



ISSN 2047-3338

Performance-Optimization and Privacy-Preservation Using Federated Learning for IoT Anomaly Detection: A Structured Review

Marriam Salman¹, Junaid Arshad²

^{1,2}Department of Computer Science, University of Engineering and Technology, Lahore, Pakistan

¹marriamsalmanqureshi@gmail.com

Abstract— The rapid rise of the Internet of Things (IoT) has been a major factor in the demand for detection mechanisms that are not only powerful but also ensure the privacy of users. The conventional models that require raw device data to be continuously sent to them suffer from scalability issues, high demand for bandwidth, and most importantly, high privacy risks that come with the continuous transfer of raw data. FL has emerged as the distributed alternative to collaborative model training thus reducing communication costs and at the same time increasing privacy. Among the recent advances, the application of structured sparsity, hierarchical aggregation, encrypted model update, adaptive client selection, optimization-based ensembles can effectively increase the accuracy of detection and at the same time preserve efficiency in IoT networks. They have given structured reviews along with the combination of these advances, how they are viewed and how they are described, pointing out their strengths and weaknesses simultaneously. The review presents very useful insights for the making of FL-based IoT anomaly detection systems that are not only very effective but also extremely robust in terms of data confidentiality against large scale distributions.

Index Terms— Federated Learning, Anomaly Detection, Privacy Preservation, Distributed Machine Learning, Edge Computing and Intrusion Detection

I. INTRODUCTION

SOME of the new possibilities in Internet of Things ecosystems long before IoT are also constant monitoring, real time automation, and even intelligent decision making in many aspects of life, including healthcare, smart homes, automated systems of industries and even critical infrastructure. Such devices also offer highly sensitive, sensitive data, which are heterogeneous, extremely sensitive, and available in many volumes. Thus, these IoT networks are vulnerable to attacks, viruses, unauthorized access, and suspicious activity that may cause reliability and operation safety problems. Thus, effective anomaly detection has become a prerequisite in providing a secure and reliable environment for IoT.

Instead, traditional centralized anomaly detection approaches gather raw information from a central distributed device server and train the model. There is some reassurance to this, but, that it has proved equally ideal in relatively restricted environments,

and it is now becoming more of a gimmick for modern IoT deployments due to the limitations of bandwidth, latency, privacy, privacy, scaling, and the risk of failure if it can fail. These limitations have led to the movement toward decentralized learning that maintains privacy and allows for large, real world IoT deployments.

Federated Learning (FL) offers a viable solution by allowing several devices or edge nodes to work together and train detection models without sharing raw data. Just model updates are exchanged, significantly reducing privacy risk while improving communication effectiveness and accommodating non-IID and device-specific data distributions. Although FL improves scalability and privacy, challenges remain, including vulnerability to poisoning attacks, performance degradation under data heterogeneity, restricted computational ability of edge devices, and the necessity to find a balance between accuracy and strict privacy requirements.

It has also been disclosed through the systematic review that out of the recent advancements in the Federated Learning (FL)-based anomaly detection applications, five key aspects have been constantly proposed by researchers as the key contributors to the enhanced performance and privacy; the aspects are structured sparsity, hierarchical aggregation, encrypted model sharing, metaheuristic optimization, and hybrid architectures. However, in the bulk of the literature reviews, they are mostly concerned with privacy or detection processes, and in this way, they rob the practitioners of a necessary insight at the combined optimization approach to the real IoT deployments.

Not only does this systematic review address the gap but it also does so by digging out the cutting-edge FL-based IoT anomaly detection methods in the contexts of a performance-optimization and privacy-preservation perspective. It outlines the new-generation structures, measures their strengths and weaknesses, and brings forth the problems and questions that are not answered. The knowledge gained is expected to be used for the design of the forthcoming FL systems which would be capable of performing anomaly detection in a melting-pot of characteristics like, expansive, safe, and eco-friendly, in real-world IoT networks.

II. BACKGROUND

The Internet of Things (IoT) is a colossal system of heterogeneous and networked apparatus that generates streams of data to watch and manage. Such devices often lack processing resources, memory and bandwidth and traditional security mechanisms are often not applicable on such platforms. The IoT anomaly detection is designed to identify the existence of malicious activity such as suspicious traffic patterns, inappropriate access, or device malfunction that may indicate cyberattacks or faulty operation. Reliable detection entails ensuring that solutions are high accuracy and do not violate latency, privacy and resources constraints.

A) Centralized vs. Federated Learning

In centralized learning, sensor information of every device is gathered and saved in a centralized core server that is utilized to train the model. Although this method can be highly accurate with adequately curated data sets, it cannot work significant problems in the IoT setting:

- high communication overhead from continuous data transmission
- increased exposure of sensitive device-level information
- scalability limitations as device numbers grow
- dependence on an individual point of malfunction

Federated Learning (FL) overcomes these limitations by allowing devices to train models locally and just send model updates with an aggregator. This minimizes bandwidth, ensures privacy and more appropriately accommodates non-IID (non-identically distributed) data. FL can thus be used efficiently in big and heterogeneous IoT networks. However, it introduces new challenges, such as robustness against unreliable devices, varying network conditions, and potential poisoning of model updates.

Table I: Comparison Table: Centralized vs Federated Learning

Feature	Centralized Learning	Federated Learning
Data Storage	Collected at server	Remains on devices
Privacy	High risk of data leakage	Strongly preserved
Scalability	Low to moderate	High, suitable for IoT
Communication Load	Heavy (raw data)	Light (model updates only)
Fault Tolerance	Single point of failure	More robust with distributed aggregation
Adaptation to heterogeneous/non-IID data	Limited	Supports personalized and hierarchical models
Real-time suitability	Often unsuitable	High potential with edge aggregation

B) Privacy-Preserving Techniques in FL

Because IoT data may contain sensitive personal, industrial, or operational information, FL systems often integrate specialized privacy mechanisms:

Homomorphic Encryption (HE): Allows encrypted model updates to be aggregated without decryption, preventing exposure of device-level gradients.

Differential Privacy (DP): Injects controlled noise into model updates to prevent inference of individual contributions while maintaining analytical utility.

Secure Multi-Party Computation (SMPC): Distributes computations across multiple devices to ensure that no single participant learns another's data.

Blockchain-based Aggregation: Provides tamper-proof and transparent tracking of model updates, ensuring integrity and accountability in multi-party IoT environments.

These techniques strengthen confidentiality but may increase computational or communication overhead.

C) Optimization and Performance Enhancement

Since many IoT devices have limited processing power, recent research focuses on optimizing FL to reduce cost while improving detection accuracy:

Structured Sparsity: Reduces model size and communication by transmitting only important parameters. Structured Sparsity: Techniques such as Federated Structured Sparse PCA (FSSPCA) [1] reduce model size, communication cost, and computational overhead while preserving anomaly detection performance. Sparse updates reduce the load on edge networks and accelerate convergence.

Model Compression and Quantization: Decrease the precision or size of model updates to save bandwidth and device energy. Several FL frameworks, including FSSPCA [1], incorporate these strategies to enhance scalability in large IoT deployments.

Metaheuristic Optimization: Algorithms such as genetic search or Sailfish Optimization tune hyperparameters to improve detection reliability on edge devices. Sailfish Optimization (SFO) [5] dynamically optimize detection model parameters (e.g., Isolation Forest) on edge devices, balancing exploration and exploitation and improving both accuracy and efficiency without excessive communication.

Client Selection and Scheduling: Choosing clients considering resource availability, data quality, or network conditions reduces training time plus enhances overall accuracy.

Hybrid and Hierarchical Architectures: Combining device-level, edge-level, and cloud-level aggregation supports low latency and efficient scaling. Combining local computation with edge or cloud aggregation balances latency and accuracy.

III. LITERATURE REVIEW

The literature discussing Federated Learning (FL) as a technique for anomaly detection in the IoT domain indicates a very high level of support in terms of privacy protection, communication, and model simplification. However, a deep

analysis shows that the large majority of previous work has been concentrated on general FL processes rather than security-based FL processes which are specially designed for intrusion or anomaly detection.

A) FL-Based IoT Anomaly and Intrusion Detection Models

This category addresses the fundamental constraints of IoT. A key part of the literature directly examines federated learning for intrusion detection systems (FL-IDS). Research proposes a federated intrusion detection architecture for IoT networks that improves attack detection accuracy while preventing data exposure between clients [6]. Their results validate that FL reduces communication costs while achieving strong performance across attacks such as DoS, probing, and botnet activity.

Similarly, a federated anomaly detection method introduces based on deep autoencoders for edge cloud IoT environments, showing that distributed local training yields comparable accuracy to centralized learning but offers significantly stronger privacy guarantees [7]. In another study designs Fed-IDS, a lightweight FL-based intrusion detection model optimized for low-resource IoT devices, demonstrating resilience against poisoning attempts and non-IID data [8].

Structured sparsity reduces communication load while maintaining detection accuracy. Huang et al. [1] introduced a hierarchical Federated Structured Sparse PCA (FSSPCA) framework that efficiently captures device-level variations while preserving privacy; however, its linear PCA model may not adequately detect complex anomalies, particularly in highly heterogeneous IoT environments. The approach captures device-level variations efficiently and enables hierarchical aggregation at edge servers, improving detection speed and reducing network load. By reducing communication overhead and maintaining accuracy across devices, this approach supports scalable deployments in heterogeneous environments.

Key points:

- FL-IDS models reduce data exposure and support distributed IoT deployments.
- Autoencoders and Deep Neural Network based distributed IDS models achieve high accuracy on IoT attacks.
- Lightweight FL architecture achieves high detection accuracy despite constrained device hardware.
- Hierarchical aggregation reduces network overhead.

B) Privacy-Preservation Techniques Relevant to IoT Security

Privacy mechanisms in FL play a central role in enabling secure anomaly detection. Secure Aggregation Introduce a protocol that allows a server to aggregate encrypted model updates without accessing individual gradients critically important for sensitive IoT information [9]. Differential Privacy for federated learning is study showing that controlled noise injection prevents device-level data inference but must be carefully tuned to avoid degrading anomaly detection accuracy [10].

Additionally, feature leakage vulnerabilities in collaborative learning systems demonstrate motivating the need for

encryption-enhanced FL-IDS frameworks. Blockchain-supported FL has also been explored analyzing how distributed ledgers provide auditability and tamper resistance for federated anomaly detection pipelines [11].

It provides decentralized trust, ensuring tamper-proof tracking of model updates and preventing malicious alterations. This approach is especially useful for large-scale IoT networks with multiple stakeholders. Arazzi et al. [4] demonstrate that blockchain-enabled FL can maintain privacy, integrity, and accountability in anomaly detection.

Key points:

- Secure Aggregation ensures encrypted, privacy-safe update sharing.
- Differential Privacy provides provable protection but may reduce anomaly detection accuracy if noise is excessive.
- The leakage of features as discussed in the feature leakage studies points out the necessity of a more secure encryption and aggregation scheme.
- Blockchain systems increase trust, however, at the cost of latency.

C) Performance Optimization and Resource Efficiency

At the IoT anomaly detection, models need to be run with little power, memory, and network bandwidth. The Federated Averaging (FedAvg) algorithm, forming the foundation for communication-efficient FL. Building on this show that edge-assisted federated learning improves real-time detection capabilities by reducing device workload [12].

Sparse and personalized FL methods help overcome non-IID IoT data distributions, as discussed [13]. Optimization techniques examined further highlight the importance of adaptive aggregation and resource-aware client selection, both of which improve convergence and maintain accuracy under IoT constraints [14].

Hybrid FL frameworks combine local, edge, and cloud-level aggregation to achieve near-real-time anomaly detection. Rampone et al. proposed a framework that integrates centralized and federated training, achieving low latency while maintaining accuracy across heterogeneous IoT environments [2]. Hierarchical Federated Learning (HFL) introduces multiple levels of aggregation for scalable and efficient processing.

Sun et al. [3] devised FedMADE, a dynamic aggregation framework robust against non-IID data and malicious updates, providing improved recall and F1-scores, but with increasing aggregation complexity as the network scales.

Key points:

- Sparse and personalized FL is effective for heterogeneous IoT devices.
- Client selection and hierarchical edge aggregation reduce latency.
- FedAvg and optimized aggregation minimize bandwidth costs while supporting large-scale deployments.

D) Datasets, Threat Models, and Evaluation Gaps

This research focuses on improving FL-based detectors major limitation in previous work is the lack of consistent benchmarking across FL-based anomaly detection studies.

Common datasets used for evaluation include:

- NSL-KDD and KDD-Cup99 (traditional IDS datasets)
- UNSW-NB15 (modern network attacks)
- TON_IoT (IoT telemetry, logs, and network attacks)
- BoT-IoT (IoT botnet and DDoS attacks)

However, many studies evaluate performance using small-scale or simulated settings, which do not fully capture realistic IoT conditions such as unstable connectivity, device failure, and heterogeneous data distributions. Moreover, only a limited number of studies evaluate robustness against poisoning attacks or adversarial updates, which are critical threats to FL-IDS.

Key points:

- TON_IoT and BoT-IoT provide realistic IoT attack data.
- Inconsistent datasets and metrics limit comparability of research results.
- Few studies explore FL robustness against poisoning or adversarial manipulation.
- Real-world IoT evaluation remains limited.

Table II: Summary Table of Key Literature

Study	Approach	Metrics	Key Contribution
[1] Huang et al., 2025	FSSPCA, hierarchical FL	F1-score, False Positives	Enhances anomaly detection by capturing device-level variations, reducing network load, and preserving privacy.
[2] Rampone et al., 2025	Hybrid FL framework	Accuracy, Latency	Enables near-real-time anomaly detection through hybrid edge-cloud aggregation.
[4] Arazzi et al., 2024	HE + Blockchain FL	Privacy metrics, Accuracy	Provides strong privacy and tamper-proof FL aggregation using homomorphic encryption and blockchain.
[5] Aravam Babu & Bagubali, 2025	Sailfish-optimized ensemble models	Recall, F1-score	Dynamically tunes detection parameters to improve recall and F1-score, balancing performance and privacy.
[6] Nguyen et al., 2021	FL-IDS for IoT (D ² IoT)	Detection Rate, False Alarm Rate	Demonstrated distributed intrusion detection with improved accuracy for IoT attacks.
[7] Chen et al., 2020	Federated Autoencoder	Reconstruction Error, Accuracy	Provided privacy-preserving anomaly detection using

	for Anomaly Detection		distributed deep learning.
[8] Salam et al., 2022	Lightweight FL for IoT intrusion detection	Accuracy, Resilience metrics	Optimized for low-resource IoT devices, resilient to poisoning.
[9] Behnia et al., 2024	Secure Aggregation for FL	Privacy Level, Encryption Overhead	Prevents gradient leakage in federated updates.
[10] Fu et al., 2025	Differential Privacy in FL	Privacy-Accuracy Trade-off	Systematically reviewed DP techniques in FL, highlighting tuning challenges.
[11] Mothukuri et al., 2021	Blockchain-Enhanced FL	Integrity, Latency	Analyzed blockchain-enabled trust mechanisms for FL security.
[12] Imteaj et al., 2022	Edge-Assisted FL	Latency, Energy Use	Demonstrated performance benefits of edge-tier aggregation for IoT.
[13] Kairouz et al., 2021	Advances in FL (survey)	General FL metrics	Comprehensive survey on FL advances and open problems.
[14] Li et al., 2020	Federated optimization in heterogeneous networks	Convergence, Accuracy	Optimization methods for non-IID data in FL.

IV. DISCUSSION

As the review of recent literature reveals, Federated Learning (FL) has significantly contributed to advancing privacy-preserving and scalable Internet of Things (IoT)-based abnormality detection. Although this has been achieved, implementation of FL-based anomaly detection in a real-world IoT system is still not feasible because of correlated limitations on performance, privacy, communication overhead, data heterogeneity, and system robustness. Another common finding is the challenge faced by ensuring tight privacy protection as well as high detection rate in resource-limited and dynamic IoT networks.

The most popular approaches to preventing gradient inference attacks in models are homomorphic Encryption (HE) and Secure Aggregation. Despite these methods providing good privacy assurance, they come with intensive computational and communication expenses that low-power devices in an IoT cannot handle. HE-based techniques have mathematically valid protection but will generally slow down training and undermine real-time detection of anomalies. Differential Privacy (DP) provides weaker protection by introducing noise to updates, but too much noise can greatly decrease accuracy, especially in detecting subtle or emerging attack patterns, e.g., those in TON_IoT and BoT-IoT datasets. The trade-off that occur

among privacy and model usability is one of the biggest shortcomings around FL-based IoT security studies that persist to this date.

Among these solutions was sparsity, resource-aware client selection, and hierarchical edge cloud aggregation to overcome the limited bandwidth and high latency problems. For instance, structured sparsity eliminates communication by merely passing on the necessary variables. Thus, it can be beneficial for limited devices but often not enough for complex anomaly signatures. Another way to improve efficiency by focusing on reliable devices, is by choosing not to leave out nodes that contain rare or vital data. Another performance approach, hierarchical edge aggregation is less burden-draining and efficient training and has edge bottlenecks and vulnerability to device dropouts.

These Federated IDS models that weigh less have been created to tackle IoT limitations by using smaller architectures such as autoencoders or shallow neural networks to lessen the intrinsic drawbacks of FL-IDS. These methods show a very high performance for detection; however, the accuracy varies with the diversity of attacks, complexity of data, and non-IID nature of IoT devices. One of the most severe difficulties is non-IID distributions, because they cause the model training to be slow, the resulting models to be biased, and the ability of the models to generalize to different devices to decrease.

Another critical area is system resilience where the blockchain-backed FL has become popular owing to its resistance to tampering and its transparent aggregation methods. Such systems are believed to be highly trustworthy and maintain integrity, but they come along with high energy consumption and latencies. The overhead coming from the high consensus is a limiting factor for the large-scale IoT areas as it impedes the rapid deployment of these systems.

One more limitation that was common across the studies is the inconsistent application of datasets, evaluation criteria, and threat models. The different datasets like NSL-KDD, UNSW-NB15, TON_IoT, and BoT-IoT have huge differences related to their scale, features, and how realistic they are, thus making it very difficult for researchers to compare the results across studies. Furthermore, although decentralized systems with malicious devices that can tamper with model updates should consider poisoning resistance and adversarial robustness as critical points, these topics are still terribly under-researched.

Table III: Comparative Analysis of Federated Learning Techniques for IoT Anomaly Detection

Technique / Model	Privacy Mechanism	Strengths	Limitations
Structured Sparse FL	Localized model updates; minimal data exposure	Low communication cost; suitable for constrained IoT; preserves privacy; good for non-IID personalization	Linear feature extraction may miss complex attack patterns; tested on small-scale IoT data

Hybrid Edge-Cloud FL	Edge-level aggregation; selective update sharing	Low latency; scalable; reduces device workload; near-real-time detection; balances performance & privacy	Edge bottlenecks possible; sensitive to client dropouts; edge-device resource use not fully optimized
Encrypted & Blockchain-Enhanced FL	Secure aggregation; homomorphic encryption; blockchain logging	Strong confidentiality; dual-layer privacy; tamper-proof model update logging; ensures integrity & accountability	High computational overhead; slower convergence; needs powerful edge devices; increased latency
Optimization-Driven FL (Ensembles, Adaptive Aggregation)	Differential Privacy; weighted aggregation	High detection accuracy; adapts to device data; efficient communication; improves recall & F1-score	Requires careful tuning; may degrade under heavy noise or adversarial updates; sensitive to non-IID distributions

Table III provides a structured comparison of the main Federated Learning technologies used to detect anomaly detection for IoTs that reveal how the two types of approaches differ in privacy, efficiency, and scalability. Structured Sparse FL is designed to reduce communication, suitable for narrowly limited IoT devices but less suited to the modeling of complex attacks. Hybrid Edge-Cloud FL's speed and lower workload can be made the advantage but is not suited to scalability and could suffer from edge-level bottlenecks. In encrypted, blockchain-enhancement FL provides strong confidentiality and tamper resistance but reduces computational complexity and slow training. The optimization-driven FL techniques improve accuracy and adaptive response to heterogeneous device data, but must be controlled, are vulnerable to noise or adversarial updates and may respond in difficult IoT environments to sound or adversarial updates.

V. CONCLUSION

This systematic literature review has assessed the application Federated Learning (FL) to optimize performance and privacy in IoT based anomaly detection. Although FL provides a radical move in terms of centralized models (reducing the exposure of data and lowering the cost of communication) its practical implementation is inherently trade-off. There is no one existing solution that is a perfect balance between detection accuracy, privacy guarantee and efficient operation within the limitations of the heterogeneous and resource-limited IoT networks. Structured sparsity, encrypted aggregation and adaptive optimization are techniques that solve one or more dimensions of this trilemma at the cost of another. The future will rely on hybrid designs that will balance these factors and the stricter standardization of the benchmarks and emphasis on the

resistance to adversarial threats. Finally, FL will continue to be a foundation toward designing scalable, secure, and smart IoT ecosystems as long as these overlapping issues are addressed through combined solutions.

REFERENCES

- [1]. X. C. Huang, X. Li, and X. Xiu, "Federated Structured Sparse PCA for Anomaly Detection in IoT Networks," arXiv preprint, arXiv:2503.23981, 2025.
- [2]. G. Rampone, T. Ivaniv, and S. Rampone, "A Hybrid Federated Learning Framework for Privacy-Preserving Near-Real-Time Intrusion Detection in IoT Environments," *Electronics*, vol. 14, no. 7, art. no. 1430, Apr. 2025, doi: 10.3390/electronics14071430.
- [3]. S. Sun, P. Sharma, K. Nwodo, A. Stavrou, and H. Wang, "FedMADE: Robust Federated Learning for Intrusion Detection in IoT Networks Using a Dynamic Aggregation Method," arXiv preprint, arXiv:2408.07152, 2024.
- [4]. M. Arazzi, S. Nicolazzo, and A. Nocera, "A Fully Privacy-Preserving Solution for Anomaly Detection in IoT using Federated Learning and Homomorphic Encryption," *Information Systems Frontiers*, Springer, 2024.
- [5]. A. Aravam Babu and A. Bagubali, "Federated Learning With Sailfish-Optimized Ensemble Models for Anomaly Detection in IoT Edge Computing Environment," *IEEE Access*, vol. 13, pp. 53171–53187, Mar. 24, 2025, doi: 10.1109/ACCESS.2025.3554301.
- [6]. T. D. Nguyen, S. Marchal, M. Miettinen, H. Fereidooni, N. Asokan, and A.-R. Sadeghi, "D²IoT: A federated self-learning anomaly detection system for IoT," in *Proc. IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*, 2021, pp. 122–132.
- [7]. Y. Chen, X. Qin, J. Wang, C. Yu, and W. Gao, "FedHealth: A federated transfer learning framework for wearable healthcare," *IEEE Intelligent Systems*, vol. 35, no. 4, pp. 83–93, 2020.
- [8]. M. A. Salam, S. H. U. Abidi, and A. A. A. Al-Fuqaha, "FedIDS: A federated learning approach for intrusion detection in IoT networks," in *Proc. IEEE International Conference on Communications (ICC)*, 2022, pp. 1–6.
- [9]. Behnia, R., Riasi, A., Ebrahimi, R., Chow, S.S.M., Padmanabhan, B. & Hoang, T. *Efficient secure aggregation for privacy-preserving federated machine learning* (2024). arXiv:2304.03841
- [10]. J. Fu, Y. Hong, X. Ling, L. Wang, X. Ran, Z. Sun, W. H. Wang, Z. Chen, and Y. Cao, "Differentially private federated learning: A systematic review," 2025. <https://arxiv.org/abs/2405.08299>.
- [11]. V. Mothukuri, R. M. Parizi, S. Pouriyeh, Y. Huang, A. Dehghantaha, and G. Srivastava, "A survey on security and privacy of federated learning," *Future Generation Computer Systems*, vol. 115, pp. 619–640, 2021.
- [12]. A. Imteaj, U. Thakker, S. Wang, J. Li, and M. H. Amini, "A survey on federated learning for resource-constrained IoT devices," *IEEE Internet of Things Journal*, vol. 9, no. 1, pp. 1–24, 2022.
- [13]. P. Kairouz et al., "Advances and open problems in federated learning," *Foundations and Trends® in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [14]. T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine Learning and Systems*, vol. 2, pp. 429–450, 2020.