International Journal of Artificial Intelligence, Data Science, and Machine Learning
Grace Horizon Publication | Volume 6, Issue 4, 87-90, 2025
ISSN: 3050-9262 | https://doi.org/10.63282/3050-9262.IJAIDSML-V6I4P112

Original Article

Use of Federated Learning in Iiot (Industrial Internet of Things) Security

Bharathram Nagaiah
Independent Researcher, USA.

Received On: 13/09/2025 Revised On: 17/10/2025 Accepted On: 24/10/2025 Published On: 13/11/2025

Abstract *- Federated learning (FL) is a distributed machine learning paradigm that enables devices or edge nodes to collaboratively train a global model without sharing raw data. This enhances privacy and reduces communication overhead. As Industrial Internet of Things (IloT) systems expand across manufacturing, energy, and logistics, security has become critical. Centralized learning approaches struggle with privacy risks, latency, and single points of failure. This article explores the application of FL to strengthen cyber-physical system (CPS) security in IIoT. We present a framework for threat detection, anomaly recognition, and intrusion prevention that utilizes lightweight models, secure aggregation, and differential privacy for resource-constrained industrial devices. Experiments on a simulated smart-factory testbed show that our approach achieves accuracy comparable to centralized methods while maintaining data locality and resilience. We also analyze trade-offs, challenges, and future directions for deploying FL in real-global lloT environments.*

Keywords *- Federated Learning (FL), Industrial Internet of Things (IIoT), Cyber-Physical Systems (CPS), Intrusion Detection Systems (IDS), Privacy-Preserving Machine Learning.*

1. Introduction

Industrial Internet of Things (IloT) is a convergence of industrial systems that control advanced sensing, communication, and computation systems. Programmable logic controllers, sensors, and actuators, as well as vision systems, are now being utilized in factories, energy plants, and logistics networks to support intelligent manufacturing and predictive maintenance, and to perform real-time monitoring. Although these technologies cause efficiency and automation, they increase the attack surface, making lloT systems vulnerable to malware, data manipulation, denial-of-service attacks, and unauthorized access. [1] Machine learning has proven a viable solution to improving IloT security, especially intrusion detection systems and anomaly detection models, which detect differences in traffic or device behaviour. Yet, most of these systems rely on centralized training, which requires aggregating sensitive data from distributed endpoints. In the industry setting, this method presents a number of challenges. Proprietary information is often sensitive data of operation that companies are not willing to share. Delays of many devices can also be intolerable, and sending raw data to a central point consumes bandwidth. Single points of failure are generated by central servers and are avenues of attack by enemies. Moreover, stringent regulatory systems in areas like energy and defense often limit the way industrial data is transmitted. [2]

Federated learning (FL), first introduced by Google in 2016, offers an alternative that addresses these challenges. Rather than merging raw data, devices learn models locally and transmit model changes only to an aggregator, where they are privately aggregated. This is a natural fit with the distributed nature of IloT and minimizes privacy vulnerabilities, enhances efficiency, and heightens attack resilience. [3] This paper discusses the use of federated learning in the context of IloT security, considering how the technology can enhance how threats, anomalies, and intrusions are detected. We develop a federated system that is well-adjusted to the limitations and heterogeneity of industrial equipment, combine secure aggregation and differential privacy to protect against adversarial manipulation and inference, and experiment with the architecture in a simulated smart-factory environment. These results demonstrate the accuracy, efficiency in communication, and resilience of the proposed approach, but also illuminate the trade-offs and real-global challenges that need to be overcome to become practical and implementable in real-global industrial applications. [4]

2. Literature Review

Federated learning has rapidly evolved since its formalization through the Federated Averaging algorithm, which allowed distributed clients to perform local training and submit only model updates for centralized aggregation. This method tackled core issues concerning privacy protection and communication expenses that had historically prevented large-scale collaborative machine learning efforts. [5] Initial investigations in this domain concentrated primarily on enhancing communication efficiency, resulting in strategies including model size reduction through pruning and quantization techniques, gradient compression methods, and streamlined communication protocols. Parallel to these advancements, significant attention was devoted to privacy and security considerations. The solutions to the threats, incorporating protective mechanisms such as differential privacy mechanisms, secure multi-party computation protocols, and homomorphic encryption systems, were presented [6]. A significant technical challenge was dealing with the diversity of data, namely, in dealing with heterogeneous and non-independently distributed data among different clients involved. Solutions, including enhanced optimization algorithms tailored for federated environments, clustered federated learning methodologies, and personalized model strategies, were created to accommodate varying data distributions, different computational resources, and inconsistent client participation throughout the federated network [7]. Together, these advancements demonstrate how federated learning has transitioned from a concept rooted in mobile devices to a versatile framework applicable across diverse domains.

Security of the Industrial Internet of Things environment has been among the burning issues since the merging of operational technology and information technology in highly sensitive environments [8]. The IIoT systems integrate real-time control needs, safety-critical operations, and various industrial communication protocols, such as Modbus, Profinet, and OPC UA [9]. The classic security tools have generally been signature-based intrusion detection systems, which work well on familiar threats but provide minimal defense against new or zero-day attacks [10]. In order to address these shortcomings, semi-supervised and unsupervised anomaly detection techniques have been proposed to identify deviations in traffic or system behavior that can indicate malicious intent.

More recently, intrusion detection based on machine learning has become a promising defense strategy, with support vector machines, random forests, and deep neural networks being examples of such classifiers used to detect attack trends on historical data [11]. Although socially successful, the vast majority of machine learning methods rely on centralized aggregation of data, thus presenting a risk to data privacy, communication overheads, and exposing the system to centralized failure. Such deficiencies are especially acute in the industrial context, where the sensitivity of working information, on the one hand, and the limitations of real-time communications, on the other hand, have a harsh effect on the structure of the system.

With the understanding of these issues, scientists have started to investigate the implementation of federated learning in the field of IIoT security. One of the proposed frameworks shows how federated learning can ensure data locality in addition to facilitating collective defense across distributed systems [12]. Applications in smart manufacturing show that the approach improves scalability and maintains privacy when compared to centralized models [13]. Other efforts in the automotive and energy domains demonstrate how federated intrusion detection systems can be enhanced with secure aggregation to protect data during training [14]. Although these studies are rather strong pieces of evidence concerning the possibility of using federated learning to improve the security of IIoT and cyber-physical systems, most of them are limited in scope [15]. Not many publications consider all challenges of resource limitations, communication expenses, resistance of federated learning as a whole to adversarial attacks, and tight real-time constraints of the industrial context [16].

This accumulating literature posits that federated learning is a privacy-preserving and efficient solution to IIoT security when compared to centralized learning [17]. But it also throws up serious shortcomings in aspects of actual application and overall assessment. To fill these gaps, special attention should be paid to the technical aspects of federated optimization and privacy mechanisms, as well as the practicalities of the industrial environment, such as the heterogeneity of the capabilities of various devices and networks, their reliability, and regulatory disadvantages.

3. Methodology

Our framework consists of the following key elements:

3.1. System Architecture

- **IIoT Edge Nodes**: Basic computing and networking units (sensors, PLCs, gateways). Each node maintains its own statistics (e.g., packet traces, operational metrics) and runs a lightweight anomaly detection model, such as a shallow CNN or RNN.
- **Cloud/Fog Aggregator**: Coordinates the federated learning process. This can be a central server or a distributed fog aggregator. Its responsibilities include selecting participants, broadcasting the global model, receiving local updates, securely aggregating them, and redistributing the improved model.
- **Security Modules**: Protect individual updates using methods like additive secret sharing.
- **Differential Privacy**: Each node injects controlled noise into its updates to prevent sensitive information leakage.
- **Byzantine-Resilient Aggregation**: Robust algorithms such as *Krum* or *trimmed mean* safeguard against malicious or corrupted updates.

3.2. Dataset and Simulation

We simulated a smart manufacturing setup using **Modbus/TCP traffic**. The dataset included both:

- Normal operational data (e.g., status updates and control commands)
- Injected attacks (spoofing, replay, and command injection)

To mimic realistic non-IID conditions, data was unevenly distributed across nodes—some handled heavy traffic while others processed very little.

3.3. Model Design and Training

- **Local Model**: Each edge node trained a small RNN (~100 KB) to classify data time-series as *normal* or *anomalous*, optimized for constrained memory.

Training Rounds:

- Model initialized at the aggregator
- Local training for a few epochs on each node's data
- Noise and masking applied for privacy and security
- Secure upload of gradients
- Aggregation and update at the server
- Updated model broadcast back to nodes

3.4. Evaluation Metrics

We evaluated performance across four dimensions:

- **Detection performance**: accuracy, precision, recall, F1
- **Communication overhead**: bandwidth per round, convergence speed
- **Privacy and robustness**: resistance to poisoning, inversion
- **Resource usage on edge devices**: memory, CPU, latency

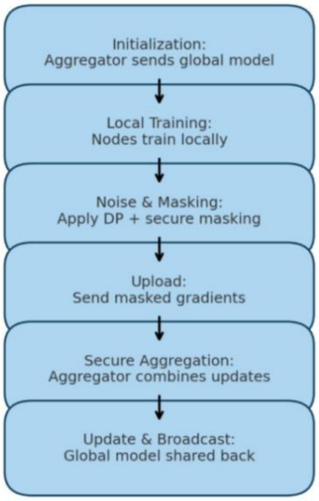

Figure 1: Proposed Methodology

4. Results

4.1. Detection Performance

The federated model achieved 94.2% accuracy, compared to 95.5% for a centralized approach—a negligible drop. Precision (92.8%), recall (93.7%), and F1-score (93.2%) were also very close to centralized benchmarks.

4.2. Communication Efficiency

Each node transmitted about 200 KB per round (model updates only), compared to several MB in centralized data uploads. Federated learning converged in ~20 rounds, which is slower than 10 centralized epochs but lighter per round.

4.3. Privacy and Security

Adding Gaussian noise (σ=0.1) kept accuracy above 93%, showing a strong privacy–utility balance. Secure aggregation prevented the server from seeing individual updates. Under simulated Byzantine attacks (10% of nodes sending poisoned data), robust aggregation limited accuracy loss to 3%, while naive averaging lost more than 15%.

4.4. Resource Utilization

The model required ~120 KB of memory, with privacy mechanisms adding about 30 KB. This fits within the memory range of typical gateways (512 KB–1 MB). Local training took ~200 ms per round, with communication adding ~150 ms—fast enough for batch updates but not for strict real-time detection.

5. Discussion

5.1. Benefits and Trade-offs

Federated learning offers several benefits:

- **Privacy protection**: raw data never leaves the device
- **Reduced bandwidth usage**: only model updates are transmitted
- **Improved fault tolerance**: no single point of failure
- **Regulatory compliance**: aligns with industrial confidentiality requirements

However, there are trade-offs:

- **Slight accuracy gap** compared to centralized training
- **Complexity of privacy mechanisms**, such as differential privacy and secure aggregation
- **Higher communication cost**, as multiple cycles may delay updates
- **Non-IID data challenges**, which may bias models—though personalization can mitigate this

5.2. Deployment Challenges

- **Device diversity**: weaker IIoT nodes may have limited computational capacity, requiring adaptive model designs and training schedules
- **Network reliability**: intermittent connectivity must be handled gracefully, allowing for lagging or offline nodes
- **Aggregator security**: the central server is a sensitive target; mitigation can involve multiple or distributed aggregators

- **Real-time requirements**: in some industrial use cases, urgent alerts must be handled locally without waiting for global model updates

5.3. Future Directions

- **Personalized federated models** tailored to local device patterns
- **Hierarchical federated learning** with multiple layers of aggregators
- **Lightweight encryption** methods to secure updates with minimal overhead
- **Adaptive communication**, where updates are sent only when necessary
- **Deployment trials** in real-world industrial plants and factories

6. Conclusion

Federated learning shows strong potential for securing IIoT systems. In our experiments, it achieved accuracy levels close to centralized models while consuming significantly less bandwidth and offering stronger resistance to adversarial attacks. The resource footprint was small enough for practical deployment on edge devices. For mainstream adoption, several challenges must be addressed: supporting heterogeneous device capabilities, ensuring robust aggregation security, and managing deployment complexity. Future work should focus on refining personalized models, developing top-down hierarchical architectures, and testing in live industrial environments.

References

[1] Wikipedia contributors. (2025, August). *Industrial internet of things*. In *Wikipedia*. Retrieved September 7, 2025, from https://en.wikipedia.org/wiki/Industrial_internet_of_things

[2] ArXiv contributors. (2023, June 5). *Federated deep learning for intrusion detection in IoT networks* [Preprint]. arXiv. https://arxiv.org/abs/2306.02715

[3] Wikipedia contributors. (2025, August). *Federated learning*. In *Wikipedia*. Retrieved September 7, 2025, from https://en.wikipedia.org/wiki/Federated_learning

[4] Potharaju, R., & Farooq, M. (2022). Securing Internet of Things devices with federated learning: A privacy-preserving approach for distributed intrusion detection. *CMC—Computers, Materials & Continua*, 83(3), 893–908. https://www.techscience.com/cmc/v83n3/61057/html

[5] McMahan, B., Moore, E., Ramage, D., Hampson, S., & Arcas, B. A. (2017). *Communication-Efficient Learning of Deep Networks from Decentralized Data*. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS), PMLR 54, 1273–1282. arXiv preprint arXiv:1602.05629

[6] Konečný, J., McMahan, H. B., Yu, F. X., Richtárik, P., Suresh, A. T., & Bacon, D. (2016). *Federated Learning: Strategies for Improving Communication Efficiency*. arXiv preprint arXiv:1610.05492. arXiv

[7] Sattler, F., Müller, K.-R., & Samek, W. (2019). *Clustered federated learning: Model-agnostic distributed multi-task optimization under privacy constraints. arXiv preprint arXiv:1910.01991*

[8] “Operational Technology,” Wikipedia. Retrieved September 7, 2025, from https://en.wikipedia.org/wiki/Operational_technology

[9] Duque Anton, S., Kanoor, S., Fraunholz, D., & Schotten, H. D. (2019). *Evaluation of machine learning-based anomaly detection algorithms on an industrial Modbus/TCP data set*. arXiv. https://arxiv.org/abs/1905.11757

[10] “Intrusion Detection System,” Wikipedia. Retrieved September 7, 2025, from https://en.wikipedia.org/wiki/Intrusion_detection_system

[11] Duque Anton, S., Sinha, S., & Schotten, H. D. (2019). *Anomaly-based intrusion detection in industrial data with SVM and random forests*. arXiv. https://arxiv.org/abs/1907.10374

[12] Ali, A., Husain, M., & Hans, P. (2025, May 21). *Federated Learning-Enhanced Blockchain Framework for Privacy-Preserving Intrusion Detection in Industrial IoT* [Preprint]. arXiv. Retrieved from https://arxiv.org/abs/2505.15376 arXiv

[13] Evaluation of federated intrusion detection in smart manufacturing industries: *FLDID: Federated Learning Enabled Deep Intrusion Detection in Smart Manufacturing Industries*. PubMed. Retrieved from https://pubmed.ncbi.nlm.nih.gov/36433569/ PubMed

[14] Chatterjee, S., & Hanawal, M. K. (2021, June 25). *Federated Learning for Intrusion Detection in IoT Security: A Hybrid Ensemble Approach* [Preprint]. arXiv. Retrieved from https://arxiv.org/abs/2106.15349 arXiv

[15] Sarhan, M., Lo, W. W., Layeghy, S., & Portmann, M. (2022, April 8). *HBFL: A Hierarchical Blockchain-based Federated Learning Framework for a Collaborative IoT Intrusion Detection* [Preprint]. arXiv. Retrieved from https://arxiv.org/abs/2204.04254

[16] *Federated learning in intrusion detection: advancements, applications, and future directions*. Cluster Computing. Retrieved from https://link.springer.com/article/10.1007/s10586-025-05325-w

[17] *Federated Learning-Based Intrusion Detection in IoT Networks: Performance Evaluation and Data Scaling Study*. MDPI. Retrieved from https://www.mdpi.com/2224-2708/14/4/78.