CITE PUBLICATIONS

Journal of Analytical and Applied Computer Sciences
Volume: 2, Issue: 2, 2026 (July - December), Pages: 22-31
ISSN (Online): 3139-3349 | DOI: https://doi.org/10.67914/jaacs/rsa/2.2/22-31
Research Article

Privacy-Preserving Intrusion Detection in Internet of Things Networks Using Federated Learning and Federated Averaging

Umar Muhammad Badikko[1*], Peter Buba Zirra[2]

[1]Department of Computer Science, Federal University of Kashere, Nigeria.
[2]Department of Computer Science, Federal University of Agriculture, Mubi, Adamawa State, Nigeria.

[*]**Correspondence**
Umar Muhammad Badikko
muhdbadikko@gmail.com

Received: 15 April 2026 / Accepted: 28 July 2026 / Published: 31 August 2026

The increasing deployment of Internet of Things (IoT) devices has created significant security and privacy concerns because connected devices continuously generate sensitive network data. Conventional centralized intrusion-detection systems require these data to be transmitted to a central server, potentially exposing confidential information and creating communication and scalability challenges. This study investigated a privacy-preserving intrusion-detection approach based on horizontal Federated Learning (FL), Random Forest (RF), and the Federated Averaging (FedAvg) algorithm. The proposed architecture enables participating IoT devices to retain raw network traffic locally while transmitting model parameters to a central aggregation server. A model-security validation stage was incorporated to assess the trustworthiness of local model parameters before their aggregation into the global model. The Edge-IIoTset cybersecurity dataset was used for model development and evaluation. The dataset contains 2,219,201 samples, 63 original features, a binary attack label, and 14 labelled attack classes, with a reported distribution of 73% normal and 27% attack traffic. Data preprocessing involved feature cleaning, removal of redundant and high-cardinality attributes, elimination of duplicate and erroneous records, and one-hot encoding of categorical variables. Experimental findings showed that the global federated model achieved 99.7% accuracy and an F1-score of 0.90 in the reported binary evaluation. Accuracy across five participating IoT device categories ranged from 99.3% to 99.9%. The results indicate that the federated approach can provide effective intrusion detection while retaining raw data at the edge. Nevertheless, the study identifies data heterogeneity, minority-class detection, computational overhead, and poisoning attacks as continuing challenges. The findings support Federated Learning as a promising architecture for privacy-preserving and scalable IoT intrusion detection.

Keywords: *Internet of Things, privacy preservation, Federated Learning, FedAvg, intrusion detection, Random Forest, Edge-IIoTset*

Introduction

Artificial Intelligence has emerged as a cornerstone of modern computing, enabling systems to learn from data, recognize patterns, and make autonomous decisions (Russell & Norvig, 2021). The Internet of Things has transformed computing by connecting physical objects capable of sensing, processing, communicating, and exchanging information. IoT devices

are now deployed across industrial systems, transportation, healthcare, households, and other application domains. The resulting interconnected environment enables real-time monitoring and automated decision-making but simultaneously expands the potential attack surface of networked systems. The dissertation identifies the compromise of sensitive data generated by IoT devices as a major security and privacy concern. IoT security is particularly challenging because devices continuously produce heterogeneous data and may possess different computational and communication capabilities. Anomalies can arise from sensor failures, environmental changes, human errors, or deliberate cyberattacks. Consequently, effective intrusion-detection systems are required to identify abnormal network behavior while maintaining the confidentiality and integrity of IoT data. Traditional machine-learning intrusion-detection systems generally follow a centralized architecture. Data from multiple devices are collected and transmitted to a central server, where the machine-learning model is trained. Although this architecture facilitates centralized data processing, transferring sensitive network traffic creates privacy and security concerns. The dissertation notes that conventional centralized approaches can expose data to interception or misuse. Federated Learning offers a different paradigm. Instead of moving raw data to a central location, each participating device trains a local model using its own data. Model parameters are subsequently shared with an aggregation server, which combines the updates into a global model. Thus, the raw data remain at the originating devices. The dissertation identifies this arrangement as a central advantage of FL for IoT privacy protection. However, Federated Learning does not automatically eliminate security threats. A compromised client can potentially contribute malicious or poisoned updates to the federated system. The dissertation specifically identifies model poisoning attacks as a research gap, where maliciously manipulated local data can produce harmful model updates and ultimately affect the global model. Therefore, this study focuses not only on the detection performance of a federated intrusion-detection system but also on the architectural benefits of decentralization, privacy preservation, and validation of model updates.

Privacy and Security Challenges in IoT Intrusion Detection

Centralized Intrusion Detection

In a centralized machine-learning architecture, IoT devices transmit their generated network data to a central server. The server then processes the data, trains a model, and distributes the resulting knowledge to the relevant devices. Although this arrangement simplifies centralized model management, the movement and storage of raw IoT traffic create privacy risks. The dissertation contrasts centralized systems with FL because centralized systems require sensitive data to be transmitted to a central location, whereas FL allows models to be trained directly at the edge. For large-scale IoT environments, centralized processing can also create communication and processing bottlenecks. The dissertation associates centralized architectures with scalability limitations arising from the need to continuously collect large quantities of IoT data at central servers.

Federated Learning as a Privacy-Preserving Alternative

Federated Learning reverses the conventional data-flow model. Instead of transferring raw training data, the model is distributed to participating clients. Local training is performed on the client, after which model parameters are sent to the aggregation server. The proposed architecture in the dissertation specifically uses a horizontal FL paradigm, where participating IoT devices have similar feature spaces but possess different local observations. Each device therefore retains its raw traffic locally and periodically communicates model updates to the server. The privacy advantage arises from the fact that the central server does not directly receive the original network traces used for local training.

Security Risks in Federated Learning

Despite its privacy advantages, FL introduces its own security considerations. In particular, a malicious or compromised client may submit manipulated model parameters. The dissertation identifies model poisoning as a significant research gap. An attacker may compromise a device and cause it to generate poisoned training data, resulting in an undesirable local model. If such a model is aggregated without validation, it could affect the global model and consequently the predictions made by other participating devices. Therefore, privacy preservation must be considered together with the integrity of the federated learning process.

Materials and Methods

Research Design

The study developed an edge-based intrusion/anomaly-detection system using horizontal Federated Learning, Random Forest, and FedAvg. The architecture was designed around distributed IoT clients and a central aggregation server.

The conceptual workflow consists of:

Global model initialization → Global model distribution → Local data processing → Local model training → Model-security validation → Trusted parameter upload → FedAvg aggregation → Global model redistribution.
The process continues until the specified accuracy condition is reached.

Edge-IIoTset Dataset

The study used the Edge-IIoTset cybersecurity dataset because it covers IoT and Industrial IoT environments and includes multiple attack scenarios. The dataset characteristics reported in the dissertation are:

Dataset characteristic	Description
Total samples	2,219,201
Original features	63
Numeric features	38
Categorical features	10
Flag features	15
Binary target	Attack_label
Multi-class target	Attack_type
Normal traffic	73%
Attack traffic	27%

The binary label identifies normal traffic as 0 and attack traffic as 1, while the multi-class target contains 14 labelled attack classes.

Data Preprocessing

Data preprocessing was undertaken before model training to improve the reliability and suitability of the input data.

The procedures included:
a. Harmonizing timestamps;
b. Removing high-cardinality identifiers;
c. Discarding raw payload information;
d. Eliminating redundant features;
e. Removing duplicate records;
f. Removing empty and erroneous columns; and
g. Converting categorical attributes into numerical form using one-hot encoding.

These procedures were intended to reduce unnecessary data complexity and improve the quality of the model input.

Federated Architecture

The proposed architecture employs multiple IoT edge clients and an aggregation server. Each client maintains its local network traffic and trains its local anomaly-detection model. The client does not send raw network data to the server. Instead, the trained model parameters are communicated to the aggregation server. The server combines trusted updates using FedAvg and redistributes the resulting global model. This architecture provides the principal privacy-preserving mechanism of the proposed system.

Model-Update Security Validation

A notable feature of the proposed system is the inclusion of a security-validation stage between local training and global aggregation. After local training, the model parameters are examined for reliability before they are incorporated into the global model. The dissertation describes this process as a mechanism for checking the trustworthiness of retrained parameters before aggregation. The intended purpose is to reduce the possibility that malicious or infected edge devices will introduce harmful model updates into the global model. The dissertation specifically states that this approach is intended to improve model reliability and security.

Federated Averaging

FedAvg was used to combine local model parameters into the global model. The process can be summarized as:

a. Initialize the global model.
b. Distribute the global model to participating clients.
c. Train the local model using local IoT data.
d. Validate the local model update.
e. Upload trusted model parameters.
f. Aggregate the parameters using FedAvg.
g. Generate the updated global model.
h. Redistribute the global model.
i. Repeat the process.

This allows the participating devices to collaboratively improve the shared model without centralizing their raw training data.

Classification

The system performed both binary and multi-class classification. For binary classification, Attack_label was used to distinguish:
0 = Normal
1 = Attack
The framework was subsequently extended to a 15-class classification arrangement using a softmax classifier, with normal traffic and different attack categories represented separately.

Software and Experimental Environment

The study was implemented using Python 3.11. The reported tools included TensorFlow 2.16, Keras 3, Pandas, NumPy, scikit-learn, and Matplotlib. The final experiment was executed using Google Colab.

Results

Multi-Class Classification

The reported multi-class classification achieved an overall accuracy of 97%. However, the macro-average F1-score was 0.69. The difference between the high overall accuracy and lower macro-F1 is important. It indicates that the model performed strongly on classes with substantial representation but had comparatively weaker recognition of some minority attack categories. Thus, while the model is effective at distinguishing the dominant patterns in the dataset, the result also demonstrates the challenge of detecting less-represented attack types.

Performance Across IoT Clients

The reported performance of the participating IoT devices is presented below.

IoT Device	Accuracy	F1-Score
Smart Camera	0.996	0.91
Smart Thermostat	0.993	0.89
Smart Door Lock	0.999	0.92
Industrial Sensor	0.995	0.90
Smart Lighting System	0.994	0.89
Global FL-Aggregated Model	**0.997**	**0.90**

The individual devices achieved accuracy between 99.3% and 99.9%, while the global model achieved 99.7% accuracy and an F1-score of 0.90. The consistency of performance across the clients suggests that the federated model was able to generalize across different IoT traffic sources. The dissertation attributes small differences between devices to variations in traffic volume, network behavior, and communication patterns.

Privacy-Preserving Performance

Binary Detection Performance

The binary evaluation produced the following results:

Class	Accuracy	Precision	Recall	F1-Score
Normal	1.0000	0.9400	0.9100	0.8900
Attack	0.9941	0.9286	0.9086	0.9186
Binary Average	**0.9970**	**0.9343**	**0.9093**	**0.9043**

The global binary evaluation therefore reported 99.7% accuracy, 0.9343 precision, 0.9093 recall, and 0.9043 F1-score. These results indicate that the system was capable of identifying malicious traffic while maintaining strong overall classification performance.

Privacy Advantage

The principal contribution of the architecture is that the high detection performance is obtained without requiring the raw client data to be transferred to the central aggregation server. The dissertation explicitly describes the proposed system as locally training models on raw data stored at different IoT edge devices, while only model parameters are sent to the aggregation server. This creates an important distinction from conventional centralized intrusion detection.

Characteristic	Centralized approach	Proposed FL approach
Raw IoT data	Sent to central server	Retained locally
Model training	Centralized	Local/client-side
Model aggregation	Central	FedAvg
Privacy protection	Limited	Improved
Client data exposure	Higher	Reduced
Scalability	Potentially constrained	Distributed
Global model	Centrally trained	Federatively generated

The dissertation argues that decentralized training can reduce communication overhead and transmission latency while allowing the architecture to scale across distributed IoT devices.

Discussion

The findings demonstrate that the proposed Federated Learning (FL)-based intrusion detection architecture can achieve strong detection performance while maintaining the decentralized nature of IoT data. The global federated model achieved 99.7% accuracy and an F1-score of 0.90 in binary classification, while the participating IoT devices achieved accuracy values ranging from 99.3% to 99.9%. These findings suggest that collaborative learning can be performed across distributed IoT environments without requiring the participating devices to transfer their raw network traffic to a central server. This is particularly important in IoT environments, where centralized data collection may create privacy risks, increase communication requirements, and create scalability challenges. The high binary classification accuracy obtained in this study indicates that the proposed model was effective in distinguishing normal from malicious network traffic. This finding supports the growing evidence that Federated Learning can provide an effective alternative to centralized machine-learning approaches for cybersecurity applications. Weinger et al. (2022) similarly demonstrated that Federated Learning can improve anomaly detection in IoT environments by allowing distributed devices to collaboratively learn from decentralized data. Their study emphasized the importance of FL in situations where data are naturally distributed among multiple devices and where centralized data collection may be undesirable. The present findings extend this perspective by demonstrating that a federated architecture based on local model training and global aggregation can achieve high binary detection performance using the Edge-IIoTset dataset. An important contribution of the present study is the demonstration that privacy-oriented decentralized learning does not necessarily result in poor detection performance. The proposed architecture retains raw traffic data at participating IoT clients and communicates model parameters rather than the original data. This approach is consistent with the wider development of FL for IoT applications, where privacy preservation and distributed intelligence are increasingly considered essential requirements. Nguyen et al. (2021) observed that FL is particularly suitable for IoT systems because IoT data are inherently decentralized and often generated by devices with different ownership and operational characteristics. Similarly, the high performance of the present global model suggests that collaborative learning can be achieved while reducing the exposure associated with centralized storage

of sensitive network data. The device-level results further indicate that the proposed architecture was relatively stable across different categories of IoT devices. The reported accuracy values ranged from 99.3% for the smart thermostat to 99.9% for the smart door lock, while the aggregated global model achieved 99.7%. This relatively narrow performance range suggests that the federated aggregation process was able to develop a global model that maintained strong predictive performance across heterogeneous IoT traffic sources. This is an important observation because data heterogeneity is one of the principal challenges associated with Federated Learning. In practical IoT environments, different devices may generate substantially different traffic volumes, communication patterns, and operational behaviours. Such non-independent and identically distributed (non-IID) data can affect local model convergence and reduce the generalizability of the global model. Therefore, the strong device-level performance reported in this study indicates that the proposed architecture was capable of maintaining useful collaborative learning despite variations among participating IoT clients.
However, the multi-class results provide a more nuanced interpretation of the proposed model's performance. Although the overall multi-class accuracy was 97%, the macro-F1 score was 0.69. This difference demonstrates that overall accuracy alone may provide an overly optimistic representation of intrusion-detection performance, particularly when the dataset is imbalanced. The Edge-IIoTset dataset used in the study contains a higher proportion of normal traffic than attack traffic, and some attack categories have substantially fewer observations than others. Consequently, a model may achieve high overall accuracy by performing particularly well on dominant classes while showing weaker performance on minority attack categories. This explains why the present study achieved strong overall multi-class accuracy but a considerably lower macro-F1 score. The difference between accuracy and macro-F1 is particularly important from a cybersecurity perspective. In real-world intrusion detection, failure to identify a rare attack may have serious consequences even when the overall classification accuracy is high. Therefore, the macro-F1 score provides a more balanced indication of how effectively the model performs across different attack categories. The present findings consequently suggest that future development of the system should focus on improving the recognition of minority attacks rather than relying exclusively on improvements in overall accuracy. Approaches such as class balancing, cost-sensitive learning, adaptive sampling, and specialized aggregation strategies may improve detection performance for underrepresented attack categories. The strong binary classification performance should also be interpreted in relation to the results of the multi-class experiment. Binary classification only requires the model to distinguish between normal and malicious traffic, whereas multi-class classification requires the model to identify the specific category of attack. The latter is inherently more difficult because several attack types may share similar network characteristics. Consequently, the reduction in macro-F1 observed in the multi-class experiment is not necessarily an indication that the proposed federated architecture failed; rather, it highlights the increased complexity associated with distinguishing among multiple heterogeneous attack categories. This observation supports the argument that intrusion-detection studies should report multiple evaluation metrics rather than relying exclusively on accuracy. Another important aspect of the proposed model is the integration of a model-security validation stage before local parameters are incorporated into the global model. While Federated Learning improves data privacy by retaining raw data at participating clients, it also introduces new security threats. A compromised client may attempt to manipulate local training data or submit malicious model parameters to influence the global model. This type of threat is generally described as data poisoning or model poisoning. Therefore, a privacy-preserving federated system must also consider the integrity and trustworthiness of client updates. The inclusion of a validation stage in the proposed architecture represents an important security-oriented design feature because it recognizes that decentralized learning does not automatically guarantee secure learning. The need to secure model aggregation is increasingly important as FL moves from controlled experimental environments to large-scale real-world applications. The ability of a malicious participant to influence a global model remains one of the major security challenges associated with collaborative learning. Therefore, future implementations of the proposed architecture could be strengthened through techniques such as robust aggregation, anomaly detection on client updates, secure aggregation, differential privacy, and cryptographic mechanisms. These approaches could provide additional protection against malicious participants while preserving the decentralized nature of the system.

The comparison with existing federated intrusion-detection approaches also indicates that the proposed model achieved competitive performance. The study benchmarked its performance against the federated intrusion-detection approach reported by Rashid et al. (2023), which achieved approximately 92.49% accuracy in its experimental setting. The 99.7% binary accuracy obtained by the proposed model therefore suggests substantial improvement within the context of the present experiment. However, such comparisons should be interpreted carefully because differences in datasets, preprocessing procedures, class distributions, experimental configurations, algorithms, and evaluation protocols can substantially influence reported performance. Consequently, the findings should not be interpreted to mean that the proposed architecture will necessarily achieve the same improvement under all IoT environments. A more important contribution of the proposed system is therefore the combination of strong detection performance with privacy preservation. In conventional centralized learning, IoT devices must transmit potentially sensitive traffic data to a central server. In contrast, the proposed FL architecture allows the devices to retain their raw data while participating in collaborative model development. This reduces the need for centralized data storage and may also reduce the communication burden associated with continuously transferring large volumes of raw network traffic. As IoT

deployments continue to expand, this distributed learning approach may become increasingly relevant for environments where privacy regulations, network limitations, or organizational boundaries prevent unrestricted data sharing.

Nevertheless, Federated Learning also introduces practical deployment challenges. Participating IoT devices may have different computational capabilities, communication bandwidths, data distributions, and availability. These factors can result in delayed updates, communication bottlenecks, and inconsistent local model quality. In addition, conventional FedAvg may not always perform optimally when client data are strongly non-IID. Future studies should therefore investigate alternative aggregation algorithms, including FedProx, FedNova, and other adaptive aggregation mechanisms that may better accommodate heterogeneous IoT environments. The findings indicate that Federated Learning provides a promising foundation for privacy-preserving intrusion detection in IoT environments. The proposed architecture achieved strong binary detection performance and maintained high accuracy across multiple participating IoT devices while retaining raw data locally. However, the multi-class results also demonstrate that high overall accuracy should not be interpreted as evidence of equally strong performance across all attack categories. The lower macro-F1 score highlights the continuing challenge of minority-class detection and suggests that future research should focus on class imbalance, heterogeneous client data, robust aggregation, and protection against poisoning attacks. Therefore, the principal contribution of this study lies not only in its high detection accuracy but also in demonstrating the potential of a decentralized architecture to balance intrusion-detection effectiveness with the privacy and scalability requirements of modern IoT environments.

Comparison with Previous Federated Approaches

The dissertation reports that Rashid et al. (2023) achieved approximately 92.49% accuracy using a federated intrusion-detection approach, compared with approximately 93.92% for their conventional centralized model. In comparison, the present study reports 99.7% binary accuracy in its detailed evaluation. The improvement should, however, be interpreted within the specific experimental setup and dataset used in each study rather than as evidence that FL universally outperforms centralized learning.

Privacy–Utility Trade-Off

The dissertation acknowledges that the federated approach faces a utility trade-off resulting from biased or heterogeneous data. This is particularly important because different IoT devices generate different traffic patterns. Although the global model achieved high binary accuracy, the multi-class macro-F1 of 0.69 indicates that classification performance is not uniform across all attack categories. Therefore, the privacy-preserving architecture is promising, but improved handling of non-IID data and minority attack classes remains necessary.

Limitations

The study identifies several limitations relevant to deployment.

Model Poisoning

The current architecture remains exposed to sophisticated FL-specific attacks, particularly data poisoning and model poisoning.

Heterogeneous Data

Differences between client datasets can reduce model utility. This is particularly relevant in IoT networks where devices have different traffic patterns and operational roles.

Minority Attack Classes

Although binary detection was highly successful, multi-class classification showed weaker performance for some underrepresented attack categories, reflected by the macro-F1 of 0.69.

Computational Overhead

The dissertation identifies computational overhead associated with real-time federated operation and the capabilities and constraints of FL algorithms as limitations.

Future Research

Future research should concentrate on strengthening the security and practical deployment of the federated intrusion-detection architecture. First, secure aggregation and stronger protection against malicious client updates should be investigated. The dissertation specifically suggests approaches including secure aggregation and homomorphic encryption. Second, alternative aggregation approaches such as FedProx, FedNova, and FedWeight should be evaluated to better accommodate non-IID IoT data. Third, future work should investigate vertical and distributed Federated Learning rather than relying exclusively on horizontal FL and FedAvg. Finally, improved sampling, semi-supervised approaches, and more sophisticated deep-learning architectures could be investigated to improve the detection of underrepresented attack categories.

Conclusion

This study examined Federated Learning as a privacy-preserving architecture for intrusion detection in distributed IoT environments. Unlike conventional centralized approaches that require raw network data to be transferred to a central server, the proposed architecture keeps training data at participating IoT devices and communicates model parameters for collaborative learning. The system combines Random Forest with horizontal Federated Learning and FedAvg. A model-security validation stage was incorporated before trusted local parameters were aggregated into the global model. This design directly addresses two important requirements of IoT security: protection of sensitive device-level data and preservation of the integrity of collaborative model training. The experimental results demonstrate strong detection capability. The global federated model achieved 99.7% binary accuracy and a 0.90 F1-score, while the participating IoT devices achieved accuracies between 99.3% and 99.9%. At the same time, the multi-class macro-F1 of 0.69 demonstrates that the system still faces challenges in identifying some minority attack categories. Finally, the findings support the use of Federated Learning as a promising approach for privacy-preserving and scalable IoT intrusion detection, while highlighting the need for stronger defenses against poisoning attacks and improved approaches for heterogeneous and imbalanced IoT data.

Acknowledgement

The work is self-sponsored by the corresponding author under the supervision of the co-author.

Author contributions

Corresponding author wrote the primary draft of the manuscript, carried out the tests, gathered and examined the data, and conducted the study. Co-author read and edited the manuscript, oversaw the research, and offered advice on technique and interpretation, among other things. The final draft of the work has been reviewed and approved by all authors.

Data Availability

The study's data came from an open-source repository, Edge-IIoTset. The author can provide the datasets created and examined for this study upon reasonable request.

Funding

No funding.

Conflict of interest

The authors declare that there is no conflict of interest, financial or otherwise, related to the publication of this research paper.

Ethics approval

N/A.

AI tool usage declaration

The authors declare that No AI-assisted tools were used and No AI system was used to generate original scientific claims, data, or research results.

References

Aashmi, R., & Jaya, T. (2022). Intrusion detection using federated learning for computing.

Almalki, L., Alnahdi, A., & Albalawi, T. (2023). Role-driven clustering of stakeholders: A study of IoT security improvement. *Sensors, 23*, 5578.

Bhatnagar, A. (2024). Network intrusion detection system using federated learning. *GLIMPSE—Journal of Computer Science, 3*(1), 4–9.

Bhavsar, M., Roy, K., Kelly, J., & Olusola, O. (2023). Anomaly-based intrusion detection system for IoT application. *Discover Internet of Things, 3*, 1–23. https://doi.org/10.1007/s43926-023-00034-5.

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. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1175–1191).

Chaurasia, N., Ram, M., Verma, P., Mehta, N., & Bharot, N. (2024). A federated learning approach to network intrusion detection using residual networks in industrial IoT networks. *The Journal of Supercomputing, 80*, 18326–18346. https://doi.org/10.1007/s11227-024-06153-2.

Dash, N., Chakravarty, S., & Rath, A. K. (2024). Deep learning model for elevating Internet of Things intrusion detection. *International Journal of Electrical and Computer Engineering, 14*(5), 5874–5883. https://doi.org/10.11591/ijece.v14i5.pp5874-5883.

Emmanni, P. S. (2021). Federated learning for cybersecurity in edge and cloud computing. *International Journal of Computing and Engineering*, 14–25.

Feng, Z. (2024). Federated learning security threats and defense approaches. *Highlights in Science, Engineering and Technology, 85*, 121–127.

Ferrag, M. A., Friha, O., Hamouda, D., Maglara, L., & Janicke, H. (2022). *Edge-IIoTset: A new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning*. TechRxiv. https://doi.org/10.36227/techrxiv.18857336.v1.

Marfo, W., Tosh, D. K., & Moore, S. V. (2022). Network anomaly detection using federated learning. In *2022 IEEE Military Communications Conference (MILCOM)*.

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.

Rashid, M. M., Khan, S. U., Eusufzai, F., Redwan, M. A., Sabuj, S. R., & Elsharief, M. (2023). A federated learning-based approach for improving intrusion detection in industrial Internet of Things networks.

Silva, P. R., Vinagre, J., & Gama, J. (2023). Towards federated learning: An overview of methods and applications. WIREs Data Mining and Knowledge Discovery, 13(4), e1486. https://doi.org/10.1002/widm.1486

Singh, V. K., Tucker, E., & Rath, S. (2024). Federated machine learning-based anomaly detection system for synchrophasor network using heterogeneous data sets. In *2024 IEEE/PES Transmission and Distribution Conference and Exposition (T&D)* (pp. 1–5). IEEE. https://doi.org/10.1109/TD47997.2024.10556259.

Tahri, R., Balouki, Y., Jarrar, A., & Lasbahani, A. (2022). Intrusion detection system using machine learning algorithms. *ITM Web of Conferences, 46*, 02003. https://doi.org/10.1051/itmconf/20224602003.

Weinger, B., Kim, J., Sim, A., Nakashima, M., Moustafa, N., & Wu, K. (2022). Enhancing IoT anomaly detection performance for federated learning. *Digital Communications and Networks, 8*(3), 314–323.

Zhang, C., Xie, Y., Bai, H., Yu, B., Li, W., & Gao, Y. (2021). A survey on federated learning. Knowledge-Based Systems, 216, 106775. https://doi.org/10.1016/j.knosys.2021.106775