International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Intrusion Detection in Industrial IoT Gateway Networks: AComparative Study of Centralized Machine Learning and Federated Learning

Anav Sobti, Kinshi Sinha[1], Ayushi Mishra[2]

[1,2]Department of Electronics and Communication Engineering, Netaji Subhas University of Technology, New Delhi, India

Abstract

Due to device heterogeneity, massive data generation, and stringent privacy regulations, the rapid growth of Industrial Internet of Things networks has increased security challenges. Although centralized machine learning based intrusion detection systems provide high detection accuracy, their reliance on centralized data aggregation poses privacy risks and limits scalability. Federated Learning addresses these issues by enabling collaborative model training without exchanging raw data. Using the X-IIoTID dataset, this paper compares centralized machine learning and Federated Learning for intrusion detection in Industrial IoT gateway networks. Under IID client partitions, several centralized machine learning models and a Federated Learning framework based on the FedAvg algorithm with a lightweight neural network are evaluated. Experimental results show that Federated Learning achieves 98.38% accuracy, while centralized machine learning achieves up to 99.85%. Federated Learning therefore offers competitive performance together with better privacy, lower communication overhead, and stronger robustness.

Keywords: Industrial Internet of Things, Intrusion Detection System, Federated Learning, Machine Learning, Privacy-Preserving Security

1. Introduction

The Industrial Internet of Things has become a cornerstone of modern industrial automation, enabling intelligent monitoring, predictive maintenance, and real time decision making across critical infrastructures. An Industrial IoT environment is composed of interconnected components such as sensors, controllers, and gateways that continuously generate large volumes of heterogeneous data. This connectivity improves operational efficiency, but it also enlarges the attack surface and makes systems more susceptible to cyber threats.

An intrusion detection system is an important mechanism for securing Industrial IoT infrastructures because it monitors network and system activities for unauthorized access or malicious behavior. Traditional rule based intrusion detection approaches are no longer sufficient against sophisticated and evolving cyber attacks. In this context, machine learning based intrusion detection has gained significant attention because it can learn complex traffic patterns automatically and detect both known and previously unseen attacks with high accuracy.

Several research works have demonstrated the effectiveness of centralized machine learning based intrusion detection in IoT and Industrial IoT settings. Sharma et al. reported strong IoT attack detection performance using traditional machine learning models, especially ensemble classifiers [5]. Likewise, Imran et al. emphasized the value of real time feature engineering in centralized intrusion detection for MQTT based IoT networks and obtained effective results using optimized decision tree models [3]. Although these works confirm the detection capability of centralized machine learning, they rely on centralized processing of raw traffic statistics gathered from multiple gateways. This design introduces concerns related to data privacy, communication latency, scalability, and stability in the event of failures.

To address these limitations, Federated Learning was introduced to support collaborative learning across multiple parties without requiring them to share raw data. Privacy protection is therefore improved and communication requirements are reduced. CYBRIA was introduced in [1] as an example of privacy aware cybersecurity enabled by federated learning. The practical implementation of such learning was explored further in [2], where an asynchronous transmission and update strategy was studied to improve efficiency with minimal loss of accuracy.

For realistic evaluation of intrusion detection methods in Industrial IoT networks, Al Hawawreh et al. proposed the X-IIoTID benchmark dataset, which is connectivity agnostic and device agnostic and was designed specifically for Industrial IoT intrusion detection [4]. Using this benchmark, Singh and Gupta studied the effect of local data imbalance on federated learning based anomaly detection in IoT networks and concluded that the absence of IID client data can significantly reduce federated learning performance [6]. However, that study did not provide a direct comparison between centralized machine learning and federated learning on a common Industrial IoT benchmark.

To address this research gap, the present work provides a comparative study of centralized machine learning and federated learning for intrusion detection in Industrial IoT gateway networks using the X-IIoTID dataset. Several centralized machine learning models are evaluated and compared with a federated learning framework based on IID client partitioning and the FedAvg aggregation strategy. The analysis focuses on detection performance, privacy, communication cost, scalability, and implementation practicality, thereby providing a realistic benchmark for real world Industrial IoT intrusion detection.

2. Related Work

Machine learning based intrusion detection systems have received significant attention in IoT and Industrial IoT environments because cyber attacks are becoming increasingly sophisticated and traditional rule based security mechanisms have clear limitations. Supervised learning models have shown strong capability in identifying malicious traffic patterns. However, most existing approaches still rely on centralized data aggregation, which raises concerns about privacy, scalability, and communication overhead.

Sharma et al. investigated intrusion detection in IoT devices using traditional machine learning techniques on the N BaIoT dataset, focusing on botnet attacks such as Mirai and Bashlite [5]. Their results showed detection accuracy close to 99 percent using classifiers such as Decision Tree and Random Forest. Despite the strong performance, the study was limited to device specific traffic and did not address the privacy and scalability issues that arise in distributed Industrial IoT systems.

Imran et al. emphasized the significance of real time feature engineering for MQTT based IoT intrusion detection [3]. Decision tree based models achieved accuracy and F1 scores greater than 99 percent after the inclusion of source level features in MQTTset traffic. Nevertheless, the protocol specific nature of the dataset and the centralized learning assumption reduce the general applicability of the study to broad Industrial IoT environments.

To improve realism and generalizability, Al Hawawreh et al. introduced the X-IIoTID dataset as a connectivity agnostic and device agnostic benchmark for Industrial IoT environments [4]. The dataset supports a wide variety of attack scenarios and incorporates network, host, and system level features. Strong centralized machine learning and deep learning results reported on this dataset make it a reliable benchmark for Industrial IoT intrusion detection studies.

Federated Learning has emerged as a decentralized alternative because of the growing concern over data privacy and communication overhead. Thantharate et al. proposed the CYBRIA framework to show that federated intrusion detection can achieve competitive performance without sharing raw data [1]. However, their evaluation was limited to smaller scale settings. Overall, although there is substantial work on centralized machine learning and growing interest in Federated Learning, systematic comparisons under identical Industrial IoT conditions remain limited. In particular, baseline evaluations using IID client distributions are needed to isolate the effect of decentralization. This study fills that gap by providing a unified comparison of centralized machine learning and Federated Learning on the X-IIoTID dataset.

3. Dataset and Problem Formulation

3.1 Dataset Description

In this research, the benchmark dataset X-IIoTID is used for evaluating intrusion detection in Industrial IoT environments. Unlike many existing datasets that are device specific or protocol specific, X-IIoTID is both connectivity agnostic and device agnostic. It provides network traffic features, host level features, and system log information, and it includes diverse Industrial IoT scenarios that contain both normal traffic and attack traffic.

3.2 Problem Definition

The intrusion detection task is formulated as a binary classification problem in which each traffic instance is classified as either benign or malicious. The study evaluates this task using both centralized machine learning and federated learning in order to compare their effectiveness under a common benchmark.

4 Centralized Machine Learning Methodology

4.1 Preprocessing

The dataset undergoes a sequence of preprocessing steps including removal of redundant features, encoding of categorical attributes, Z score normalization, and stratified train test splitting. These steps support fair comparison across models and ensure that each method is trained on standardized input data.

4.2 Machine Learning Models

To examine both linear and non linear properties of Industrial IoT traffic, a range of supervised machine learning models is considered. The evaluated models include Logistic Regression, Decision Tree, K Nearest Neighbors, Support Vector Machine, Naive Bayes, Random Forest, Gradient Boosting, and

XGBoost. Simpler models such as Logistic Regression and Naive Bayes provide baseline performance, while ensemble methods help assess generalization strength and robustness.

4.3 Results

The experiments show that ensemble based approaches outperform single classifiers. XGBoost achieves the best detection capability with accuracy above 99 percent, followed by Random Forest and Gradient Boosting. These findings confirm the effectiveness of centralized ensemble learning for intrusion detection in Industrial IoT gateway networks, while also highlighting the higher data sharing and communication requirements of centralized systems. A comparison of accuracy, precision, recall, and F1 score across the evaluated centralized models is shown in Figure 1.

MODEL	Accuracy	Precision	Recall	F1 Score
XG Boosting	0.9985	0.9982	0.9984	0.9983
Gradient Boost	0.9934	0.9951	0.9896	0.9924
Random Forest	0.9980	0.9982	0.9973	0.9978
Decision Tree	0.9974	0.9964	0.9976	0.9970
KNN	0.9871	0.9842	0.9862	0.9852
SVM	0.9809	0.9965	0.9594	0.9776
Logistic Regression	0.9550	0.9722	0.9230	0.9470
Naive Bayes	0.5684	0.5021	0.9923	0.6668

Figure 1: Performance Evaluation of Centralized Machine Learning Models on the X-IIoTID Dataset

5 Federated Learning Implementation

5.1 Architecture

Federated Learning enables collaborative intrusion detection model training across distributed Industrial IoT gateways without transferring raw data. Each gateway acts as a federated client and performs local training on its private dataset. A central server coordinates the overall training process by collecting model updates from the participating clients and aggregating them into a shared global model. This distributed design reduces privacy risk and communication burden while remaining practical for industrial settings where sensitive data and limited bandwidth are major constraints.

5.2 Client Partitioning

To enable fair comparison with centralized machine learning, the dataset is divided into 10 IID client subsets. Each client receives a balanced share of the overall dataset with similar class distribution. This design separates the effect of decentralization from the effect of data heterogeneity and creates a controlled baseline for evaluating federated learning performance.

5.3 Model Architecture

A lightweight neural network is designed specifically for federated training on resource constrained Industrial IoT gateways. The model consists of an input layer that matches the selected features, one hidden layer with ReLU activation, a dropout layer to reduce overfitting, and an output layer for binary

classification. This compact design minimizes local computation cost and communication overhead during model update exchange.

5.4 Training and Aggregation

Federated training is performed over multiple communication rounds. In each round, the server sends the current global model parameters to all participating clients. Each client performs local training for a fixed number of epochs on its local dataset and returns the updated parameters to the server. The server then combines the client updates using the FedAvg algorithm and redistributes the updated global model for the next round.

$$w(t+1) = \Sigma(k = 1 \text{ to } K) (n_k / n) w_k(t) \quad (1)$$

In Equation 1, $w(t+1)$ represents the updated global model parameters at round t plus 1, $w_k(t)$ denotes the locally trained model parameters from client k at round t, n_k is the number of training samples at client k, n is the total number of samples across all clients, and K is the total number of participating clients.

5.5 Training Configuration and Hyperparameters

The federated learning setup uses 10 federated clients that simulate Industrial IoT gateways. Training is carried out for 20 communication rounds. Each client performs two local epochs per round with a batch size of 64, and the local models are optimized with a learning rate of 5 × 10^-4. This configuration reflects the trade off among convergence speed, communication cost, and computational efficiency and remains practical for real world gateway deployment.

6 Results

The federated global model is evaluated after multiple communication rounds and achieves an accuracy of 98.38 percent, with an F1 score above 0.90. This shows that effective intrusion detection can be achieved without relying on centralized raw data aggregation. The increase in accuracy across rounds also indicates stable convergence under the IID client distribution used in this study.

6.1 Comparative Analysis

6.1.1 Detection Performance

Centralized machine learning produces the highest detection performance because the models are trained with complete access to the full dataset. Ensemble methods such as XGBoost and Gradient Boosting achieve accuracy close to 99.8 percent together with excellent precision, recall, and F1 score. Federated Learning produces slightly lower accuracy at 98.38 percent and an F1 score above 0.90 because training is decentralized and the model architecture is intentionally lightweight. Even so, it remains a viable and competitive alternative for Industrial IoT intrusion detection.

6.1.2 Privacy and Scalability

Data privacy is a major concern in Industrial IoT environments because network and operational data can be highly sensitive. Centralized machine learning requires raw data aggregation and therefore exposes the system to greater privacy risk. Federated Learning protects privacy by keeping data local and sharing only model parameters, which substantially reduces the risk of data exposure.

6.1.3 System Robustness

Centralized machine learning is constrained by server capacity and introduces a single point of failure. Federated Learning distributes computation across multiple gateways, which improves scalability and increases fault tolerance. This makes Federated Learning especially attractive for practical Industrial IoT deployment.

6.2 Quantitative Performance Comparison

The comparison between centralized machine learning and Federated Learning on the X-IIoTID dataset is summarized in Table 1 and Figures 2, 3, and 4.

Table 1: Performance Comparison Between Centralized Machine Learning and Federated Learning

Aspect	Centralized Machine Learning	Federated Learning
Training Architecture	Centralized server with full data access	Distributed clients with server side aggregation
Dataset Availability	Complete dataset at server	Local data at clients with no raw data sharing
Best Model	XGBoost	Lightweight neural network with FedAvg
Accuracy	Approximately 99.8%	Approximately 98.38%
Precision	Very High	High
Recall	Very High	High
F1 Score	Approximately 0.99	Greater than 0.90
Communication Cost	High due to raw data transfer	Low because only model parameters are exchanged
Privacy Preservation	Low	High
Scalability	Limited by server capacity and bandwidth	Highly scalable
Single Point of Failure	Yes	No
Suitability for IIoT Deployment	Limited	High

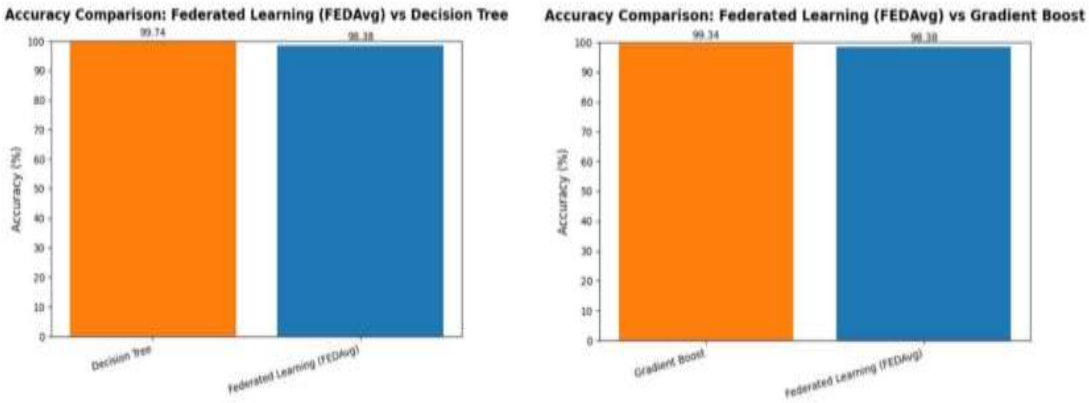

Figure 2: Comparison of Federated Learning with Decision Tree and Gradient Boost Models

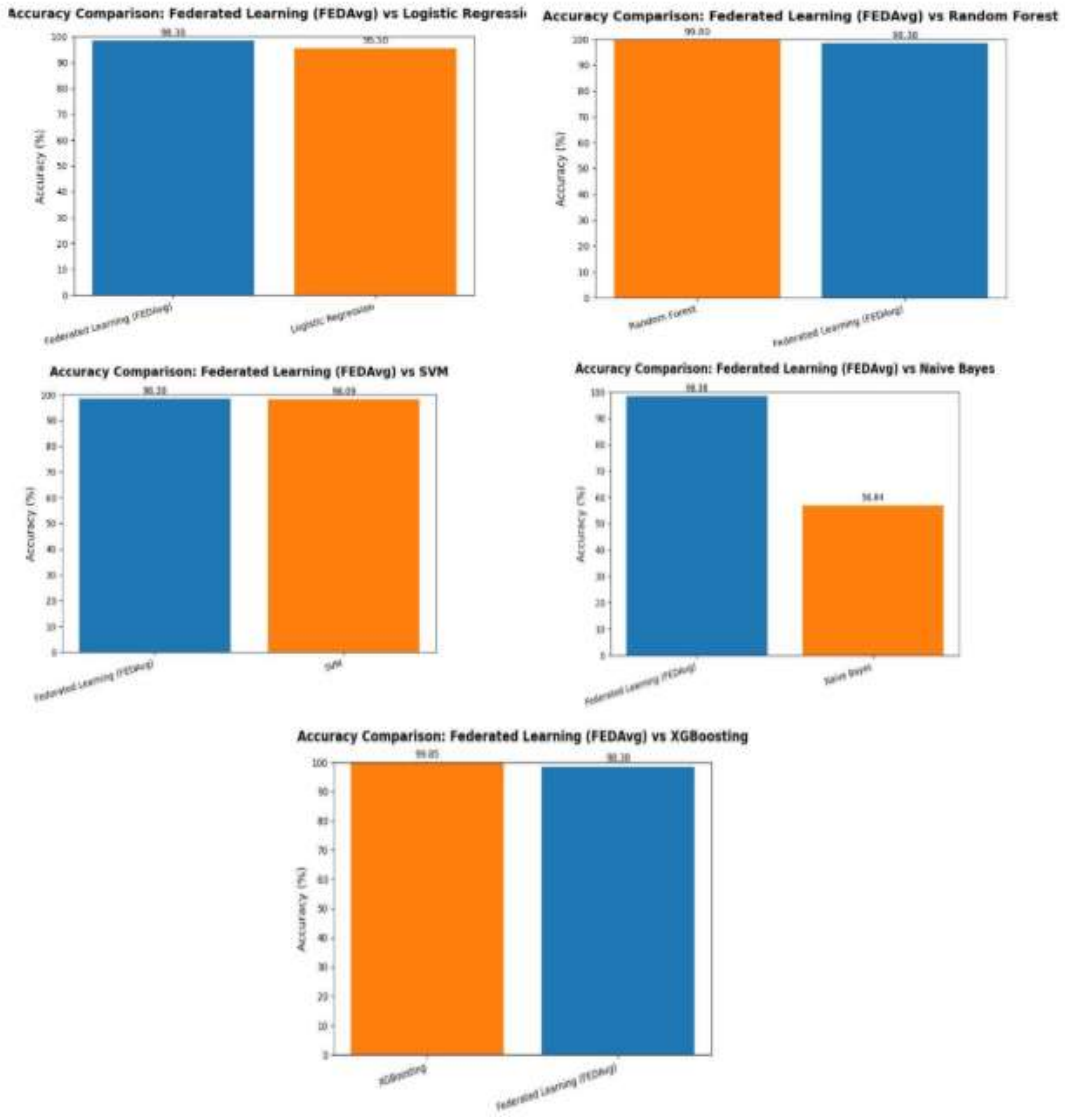

Figure 3: Comparison of Federated Learning with Logistic Regression, Random Forest, SVM, Naive Bayes, and XGBoost

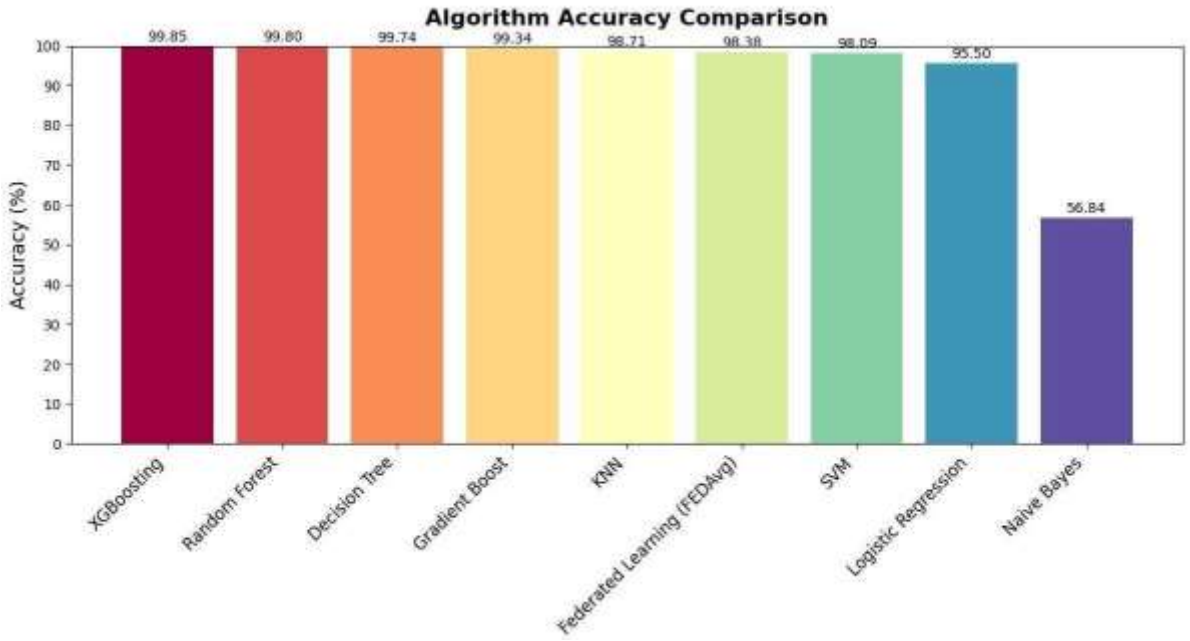

Figure 4: Overall Accuracy Comparison of Centralized Machine Learning Models and the Federated Global Model on the X-IIoTID Dataset

7. Conclusion and Future Work

The experimental results show that centralized machine learning achieves the highest detection performance because it benefits from full data visibility and can exploit powerful ensemble models such as XGBoost. These models capture complex attack patterns in the X-IIoTID dataset and achieve near perfect classification accuracy.

This performance advantage comes at the cost of high communication overhead and serious privacy concerns, because raw Industrial IoT traffic data must be transmitted to and stored at a central location. In industrial settings, where bandwidth may be limited and data security is critical, these constraints can be substantial.

Federated Learning provides slightly lower detection accuracy, but it offers clear advantages in privacy preservation, communication efficiency, scalability, and deployment practicality. Because data remains close to the gateway and only model parameters are exchanged, the federated approach aligns more naturally with the operational constraints of real world Industrial IoT systems.

Overall, the comparison shows that centralized machine learning remains a strong performance baseline, while Federated Learning offers a more balanced and deployable solution for intrusion detection in Industrial IoT gateway networks.

Future work can extend this study to multi class intrusion detection by using the hierarchical labeling structure of the X-IIoTID dataset. Such an extension would allow the framework not only to detect malicious traffic but also to identify the specific type of attack. Finer grained classification would support more context aware response mechanisms and help industrial security systems trigger more suitable mitigation strategies when an intrusion is detected.

8. Challenges and Discussion

The main challenges observed during federated learning implementation include slower convergence, modest reduction in accuracy because of lightweight models, communication delays across rounds, and

increased debugging complexity. Even so, these challenges reflect the practical issues that must be addressed for real world industrial deployment, and they do not diminish the promise of Federated Learning as a privacy conscious intrusion detection strategy.

References

1. P. Thantharate, T. Anurag, "CYBRIA - Pioneering Federated Learning for Privacy-Aware Cybersecurity with Brilliance," 2023 IEEE 20th International Conference on Smart Communities: Improving Quality of Life Using AI, Robotics and IoT (HONET), Boca Raton, FL, USA, 2023, 56-61. https://doi.org/10.1109/HONET59747.2023.10374608
2. Y. Jia, N. Zhang, "Research and Implementation of Asynchronous Transmission and Update Strategy for Federated Learning," 2022 IEEE 8th International Conference on Computer and Communications (ICCC), Chengdu, China, 2022, 1281-1286.
3. S. I. Ali et al., "Realtime Feature Engineering for Anomaly Detection in IoT Based MQTT Networks," IEEE Access, 2024, 12, 25700-25718. https://doi.org/10.1109/ACCESS.2024.3363889
4. M. Al-Hawawreh, E. Sitnikova, N. Aboutorab, "X-IIoTID: A Connectivity-Agnostic and Device-Agnostic Intrusion Data Set for Industrial Internet of Things," IEEE Internet of Things Journal, 2022, 9 (5), 3962-3977. https://doi.org/10.1109/JIOT.2021.3102056
5. Y. Sharma, V. Kumar, H. Chaudhary, "Attack Detection on Internet of Things Devices Using Machine Learning Techniques," Proceedings of the 7th International Conference on Intelligent Computing and Control Systems, 2023, 281-287.
6. J. Singh, S. Gupta, "Evaluating the Impact of Local Data Imbalance on Federated Learning Performance for IoT Anomaly Detection," Proceedings of the 14th International Conference on Computing, Communication and Networking Technologies, Delhi, India, July 2023, 1-7.