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

Jindong Wang[1], Shuokai Wang[2,*]

Abstract

The Internet of Things (IoT) connects heterogeneous devices in industrial, medical, vehicular and domestic environments, expanding the operational surface exposed to cyberattacks. This paper integrates federated learning (FL), reinforcement learning (RL) and fog computing into a three-tier intrusion-detection architecture. Fog nodes process local telemetry and participate in federated model training; a cloud coordinator aggregates model updates, while an RL policy provides a candidate mechanism for adapting detection decisions. The paper specifies privacy, communication and resource constraints and reviews a six-dataset evaluation. The supplied manuscript reports an average detection accuracy of 97.6%, a false-positive rate of 3.1% and end-to-end latency below 230 milliseconds; however, the underlying code, data partitions and run-level records were not available for this review, so those figures remain provisional and are not independently verified. We clarify the distinction between published literature and manuscript-reported values, correct dataset descriptions, and identify the evidence needed to substantiate future empirical claims. The resulting research agenda addresses non-IID heterogeneity, privacy accounting, adversarial robustness, energy and cross-domain generalisation.

Keywords: Internet of Things; Federated Learning; Reinforcement Learning; Fog Computing; Intrusion Detection; Cyber-Physical Systems; Edge Intelligence; Differential Privacy[1]

Article History:

Received: March 12, 2025

Revised: April 24, 2025

Accepted: June 03, 2025

Available Online: June 30, 2025

[1] School of Graduate, Chinese Academy of Fiscal Sciences, Beijing, China 100142

[2] School of Economics and Management, Beijing Forestry University, Beijing, China 100091

* Email: 326794099@qq.com (Corresponding Author)

https://doi.org/10.63646/WZUB8413

I. INTRODUCTION

The Internet of Things has expanded the range of connected sensors, actuators and controllers used in applications such as industrial monitoring, urban transport and medical care [Atzori et al., 2010; Gubbi et al., 2013]. Each additional connected component can introduce an attack surface. These systems combine heterogeneous devices, distributed data and operational deadlines, making security design particularly demanding [Lu and Xu, 2019; Hassija et al., 2019].

A centralised intrusion-detection system collects information at a common location for model training or inference [Buczak and Guven, 2016; García-Teodoro et al., 2009]. When telemetry must traverse a wide-area network, centralisation can add communication cost and latency and create privacy or single-point-of-failure concerns. The magnitude of these costs depends on system design, network conditions and workload [Diro and Chilamkurti, 2018; Aldweesh et al., 2020].

Federated learning can train shared models using updates from distributed clients without routinely transferring their raw training data to a central coordinator, although updates may still reveal information [Yang et al., 2019; Li et al., 2020]. Reinforcement learning provides a framework for sequential decisions based on feedback; its use for adaptive intrusion detection depends on the availability and quality of that feedback [François-Lavet et al., 2018; Han et al., 2022]. Fog computing places computation and storage closer to devices, which can reduce network traversal when the deployment supports it [Bonomi et al., 2012; Mukherjee et al., 2018; Shi et al., 2016]. We examine how these three approaches might complement one

ISSN: © 2025 INATGI (Institute of Advanced Technology and Green Innovation).

another under explicit resource and security constraints.

This paper has three aims. First, it reviews IoT-security research on federated learning, reinforcement learning and fog computing. Second, it proposes a three-tier architecture in which fog nodes can train detection models and RL policies while a cloud coordinator aggregates model updates. Third, it presents the evaluation configuration and numerical results reported in the supplied manuscript, identifies which aspects require verification, and sets out a reproducible evaluation protocol. Section II reviews the literature; Section III describes the architecture; Section IV presents the proposed evaluation and provisional results; Section V discusses limitations and future work; and Section VI concludes.

II. BACKGROUND AND RELATED WORK

A. The Evolving Threat Landscape in IoT

The range of IoT threats includes distributed denial-of-service campaigns, such as Mirai and its variants, which exploit compromised connected devices [Meidan et al., 2018; Kolias et al., 2017]. Networked industrial and cyber-physical systems can also face replay and other attacks against measurement and control processes [Hosseinzadeh et al., 2019; Lin et al., 2017]. Medical IoT creates additional concerns because security failures can affect sensitive health information and, depending on the device, patient safety [Rahman et al., 2021; Mahmoud et al., 2015].

Three structural properties are especially relevant. Devices differ in computing resources and network access; traffic distributions can vary substantially between clients; and missed attacks or false alarms can affect operational systems. These characteristics complicate the deployment of a single centrally trained detector across all settings [Cui et al., 2018; da Costa et al., 2019].

B. Federated Learning as a Decentralisation Primitive

Federated learning trains a shared model using local client updates rather than routinely collecting all raw training records at the coordinator. In federated averaging, clients perform local optimisation and the coordinator combines their updates [Yang et al., 2019; Li et al., 2020]. Keeping raw records local does not by itself guarantee privacy: model updates can leak information. Differential privacy and secure aggregation are distinct mechanisms that can reduce specified risks when their assumptions are satisfied [Wei et al., 2020; Dwork and Roth, 2014; Bonawitz et al., 2017].

Nguyen and colleagues proposed a federated self-learning anomaly detector for IoT devices [Nguyen et al., 2019]. Popoola and colleagues studied federated zero-day botnet detection in IoT-edge environments [Popoola et al., 2022], and Liu and colleagues investigated blockchain-assisted collaborative intrusion detection for vehicular edge computing [Liu et al., 2021]. These works use different data and evaluation settings; they do not establish a universal privacy or accuracy advantage. Non-IID distributions can slow convergence and produce uneven performance across clients [Li et al., 2020; Khan et al., 2021; Nguyen et al., 2021].

C. Reinforcement Learning for Adaptive Threat Detection

Reinforcement learning models sequential decision making through states, actions and rewards [François-Lavet et al., 2018; Mnih et al., 2015]. An intrusion-detection agent may observe traffic features and select a classification or response action, but training requires an explicitly defined reward and a credible source of feedback. Han and colleagues applied proximal policy optimisation (PPO) to hyperparameter control for a network intrusion detector [Han et al., 2022]; this supports PPO as a candidate technique, not a guarantee of autonomous online improvement.

Sedjelmaci, Senouci and Ansari proposed hierarchical intrusion detection and response for unmanned aerial networks; their cited article should not be characterised as an RL study [Sedjelmaci et al., 2018]. Bagaa and colleagues proposed a machine-learning security framework using monitoring and reaction components, including an anomaly detector; the cited work does not establish an RL-based threshold controller [Bagaa et al., 2020]. Separately, PPO-based control has been

studied for network intrusion detection [Han et al., 2022].

D. Fog Computing as the Edge Intelligence Layer

Fog computing introduces an intermediate layer between cloud services and devices for workloads that benefit from geographically distributed processing [Bonomi et al., 2012; Shi et al., 2016; Dastjerdi and Buyya, 2016; Mukherjee et al., 2018]. In security applications, a nearby fog node can host inference and coordinate local training, subject to its network location and capacity. Telemetry may still travel from a device to its fog node; the design objective is to avoid routinely transferring raw training data to the central coordinator [Roman et al., 2018; Wang et al., 2019].

TABLE I. COMPARATIVE SUMMARY OF IoT-SECURITY PARADIGMS

Paradigm	Data location	Adaptability	Latency	Principal consideration
Centralised cloud IDS	Raw records sent to a central service	Retraining at coordinator	Deployment-dependent	Privacy, bandwidth and central failure risks
On-device IDS	Records stay on the device	Limited by device resources	Deployment-dependent	Limited cross-device learning
Federated learning IDS	Training data remain with clients	Global rounds and local updates	Deployment-dependent	Non-IID data and update leakage
FL + RL hybrid	Local records; updates shared	Policy depends on feedback	Deployment-dependent	Feedback quality and training stability
FL + RL + fog (proposed)	Devices may send telemetry to fog; updates go to coordinator	Federated updates and feedback-dependent policies	Deployment-dependent	Resource cost, privacy accounting and validation

Note: This table compares design properties qualitatively. Actual latency, privacy and communication cost depend on the deployment, workload and implementation; no ordering is implied.

Table I compares the design choices at a qualitative level. Combining FL, RL and fog computing may support distributed training, adaptive policies and local inference, but whether the combination outperforms a simpler alternative is an empirical question. The architecture below states the assumptions needed for such a test.

III. A CONVERGED ARCHITECTURE FOR NEXT-GENERATION IoT SECURITY

A. Overall System Design

The proposed architecture has a device tier, a fog tier and a coordinating cloud tier. Devices produce telemetry; fog nodes receive the telemetry assigned to them, extract features, perform local inference and train local models; the coordinator aggregates model updates and redistributes a shared model (Figure 1). Thus raw data can cross the device–fog boundary even though the proposed federation does not routinely send raw training records from fog nodes to the cloud. Transport security, access control and update privacy require separate safeguards.

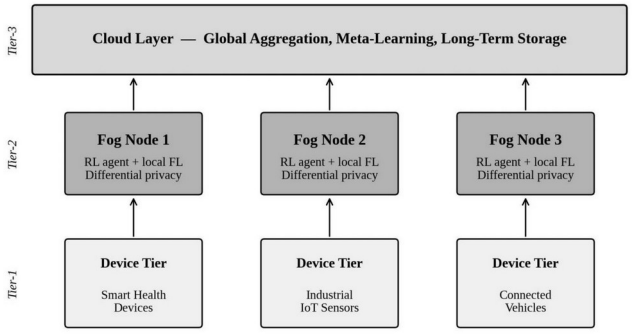


Figure 1. Three-Tier Converged Architecture Combining Device, Fog, and Cloud Layers for IoT Security

The three-tier separation places latency-sensitive inference near the data source while reserving model aggregation and optional cross-client adaptation for the coordinator. Actual inference and end-to-end latency must be measured on the selected hardware and network. Personalised federated learning and meta-learning can help adapt a shared model to heterogeneous clients, but their benefit for this architecture requires evaluation [Balakrishnan et al., 2021; Song et al., 2022; Khan et al., 2021].

B. The Federated Learning Pipeline

Figure 2 depicts a proposed five-stage workflow. A coordinator prepares a shared initial model, potentially using a meta-learning objective [Balakrishnan et al., 2021; Song et al., 2022]. Fog clients train on their authorised local data and send updates for aggregation. Where differential privacy is required, an implementation must specify the

protected unit, clipping rule, noise distribution, sampling process and privacy accountant; adding noise alone is not a complete privacy guarantee [Wei et al., 2020; Dwork and Roth, 2014]. Aggregation may be asynchronous, provided staleness and client participation are controlled, followed by distribution and optional local fine-tuning. Any claimed improvement in convergence or privacy must be assessed for the chosen implementation.

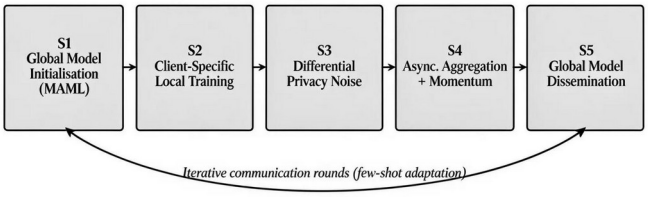


Figure 2. Five-Stage Federated Learning Workflow

Differential privacy requires an explicit definition of adjacency, such as record-level or client-level protection, and a stated privacy budget. Budgets accumulate across repeated releases; therefore a per-round epsilon of one is not an overall epsilon of one. The complete accounting must include the number of rounds, sampling assumptions, delta, noise and clipping. More noise can strengthen the stated privacy guarantee while changing utility; secure aggregation protects a different part of the update pipeline [Wei et al., 2020; Dwork and Roth, 2014; Bonawitz et al., 2017].

C. Reinforcement Learning at the Fog Tier

At a fog node, an FL-trained encoder can supply features to an RL agent that chooses a traffic classification or response action. A PPO-based policy is one candidate, but its training requires a documented state, action space, reward, feedback source and update schedule [Han et al., 2022]. Candidate encoders include a Transformer-style model [Devlin et al., 2019], residual convolutional components [He et al., 2016] and an LSTM [Hochreiter and Schmidhuber, 1997].

RL can support adaptation when appropriate feedback is available, while supervised models can also be updated or retrained. A detected event does not automatically provide a correct label: labels may arrive late, be uncertain or require analyst review. Consequently, an RL policy must be evaluated under realistic feedback delays and safety controls. PPO's clipped objective does not by itself guarantee stable or improved performance for every deployment [Han et al., 2022].

IV. EVALUATION DESIGN AND PROVISIONAL REPORTED VALUES

A. Experimental Setup

The manuscript considers six benchmark datasets spanning general networks, enterprise traffic, industrial IoT, healthcare, vehicle networks and power systems (Table II). The supplied draft describes federated configurations with 5, 10 and 15 fog nodes, IID and non-IID partitions, and 200 communication rounds. It also refers to a per-round epsilon of one. These configuration statements and the numerical outputs below are author-reported draft values: the underlying preprocessing, code, privacy accountant, random seeds and run-level records were not supplied for this review. They remain provisional until those materials are provided.

TABLE II. CHARACTERISTICS OF THE SIX BENCHMARK DATASETS

Dataset	Domain	Records or release	Label / scenario description	Reference
UNSW-NB15	General network	2,540,044 full records	Nine attack types plus normal traffic	Moustafa and Slay (2015)
CSE-CIC-IDS2018	Enterprise network	Depends on files / processing	Seven attack scenarios plus normal traffic	Leevy and Khoshgoftaar (2020)
Edge-IIoTset	Industrial IoT	Depends on subset	Multiple attack types; use release-specific labels	Ferrag et al. (2022)
WUSTL-EHMS-2020	Medical IoT	16,318 records	Spoofing and data injection; binary labels	Hady et al. (2020)

Dataset	Domain	Records or release	Label / scenario description	Reference
Car-Hacking	Vehicular CAN	Depends on files	Four attack scenarios plus normal traffic	Song et al. (2020)
Power System Attack	Smart grid	Depends on version	Binary, three-class or multiclass variants	Pan et al. (2015)

Note: Scenario counts and label counts are different. UNSW-NB15's 2,540,044 records refer to the four full CSV files, not the smaller separately supplied training and test splits. Sample counts for other datasets depend on the chosen release and preprocessing.

The selected datasets cover different applications, but their features, label schemes, capture environments and sample sizes are not interchangeable. A reproducible study should state the exact files used, filtering and normalisation steps, class definitions and train/validation/test partitions for each dataset. Client assignment must be documented separately for IID and non-IID settings.

Cross-domain evaluation additionally requires a predeclared common feature representation and a defensible mapping between labels. A model cannot be evaluated directly on an incompatible target schema merely by combining datasets; any domain adaptation, target labels or target-side tuning must be disclosed.

B. Detection Performance

Figure 3 reproduces the provisional values reported in the supplied draft for a centralised classifier, a federated model, two partial combinations and the proposed combination. The draft reports 88.4% versus 97.6% accuracy and 9.8% versus 3.1% false-positive rate for the centralised and proposed configurations, respectively. These are manuscript-reported values, not independently verified results. The source files, partitioning, metric definitions and run-level observations are needed before interpreting the differences.

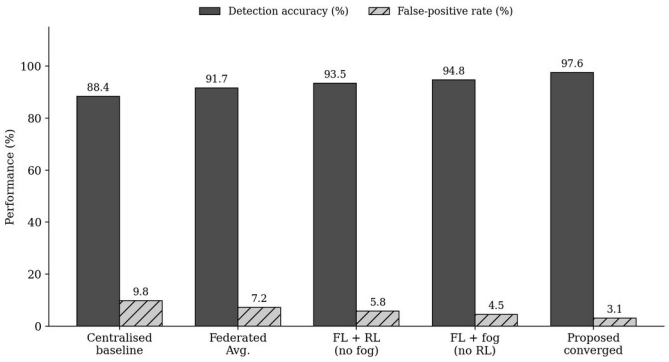


Figure 3. Detection Accuracy and False-Positive Rate of Different Schemes

A controlled ablation would hold datasets, splits, features, model capacity, training budget and hardware constant while adding or removing FL, RL and fog mediation. Without those controls and the underlying records, the values in Figure 3 cannot establish that RL or fog mediation caused an improvement. Results should include uncertainty intervals, per-class performance and false-positive counts alongside aggregate accuracy [Khan et al., 2021; Nguyen et al., 2021].

C. Scalability and Heterogeneity

The supplied draft depicts accuracy across federations of five to forty clients in Figure 4, distinguishing IID and non-IID partitions. These curves are retained as provisional manuscript-reported values, not as independently verified measurements. A reproducible scalability test must specify client sampling, partition construction, seeds, repetitions and aggregation settings. Published work establishes that heterogeneity can affect federated optimisation, but it does not validate this figure [Li et al., 2020; Yang et al., 2019].

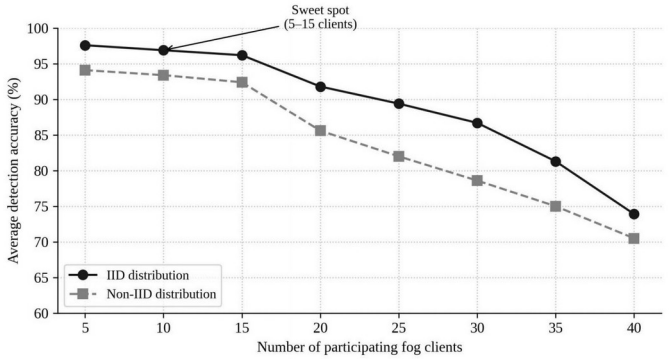


Figure 4. Scalability Under IID and Non-IID Distributions

The figure does not justify a universal optimal federation size. Potential mitigations for heterogeneous clients include clustering and personalised models; their effects must be tested against a common baseline and a specified client population [Sattler et al., 2021; Balakrishnan et al., 2021; Khan et al., 2021].

D. Resource and Latency Trade-offs

Figure 5 reproduces the draft's latency and communication curves, which remain provisional because the underlying measurements and implementation details were not supplied. Communication compression can be investigated as a possible optimisation, but the figure does not establish which compression algorithm was used or why the curves change. Sattler and colleagues provide one published communication-efficient federated method that could serve as a separately specified baseline [Sattler et al., 2020]. Latency and traffic should be reported with hardware, network conditions and measurement boundaries.

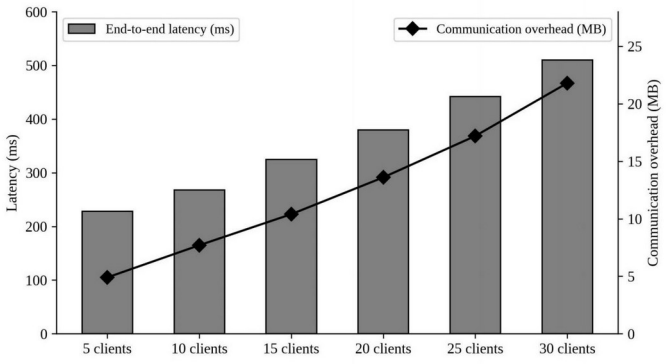


Figure 5. Latency and Communication Overhead with Varying Client Numbers

Table III retains the draft's three proposed resource profiles and numerical values for author verification. The reported latencies in these profiles are 200, 250 and 320 ms, whereas the draft reports latency below 230 ms without linking that figure to a specific profile. The available material does not identify the hardware, measurement procedure or run-level data needed to reconcile the figures or confirm that additional resources caused the reported accuracy differences.

TABLE III. RESOURCE PROFILE METRICS

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: The draft describes these values as means over three runs. The individual runs, hardware specifications and measurement records were not supplied for verification.

E. Convergence Behaviour

Figure 6 reproduces the draft's loss curves for a centralised baseline and three federated configurations. They are provisional manuscript-reported curves; the underlying run-level losses, model settings and seeds were not supplied. A proper comparison should distinguish convergence speed from final loss and quantify variation across repeated runs [Li et al., 2020; Yang et al., 2019].

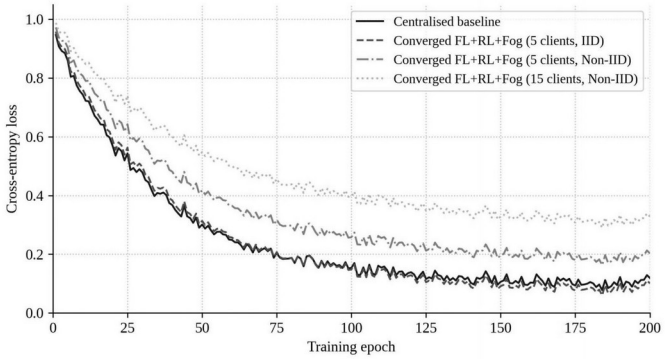


Figure 6. Training Cross-Entropy Loss Curves

Neither the curves nor the currently available records establish that federation primarily affects final loss rather than convergence rate. Both outcomes, together with per-client performance, should be tested under controlled partitions before drawing a deployment recommendation.

F. Comparative Analysis with Recent Federated IDS Literature

Table IV summarises the approaches and scope of selected published studies. Nguyen and colleagues report a 95.6% detection rate and separately report no false alarms in a real-world smart-home evaluation [Nguyen et al., 2019]. Metrics, test data and conditions differ across publications, so the values cannot be interpreted as a controlled head-to-head comparison.

TABLE IV. QUALITATIVE COMPARISON WITH INTRUSION DETECTION STUDIES

Study / system	Setting	Published contribution	Comparison limit
Nguyen et al. (2019), DÏoT	IoT	Federated self-learning anomaly detection; reports detection and false-alarm results.	Its metrics and evaluation conditions differ.
Popoola et al. (2022)	IoT edge	Federated deep learning for zero-day botnet detection.	Not an identical test to the present draft.
Liu et al. (2021)	Vehicular edge	Blockchain and FL for collaborative intrusion detection.	Different architecture and evaluation.
Bagaa et al. (2020)	IoT	ML monitoring and anomaly detection using an SDN/NFV framework.	The cited study does not establish an RL hybrid.
Han et al. (2022)	Network IDS	PPO controls intrusion-detector hyperparameters.	Does not evaluate the proposed FL–fog combination.

Note: Studies differ in datasets, threat models and metrics. A valid quantitative comparison requires matched evaluation conditions or clearly identified results from each original source.

G. Cross-Domain Generalisation

Table V reproduces the cross-domain values in the supplied draft. Their minimum and maximum are 78.2% and 92.5%, respectively; they are provisional and not independently verified. Before interpreting transfer, the study must document how features and labels were aligned, whether target labels were used, how training and test records were separated, and how results vary across repeated runs. A universal random baseline cannot be specified without the target label distribution.

TABLE V. CROSS-DOMAIN DETECTION ACCURACY (%)

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 omitted in the draft. Source and target labels, feature mappings and the underlying evaluations require verification; values should not be treated as established cross-domain results.

Deployment implications should be evaluated after the transfer protocol is specified. Personalisation and meta-initialisation are candidate approaches for new or heterogeneous clients, not guarantees of successful transfer [Balakrishnan et al., 2021; Song et al., 2022].

V. DISCUSSION AND FUTURE WORK

A. Potential Benefits and Evidentiary Limits

The proposed combination has plausible benefits: FL can support collaborative training without centralising raw records, fog placement can reduce some network travel, and an RL component can use suitable feedback to adapt a policy. None of these benefits follows automatically from combining the components. The present draft's numerical values cannot establish comparative accuracy, latency, false-positive ranking or cross-domain generalisation until the underlying experiments are verified. Excessive false positives remain an operational concern [García-Teodoro et al., 2009; Anuar et al., 2011].

B. Where the Approach Strains

Key limitations include non-IID data, communication cost, privacy–utility trade-offs and potential poisoning of client updates. The draft does not provide an independently verifiable adversarial evaluation. Published work documents poisoning threats and Byzantine-resilient approaches, but their assumptions and interaction with secure aggregation must be assessed for each deployment [Xia et al., 2023; So et al., 2021].

C. Energy and Sustainability

Energy consumption is a material deployment consideration. Fog nodes can shift computation away from devices, but some gateways are themselves power constrained and the total cost includes device, fog, network and cloud resources. Future evaluation should report hardware, power-measurement methods and the energy–accuracy trade-off rather than assume a reduction [Mahdavinejad et al., 2018; Wang et al., 2019].

D. Quantum-Era Considerations

Quantum computing presents a longer-term concern for cryptographic mechanisms that rely on

assumptions vulnerable to quantum algorithms. Appropriate post-quantum migration depends on the specific protocol and its deployment constraints [Moody and Robinson, 2022; Lu et al., 2023]. Quantum machine learning is a separate research direction; any security or speed advantage for this architecture would require a specified algorithm and direct empirical comparison [Lu et al., 2024].

E. Integration with Blockchain and Auditability

Blockchain can provide an auditable record of selected events or model updates in federated systems, while adding throughput, storage and governance costs [Xu et al., 2021; Liu et al., 2021; Alkadi et al., 2021; Sun et al., 2022; Peng et al., 2022]. The cited technical studies support auditability mechanisms; they do not by themselves establish specific regulatory requirements. Client incentives can also be designed explicitly [Feng et al., 2022]. A deployment can separate time-critical detection from periodic audit recording, provided the resulting delay, trust assumptions and verification coverage are measured.

F. Toward Real-Time Adversarial Robustness

Adversarial robustness should be evaluated against both evasive traffic and malicious federation participants. Randomised smoothing has been studied for certifiable robustness in federated adversarial learning, and Byzantine-resilient aggregation addresses a different threat model [Chen et al., 2021; So et al., 2021]. Their guarantees depend on assumptions about perturbations, attackers and protocol behaviour; neither technique establishes robustness for the proposed system without implementation and testing.

VI. CONCLUSION

The convergence of federated learning, reinforcement learning and fog computing offers a research framework for distributed IoT intrusion detection. The architecture separates local telemetry processing, collaborative model training and potential feedback-driven decisions, but its privacy, accuracy, latency, energy use and robustness depend on implementation details. The numerical results reproduced in Section IV are author-reported draft values that could not be independently verified from the supplied material. A defensible empirical conclusion requires the original data partitions, executable code, privacy accounting, hardware and network specifications, and run-level outputs. Future work should use controlled ablations and clearly defined threat models to test the proposed benefits.

Acknowledgement: The authors sincerely thank anonymous reviewers for their constructive comments, which greatly improved the rigor of the three-tier converged architecture design, the clarity of the evaluation framework discussion, and the comprehensiveness of the future research agenda. The authors also acknowledge the academic resources and research environment support provided by their respective affiliated institutions.

Author contributions: J.D.W.: Conceptualization of the FL-RL-fog converged architecture for IoT security, systematic literature review and paradigm comparison, design of the evaluation protocol, analysis of provisional performance results, Writing – Original Draft. As the first author, he led the overall research design and manuscript preparation. S.K.W.: Research supervision, validation of the methodological framework, manuscript review and polishing, project administration. As the corresponding author, he guided the research direction and finalized the manuscript.

Generative AI Statement: Generative artificial intelligence was used to assist with language polishing and drafting of the manuscript, as well as the initial generation of schematic diagrams and comparative summary tables. All AI-assisted content has been carefully reviewed and verified by the authors, who take full responsibility for the final content.

Funding: This research did not receive any specific financial grants or funding support from public, commercial, or non-profit institutions.

Declarations: The authors declare that there are no known competing financial interests or personal relationships that could have affected the objectivity of the research design, analysis process and conclusions reported in this paper.

ISSN: © 2025 INATGI (Institute of Advanced Technology and Green Innovation).

REFERENCES

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

Alkadi, O., Moustafa, N., Turnbull, B., & Choo, K. K. R. (2021). 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

Anuar, N. B., Furnell, S., Papadaki, M., & Clarke, N. (2011). A risk index model for security incident prioritisation. Security Research Institute (SRI), Edith Cowan University. https://doi.org/10.4225/75/57b52a66cd8b5

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

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

Balakrishnan, R., Akdeniz, M., Dhakal, S., Anand, A., Zeira, A., & Himayat, N. (2021). Resource Management and Model Personalization for Federated Learning over Wireless Edge Networks. Journal of Sensor and Actuator Networks, 10(1), 17. https://doi.org/10.3390/jsan10010017

Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., & Seth, K. (2017). Practical Secure Aggregation for Privacy-Preserving Machine Learning. Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, 1175–1191. https://doi.org/10.1145/3133956.3133982

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

Buczak, A. L., & Guven, E. (2016). A Survey of Data Mining and Machine Learning Methods for Cyber Security Intrusion Detection. IEEE Communications Surveys & Tutorials, 18(2), 1153–1176. https://doi.org/10.1109/COMST.2015.2494502

Chen, C., Kailkhura, B., Goldhahn, R., & Zhou, Y. (2021). Certifiably-Robust Federated Adversarial Learning via Randomized Smoothing. 2021 IEEE 18th International Conference on Mobile Ad Hoc and Smart Systems (MASS), 173–179. https://doi.org/10.1109/MASS52906.2021.00032

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

da Costa, K. A. P., 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

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

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), 4171–4186. https://doi.org/10.18653/v1/N19-1423

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

Dwork, C., & Roth, A. (2014). The Algorithmic Foundations of Differential Privacy. Foundations and Trends® in Theoretical Computer Science, 9(3-4), 211–487. https://doi.org/10.1561/0400000042

Feng, L., Yang, Z., Guo, S., Qiu, X., Li, W., & Yu, P. (2022). Two-Layered Blockchain Architecture for Federated Learning Over the Mobile Edge Network. IEEE Network, 36(1), 45–51. https://doi.org/10.1109/MNET.011.2000339

Ferrag, M. A., Friha, O., Hamouda, D., Maglaras, L., & Janicke, H. (2022). Edge-IIoTset: A New Comprehensive Realistic Cyber Security Dataset of IoT and IIoT Applications for Centralized and Federated Learning. IEEE Access, 10, 40281–40306. https://doi.org/10.1109/ACCESS.2022.3165809

François-Lavet, V., Henderson, P., Islam, R., Bellemare, M. G., & Pineau, J. (2018). An Introduction to Deep Reinforcement Learning. Foundations and Trends® in Machine Learning, 11(3-4), 219–354. https://doi.org/10.1561/2200000071

García-Teodoro, P., Díaz-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

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

Hady, A. A., Ghubaish, A., Salman, T., Unal, D., & Jain, R. (2020). Intrusion Detection System for Healthcare Systems Using Medical and Network Data: A Comparison Study. IEEE Access, 8, 106576–106584. https://doi.org/10.1109/ACCESS.2020.3000421

Han, H., Kim, H., & Kim, Y. (2022). An Efficient Hyperparameter Control Method for a Network Intrusion

Detection System Based on Proximal Policy Optimization. Symmetry, 14(1), 161. https://doi.org/10.3390/sym14010161

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

He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770 –778. https://doi.org/10.1109/CVPR.2016.90

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

Hosseinzadeh, M., Sinopoli, B., & Garone, E. (2019). Feasibility and Detection of Replay Attack in Networked Constrained Cyber-Physical Systems. 2019 57th Annual Allerton Conference on Communication, Control, and Computing (Allerton), 712–717. https://doi.org/10.1109/ALLERTON.2019.8919762

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 & Tutorials, 23(3), 1759–1799. https://doi.org/10.1109/COMST.2021.3090430

Kolias, C., Kambourakis, G., Stavrou, A., & Voas, J. (2017). DDoS in the IoT: Mirai and Other Botnets. Computer, 50(7), 80–84. https://doi.org/10.1109/MC.2017.201

Leevy, J. L., & Khoshgoftaar, T. M. (2020). A survey and analysis of intrusion detection models based on CSE-CIC-IDS2018 Big Data. Journal of Big Data, 7(1), 104. https://doi.org/10.1186/s40537-020-00382-x

Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated Learning: Challenges, Methods, and Future Directions. IEEE Signal Processing Magazine, 37(3), 50–60. https://doi.org/10.1109/MSP.2020.2975749

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

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

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

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

Lu, Y., Sigov, A., 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

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

Mahmoud, R., Yousuf, T., Aloul, F., & Zualkernan, I. (2015). Internet of things (IoT) security: Current status, challenges and prospective measures. 2015 10th International Conference for Internet Technology and Secured Transactions (ICITST), 336–341. https://doi.org/10.1109/ICITST.2015.7412116

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

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. nature, 518(7540), 529-533.https://doi.org/10.1038/nature14236

Moody, D., & Robinson, A. (2022). Cryptographic Standards in the Post-Quantum Era. IEEE Security & Privacy, 20(6), 66–72. https://doi.org/10.1109/MSEC.2022.3202589

Moustafa, N., & Slay, J. (2015). UNSW-NB15: a comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). 2015 Military Communications and Information Systems Conference (MilCIS), 1–6. https://doi.org/10.1109/MilCIS.2015.7348942

Mukherjee, M., Shu, L., & Wang, D. (2018). Survey of Fog Computing: Fundamental, Network Applications, and Research Challenges. IEEE Communications Surveys & Tutorials, 20(3), 1826–1857. https://doi.org/10.1109/COMST.2018.2814571

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 & Tutorials, 23(3), 1622–1658. https://doi.org/10.1109/COMST.2021.3075439

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. 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), 756–767. https://doi.org/10.1109/ICDCS.2019.00080

Pan, S., Morris, T., & Adhikari, U. (2015). Developing a Hybrid Intrusion Detection System Using Data Mining for Power Systems. IEEE Transactions on Smart Grid, 6(6), 3104–3113. https://doi.org/10.1109/TSG.2015.2409775

ISSN: © 2025 INATGI (Institute of Advanced Technology and Green Innovation).

Peng, Z., Xu, J., Chu, X., Gao, S., Yao, Y., Gu, R., & Tang, Y. (2022). VFChain: Enabling Verifiable and Auditable Federated Learning via Blockchain Systems. IEEE Transactions on Network Science and Engineering, 9(1), 173–186. https://doi.org/10.1109/TNSE.2021.3050781

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

Rahman, 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

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

Sattler, F., Müller, K. R., & Samek, W. (2021). Clustered Federated Learning: Model-Agnostic Distributed Multitask Optimization Under Privacy Constraints. IEEE Transactions on Neural Networks and Learning Systems, 32(8), 3710–3722. https://doi.org/10.1109/TNNLS.2020.3015958

Sattler, F., Wiedemann, S., Müller, K. R., & Samek, W. (2020). Robust and Communication-Efficient Federated Learning From Non-i.i.d. Data. IEEE Transactions on Neural Networks and Learning Systems, 31(9), 3400–3413. https://doi.org/10.1109/TNNLS.2019.2944481

Sedjelmaci, H., Senouci, S. M., & Ansari, N. (2018). 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

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

So, J., Guler, B., & Avestimehr, A. S. (2021). Byzantine-Resilient Secure Federated Learning. IEEE Journal on Selected Areas in Communications, 39(7), 2168–2181. https://doi.org/10.1109/JSAC.2020.3041404

Song, H. M., Woo, J., & Kim, H. K. (2020). In-vehicle network intrusion detection using deep convolutional neural network. Vehicular Communications, 21, 100198. https://doi.org/10.1016/j.vehcom.2019.100198

Song, J., Oh, M. H., & Kim, H. S. (2022). Personalized Federated Learning With Server-Side Information. IEEE Access, 10, 120245–120255. https://doi.org/10.1109/ACCESS.2022.3221401

Sun, Z., Wan, J., Yin, L., Cao, Z., Luo, T., & Wang, B. (2022). A blockchain-based audit approach for encrypted data in federated learning. Digital Communications and Networks, 8(5), 614–624. https://doi.org/10.1016/j.dcan.2022.05.006

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

Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., Jin, S., Quek, T. Q. S., & Poor, H. V. (2020). Federated Learning With Differential Privacy: Algorithms and Performance Analysis. IEEE Transactions on Information Forensics and Security, 15, 3454–3469. https://doi.org/10.1109/TIFS.2020.2988575

Xia, G., Chen, J., Yu, C., & Ma, J. (2023). Poisoning Attacks in Federated Learning: A Survey. IEEE Access, 11, 10708–10722. https://doi.org/10.1109/ACCESS.2023.3238823

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

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

ISSN: © 2025 INATGI (Institute of Advanced Technology and Green Innovation).