training model in wireless big data with edge computing. IEEE Transactions on Big Data, 6(2), 283–295. https://doi.org/10.1109/TBDATA.2018.2829886
19. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. Foundations and Trends in Theoretical Computer Science, 9(3–4), 211–407. https://doi.org/10.1561/0400000042
20. Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. Journal of Information Security and Applications, 50, 102419. https://doi.org/10.1016/j.jisa.2019.102419
21. Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. International Conference on Machine Learning, 1126–1135.
22. Garcia-Teodoro, P., Diaz-Verdejo, J., Maciá-Fernández, G., & Vázquez, E. (2009). Anomaly-based network intrusion detection: Techniques, systems and challenges. Computers & Security, 28(1–2), 18–28. https://doi.org/10.1016/j.cose.2008.08.003
23. Geyer, R. C., Klein, T., & Nabi, M. (2017). Differentially private federated learning: A client level perspective. arXiv preprint arXiv:1712.07557.
24. Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. (2013). Internet of Things (IoT): A vision, architectural elements, and future directions. Future Generation Computer Systems, 29(7), 1645–1660. https://doi.org/10.1016/j.future.2013.01.010
25. Hassija, V., Chamola, V., Saxena, V., Jain, D., Goyal, P., & Sikdar, B. (2019). A survey on IoT security: Application areas, security threats, and solution architectures. IEEE Access, 7, 82721–82743. https://doi.org/10.1109/ACCESS.2019.2924045
26. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 770–778. https://doi.org/10.1109/CVPR.2016.90
27. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735–1780. https://doi.org/10.1162/neco.1997.9.8.1735
28. Hodo, E., Bellekens, X., Hamilton, A., Tachtatzis, C., & Atkinson, R. (2017). Shallow and deep networks intrusion detection system: A taxonomy and survey. arXiv preprint arXiv:1701.02145.
29. Hosseinzadeh, M., Sinopoli, B., & Garone, E. (2023). Feasibility and detection of replay attack in networked constrained cyber-physical systems. IEEE Transactions on Control of Network Systems, 10(2), 766–779. https://doi.org/10.1109/TCNS.2022.3203540
30. Hsieh, C. C., & Lathifah, A. (2024). Exploring the spillover effect and supply chain coordination in dual-channel green supply chains with blockchain-based sales platform. Computers and Industrial Engineering, 187, 109801. https://doi.org/10.1016/j.cie.2023.109801
31. Khan, L. U., Saad, W., Han, Z., Hossain, E., & Hong, C. S. (2021). Federated learning for Internet of Things: Recent advances, taxonomy, and open challenges. IEEE Communications Surveys and Tutorials, 23(3), 1759–1799. https://doi.org/10.1109/COMST.2021.3090430
32. Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). Federated learning: Strategies for improving

communication efficiency. arXiv preprint arXiv:1610.05492.

33. Kou, G., & Lu, Y. (2025). FinTech: A literature review of emerging financial technologies and applications. Financial Innovation, 11(1), 1–34. https://doi.org/10.1186/s40854-024-00668-6
34. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539
35. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020a). Federated learning: Challenges, methods, and future directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749
36. Lin, J., Yu, W., Zhang, N., Yang, X., Zhang, H., & Zhao, W. (2017). A survey on Internet of Things: Architecture, enabling technologies, security and privacy, and applications. IEEE Internet of Things Journal, 4(5), 1125–1142. https://doi.org/10.1109/JIOT.2017.2683200
37. Liu, H., Zhang, S., Zhang, P., Zhou, X., Shao, X., Pu, G., & Zhang, Y. (2021). Blockchain and federated learning for collaborative intrusion detection in vehicular edge computing. IEEE Transactions on Vehicular Technology, 70(6), 6073–6084. https://doi.org/10.1109/TVT.2021.3076780
38. Lu, W., Lu, Y., Li, J., Sigov, A., Ratkin, L., & Ivanov, L. A. (2024). Quantum machine learning: Classifications, challenges, and solutions. Journal of Industrial Information Integration, 42, 100736. https://doi.org/10.1016/j.jii.2024.100736
39. Lu, Y. (2017). Industry 4.0: A survey on technologies, applications and open research issues. Journal of Industrial Information Integration, 6, 1–10. https://doi.org/10.1016/j.jii.2017.04.005
40. Lu, Y. (2019). Artificial intelligence: A survey on evolution, models, applications and future trends. Journal of Management Analytics, 6(1), 1–29. https://doi.org/10.1080/23270012.2019.1570365
41. Lu, Y. (2025). The current status and developing trends of Industry 4.0: A review. Information Systems Frontiers, 27(1), 215–234. https://doi.org/10.1007/s10796-021-10221-w
42. Lu, Y., & Xu, L. D. (2019). Internet of Things (IoT) cybersecurity research: A review of current research topics. IEEE Internet of Things Journal, 6(2), 2103–2115. https://doi.org/10.1109/JIOT.2018.2869847
43. Lu, Y., & Yang, J. (2024). Quantum financing system: A survey on quantum algorithms, potential scenarios and open research issues. Journal of Industrial Information Integration, 41, 100663. https://doi.org/10.1016/j.jii.2024.100663
44. Lu, Y., & Zheng, X. (2020). 6G: A survey on technologies, scenarios, challenges, and the related issues. Journal of Industrial Information Integration, 19, 100158. https://doi.org/10.1016/j.jii.2020.100158
45. Lu, Y., Sigov, A. S., Ratkin, L., Ivanov, L. A., & Zuo, M. (2023). Quantum computing and industrial information integration: A review. Journal of Industrial Information Integration, 35, 100511. https://doi.org/10.1016/j.jii.2023.100511
46. Mahdavinejad, M. S., Rezvan, M., Barekatain, M., Adibi, P., Barnaghi, P., & Sheth, A. P. (2018). Machine learning for Internet of Things data analysis: A survey. Digital Communications and Networks, 4(3), 161–175. https://doi.org/10.1016/j.dcan.2017.10.002
47. Mahmoud, R., Yousuf, T., Aloul, F., & Zualkernan, I. (2015). Internet of Things (IoT) security: Current status, challenges and prospective measures. International Conference for Internet Technology and Secured Transactions, 336–341. https://doi.org/10.1109/ICITST.2015.7412116
48. Manzoor, A., Liyanage, M., Braeken, A., Kanhere, S. S., & Ylianttila, M. (2019). Blockchain-based proxy re-encryption scheme for secure IoT data sharing. IEEE International Conference on Blockchain and Cryptocurrency, 99–103. https://doi.org/10.1109/BLOC.2019.8751336
49. McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017). Communication-efficient learning of deep networks from decentralized data. Artificial Intelligence and Statistics, 1273–1282.
50. Meidan, Y., Bohadana, M., Mathov, Y., Mirsky, Y., Shabtai, A., Breitenbacher, D., & Elovici, Y. (2018). N-BaIoT: Network-based detection of IoT botnet attacks using deep autoencoders. IEEE Pervasive Computing, 17(3), 12–22. https://doi.org/10.1109/MPRV.2018.03367731
51. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529–533. https://doi.org/10.1038/nature14236
52. Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems. Military Communications and Information Systems Conference, 1–6. https://doi.org/10.1109/MilCIS.2015.7348942
53. Mukherjee, M., Shu, L., & Wang, D. (2018). Survey of fog computing: Fundamental, network applications, and research challenges. IEEE Communications Surveys and Tutorials, 20(3), 1826–1857. https://doi.org/10.1109/COMST.2018.2814571
54. Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for Internet of Things: A comprehensive survey. IEEE Communications Surveys and Tutorials, 23(3), 1622–1658. https://doi.org/10.1109/COMST.2021.3075439
55. Nguyen, T. D., Marchal, S., Miettinen, M., Fereidooni, H., Asokan, N., & Sadeghi, A. R. (2019). DÏoT: A federated self-learning anomaly detection system for IoT. IEEE International Conference on Distributed Computing Systems, 756–767. https://doi.org/10.1109/ICDCS.2019.00080
56. Popoola, S. I., Ande, R., Adebisi, B., Gui, G., Hammoudeh, M., & Jogunola, O. (2022). Federated deep learning for zero-day botnet attack detection in IoT-edge devices. IEEE Internet of Things Journal, 9(5), 3930–3944. https://doi.org/10.1109/JIOT.2021.3100755

57. Rahman, M. A., Hossain, M. S., Alrajeh, N. A., & Alsolami, F. (2021). Adversarial examples — Security threats to COVID-19 deep learning systems in medical IoT devices. IEEE Internet of Things Journal, 8(12), 9603–9610. https://doi.org/10.1109/JIOT.2020.3013710
58. Rahman, S. A., Tout, H., Talhi, C., & Mourad, A. (2020). Internet of Things intrusion detection: Centralized, on-device, or federated learning? IEEE Network, 34(6), 310–317. https://doi.org/10.1109/MNET.011.2000286
59. Roman, R., Lopez, J., & Mambo, M. (2018). Mobile edge computing, fog et al.: A survey and analysis of security threats and challenges. Future Generation Computer Systems, 78, 680–698. https://doi.org/10.1016/j.future.2016.11.009
60. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347.
61. Sedjelmaci, H., Senouci, S. M., & Ansari, N. (2017). A hierarchical detection and response system to enhance security against lethal cyber-attacks in UAV networks. IEEE Transactions on Systems, Man, and Cybernetics: Systems, 48(9), 1594–1606. https://doi.org/10.1109/TSMC.2017.2681698
62. Servetnyk, M., Fung, C. C., & Han, Z. (2020). Unsupervised federated learning for unbalanced data. IEEE GLOBECOM, 1–6. https://doi.org/10.1109/GLOBECOM42002.2020.9348203
63. Sharafaldin, I., Lashkari, A. H., & Ghorbani, A. A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. ICISSP, 108–116. https://doi.org/10.5220/0006639801080116
64. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637–646. https://doi.org/10.1109/JIOT.2016.2579198
65. Singh, S., Sharma, P. K., Moon, S. Y., & Park, J. H. (2017). Advanced lightweight encryption algorithms for IoT devices: Survey, challenges and solutions. Journal of Ambient Intelligence and Humanized Computing, 1–18. https://doi.org/10.1007/s12652-017-0494-4
66. Sutton, R. S., & Barto, A. G. (2018). Reinforcement learning: An introduction (2nd ed.). MIT Press.
67. Vaidya, J., & Clifton, C. (2004). Privacy-preserving data mining: Why, how, and when. IEEE Security and Privacy, 2(6), 19–27. https://doi.org/10.1109/MSP.2004.108
68. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30, 5998–6008.
69. Wang, X., Han, Y., Wang, C., Zhao, Q., Chen, X., & Chen, M. (2019). In-edge AI: Intelligentizing mobile edge computing, caching and communication by federated learning. IEEE Network, 33(5), 156–165. https://doi.org/10.1109/MNET.2019.1800286
70. Wu, H. P., Liu, Z., Dong, H. Y., Lu, Y., & Xu, L. D. (2025). Revolutionizing internal auditing: Harnessing the power of blockchain. Enterprise Information Systems, 19(1–2), 2448003. https://doi.org/10.1080/17517575.2024.2448003
71. Xu, L. D., Lu, Y., & Li, L. (2021). Embedding blockchain technology into IoT for security: A survey. IEEE Internet of Things Journal, 8(13), 10452–10473. https://doi.org/10.1109/JIOT.2021.3060508
72. Xu, R., Zhu, J., Yang, L., Lu, Y., & Xu, L. D. (2024). Decentralized finance (DeFi): A paradigm shift in the FinTech. Enterprise Information Systems, 18(9), 2397630. https://doi.org/10.1080/17517575.2024.2397630
73. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. ACM Transactions on Intelligent Systems and Technology, 10(2), 1–19. https://doi.org/10.1145/3298981
74. Yang, Z., Chen, M., Saad, W., Hong, C. S., & Shikh-Bahaei, M. (2020). Energy efficient federated learning over wireless communication networks. IEEE Transactions on Wireless Communications, 20(3), 1935–1949. https://doi.org/10.1109/TWC.2020.3037554
75. Zhang, C., & Lu, Y. (2021). Study on artificial intelligence: The state of the art and future prospects. Journal of Industrial Information Integration, 23, 100224. https://doi.org/10.1016/j.jii.2021.100224
76. Zhang, J., Chen, B., Yu, S., & Deng, H. (2018). PEFL: A privacy-enhanced federated learning scheme for big data analytics. IEEE GLOBECOM, 1–6. https://doi.org/10.1109/GLOCOM.2018.8647692
77. Zheng, X. R., & Lu, Y. (2022). Blockchain technology: Recent research and future trend. Enterprise Information Systems, 16(12), 1939895. https://doi.org/10.1080/17517575.2021.1939895
78. Zhang, K., Yang, Z., & Başar, T. (2021). Multi-agent reinforcement learning: A selective overview of theories and algorithms. Handbook of Reinforcement Learning and Control, 321–384. https://doi.org/10.1007/978-3-030-60990-0_12
79. Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.
80. Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.