International IT Journal of Research (IITJR), ISSN: 3007-6706
Volume 3, Issue 3, July-September, 2025
Available online at: https://itjournal.org/index.php

Federated Learning Frameworks for Privacy-Preserving Intrusion Detection in Industrial IoT Networks

Dr. Wei-Chieh Lin

Associate Professor, Department of Computer Science and Information Engineering,
National Tsing Hua University, Taiwan

ABSTRACT

The rapid adoption of Industrial Internet of Things (IIoT) technologies has transformed modern industrial environments by enabling intelligent automation, real-time monitoring, interconnected devices, and data-driven decision-making. However, the increasing connectivity of industrial systems has significantly expanded their cybersecurity attack surfaces, exposing critical infrastructures to sophisticated cyber threats, including distributed denial-of-service (DDoS) attacks, malware, unauthorized access, data manipulation, and network intrusions. Traditional centralized intrusion detection systems (IDS) often require transferring large volumes of sensitive operational data to centralized servers, creating privacy risks, communication overhead, and potential single points of failure. To address these challenges, this study explores the application of Federated Learning (FL) frameworks for developing privacy-preserving intrusion detection systems in Industrial IoT networks. Federated Learning enables multiple industrial devices, gateways, and edge computing nodes to collaboratively train machine learning models without directly sharing their raw data. The proposed framework integrates distributed model training, secure parameter aggregation, privacy-enhancing mechanisms, and intelligent anomaly detection to improve cybersecurity while maintaining data confidentiality.

Different FL architectures, including centralized federated aggregation, decentralized federated learning, and hierarchical federated learning, are examined in terms of detection accuracy, communication efficiency, scalability, computational requirements, and privacy protection. Furthermore, advanced techniques such as differential privacy, secure multiparty computation, blockchain-assisted aggregation, and adaptive learning are considered to strengthen the reliability and robustness of intrusion detection mechanisms. The study also investigates challenges associated with non-independent and identically distributed (Non-IID) industrial data, heterogeneous devices, communication constraints, adversarial model updates, and dynamic cyberattack patterns. A comparative analytical approach is proposed to evaluate the effectiveness of different federated learning strategies under diverse industrial network conditions. The findings are expected to demonstrate that Federated Learning can provide a scalable, collaborative, and privacy-aware approach to intrusion detection while reducing dependence on centralized data collection. This research contributes to the development of secure, intelligent, and resilient Industrial IoT infrastructures by establishing a framework for balancing cybersecurity performance, data privacy, and operational efficiency.

Keywords: Federated Learning, Industrial Internet of Things (IIoT), Privacy-Preserving Intrusion Detection, Cybersecurity, Distributed Machine Learning.

INTRODUCTION

The Industrial Internet of Things (IIoT) has emerged as a transformative technology in modern industrial environments, integrating intelligent sensors, connected machines, embedded systems, communication networks, and cloud-based platforms to facilitate automated operations and real-time data exchange. The integration of IIoT technologies into manufacturing, energy production, transportation, healthcare, and critical infrastructure has significantly improved operational efficiency, productivity, system monitoring, and predictive maintenance. Through interconnected devices and intelligent data-processing mechanisms, industrial organizations can achieve better resource utilization, improved decision-making, and enhanced operational flexibility. However, the growing interconnectivity of industrial systems has also introduced significant cybersecurity challenges, making industrial networks increasingly vulnerable to sophisticated cyberattacks.

Industrial IoT networks generate and process enormous volumes of sensitive operational information, including machine configurations, production parameters, equipment status, communication records, and process-control data. Unauthorized access to such information can result in financial losses, operational disruptions, intellectual property theft, and potential threats to public safety. Common cyber threats targeting industrial environments include distributed

denial-of-service (DDoS) attacks, ransomware, malware infections, spoofing, unauthorized device access, man-in-the-middle attacks, and data manipulation. Traditional security mechanisms, including signature-based detection, conventional firewalls, and centralized intrusion detection systems, often struggle to identify previously unknown attacks and sophisticated threats that continuously evolve. Consequently, intelligent and adaptive intrusion detection mechanisms are essential for ensuring the confidentiality, integrity, and availability of industrial communication networks.

PROPOSED MODELS AND METHODOLOGIES

2.1 Overview of the Proposed Framework

The proposed research develops a Federated Learning (FL)-based privacy-preserving intrusion detection framework for Industrial Internet of Things (IIoT) networks. The primary objective is to establish a distributed cybersecurity architecture that enables multiple industrial devices, edge gateways, and interconnected organizations to collaboratively train intelligent intrusion detection models without sharing their original network traffic datasets.

The proposed framework combines Federated Learning, machine learning, deep learning, edge computing, and privacy-enhancing security mechanisms. It is designed to identify malicious network activities, detect abnormal communication patterns, and improve cybersecurity performance across heterogeneous industrial environments.

The framework consists of six major operational components:

1. Industrial IoT data collection and preprocessing.
2. Distributed local model training.
3. Federated model aggregation.
4. Privacy-preserving communication and secure aggregation.
5. Intelligent intrusion detection and anomaly classification.
6. Model evaluation and continuous optimization.

The architecture follows a collaborative learning approach in which individual industrial nodes train local models using their respective datasets. Instead of transmitting sensitive operational information, participating nodes share encrypted or otherwise protected model updates with an aggregation server. The aggregated global model is subsequently distributed to participating nodes for further training and improvement.

2.2 Proposed System Architecture

The proposed architecture is organized into four major layers.

1. Industrial IoT Device Layer

This layer consists of industrial sensors, programmable logic controllers (PLCs), supervisory control and data acquisition (SCADA) systems, actuators, smart meters, and industrial machines. These devices continuously generate network traffic and operational information.

2. Edge Computing and Local Processing Layer

Industrial gateways and edge computing nodes collect network traffic, perform preprocessing, extract relevant features, and train local intrusion detection models. Local processing minimizes the need to transfer raw industrial datasets to centralized infrastructure.

3. Federated Learning Coordination Layer

This layer coordinates the federated learning process. It distributes global model parameters, receives local model updates, performs aggregation, and generates an improved global intrusion detection model.

4. Security and Intrusion Detection Layer

This layer applies the trained model to incoming network traffic to identify normal and malicious activities. Additional privacy and security mechanisms are incorporated to protect model updates and reduce the risk of unauthorized access or model manipulation.

2.3 Proposed Federated Learning Models

Different Federated Learning models are considered to identify the most suitable architecture for privacy-preserving intrusion detection in industrial environments.

2.3.1 Centralized Federated Learning Model

Centralized Federated Learning (CFL) uses a central aggregation server to coordinate model training among multiple industrial participants. Each participating device independently trains a local model and sends its model updates to the central server.

The server aggregates the updates using Federated Averaging (FedAvg) and generates a global model.

The aggregation process can be represented as:

$$W_{t+1}=\sum_{k=1}^{K}\frac{n_k}{N}W_{t+1}^{k}$$

Where:

- W_{t+1} = Updated global model parameters.
- W_{t+1}^{k} = Local model parameters of participant k.
- K = Total number of participating devices.
- n_k = Number of training samples at participant k.
- N = Total number of training samples across all participants.

Advantages:

- Simple implementation and management.
- Efficient collaborative model training.
- Reduced requirement for raw-data sharing.
- Centralized monitoring and model coordination.

Limitations:

- Dependence on a central aggregation server.
- Potential communication bottlenecks.
- Vulnerability to server-side attacks and failures.

RESULTS & ANALYSIS

3.1 Overview of Experimental Results

The performance evaluation of Federated Learning (FL) frameworks for privacy-preserving intrusion detection in Industrial Internet of Things (IIoT) networks focuses on identifying the effectiveness of distributed learning architectures in detecting malicious network activities while maintaining data confidentiality. The proposed framework considers centralized federated learning, decentralized federated learning, and hierarchical federated learning architectures integrated with machine learning and deep learning algorithms.

The evaluation considers several important performance indicators, including classification accuracy, precision, recall, F1-score, false positive rate, communication overhead, computational efficiency, and privacy protection. These indicators provide a comprehensive understanding of the strengths and limitations of different federated learning approaches.

The experimental analysis is structured around simulated industrial network environments in which distributed clients independently train local intrusion detection models. Model updates are subsequently aggregated to develop a global model capable of identifying different categories of cyberattacks.

Research note: The numerical values in the following tables are illustrative results for a proposed experimental study, not measurements obtained from an executed experiment. They should be replaced with actual experimental observations before publication.

3.2 Performance Evaluation of Federated Learning Architectures

The first evaluation compares three major federated learning architectures based on their intrusion detection performance.

Federated Learning Architecture	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Centralized FL	96.20	95.80	95.40	95.60
Decentralized FL	94.80	94.20	93.90	94.05
Hierarchical FL	97.40	97.10	96.80	96.95

The illustrative results indicate that hierarchical federated learning achieves the highest overall detection performance, with an accuracy of 97.40% and an F1-score of 96.95%. This improvement can be attributed to its multi-level aggregation structure, which facilitates collaborative learning among industrial devices and regional edge gateways.

Centralized federated learning also demonstrates strong classification performance because of its coordinated global model aggregation. However, its dependence on a central server may introduce communication bottlenecks and infrastructure-related risks.

Decentralized federated learning provides an alternative approach by distributing model coordination across participating nodes. Although its illustrative detection performance is slightly lower, it offers potential advantages in reducing dependence on a central aggregation server.

3.3 Comparative Analysis of Machine Learning and Deep Learning Models

Different machine learning and deep learning algorithms are evaluated to understand their suitability for federated intrusion detection.

Algorithm	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Decision Tree	91.60	90.80	89.90	90.35
Random Forest	94.70	94.10	93.60	93.85
SVM	92.80	92.10	91.40	91.75
CNN	95.60	95.20	94.70	94.95
LSTM	96.10	95.70	95.20	95.45
CNN-LSTM	97.40	97.10	96.80	96.95

Intrusion Detection Accuracy by Algorithm
Illustrative comparison of classification accuracy.
85%88.75%92.5%96.25%100%Decision TreeRandom ForestSVMCNNLSTMCNN-LSTM
The comparison suggests that deep learning models can provide improved detection capabilities when network traffic contains complex relationships and sequential patterns.

The CNN-LSTM hybrid model demonstrates the highest illustrative accuracy of 97.40%. CNN extracts important characteristics from network traffic, while LSTM identifies temporal dependencies associated with abnormal communication patterns.

Random Forest also provides competitive performance and may be appropriate for industrial environments where computational simplicity and interpretability are important considerations.

3.4 Analysis of Privacy-Preserving Mechanisms

Privacy preservation is a fundamental requirement of federated intrusion detection because industrial datasets frequently contain confidential operational information.

The proposed framework considers differential privacy, secure aggregation, encrypted communication, and blockchain-assisted verification.

Privacy Mechanism	Privacy Protection	Computational Overhead	Major Advantage
Differential Privacy	Statistical privacy protection	Moderate	Reduces information leakage
Secure Aggregation	Protects individual model updates	Moderate to High	Confidential model aggregation
Encryption	Protects transmitted information	Moderate	Prevents unauthorized interception
Blockchain	Tamper-evident verification	High	Improves auditability and traceability
Combined Mechanisms	Multiple protection layers	High	Comprehensive security approach

Differential privacy introduces controlled noise into model updates, reducing the possibility of reconstructing sensitive training information. However, excessive noise may negatively affect classification accuracy.

Secure aggregation provides additional protection by preventing the aggregation server from directly observing individual model contributions. Encryption protects communication channels against interception and unauthorized modification.

Blockchain-assisted verification can improve the traceability of federated model updates. Nevertheless, its computational and storage requirements must be considered before deployment in resource-constrained industrial environments.

3.5 Effect of Differential Privacy on Detection Accuracy

The introduction of differential privacy creates a trade-off between privacy protection and model performance. Different privacy configurations may produce different levels of classification accuracy.

Privacy Configuration	Illustrative Accuracy (%)	Privacy Protection Level
No Differential Privacy	97.40	Lower
Low Noise	97.10	Moderate
Medium Noise	96.30	High
High Noise	94.80	Very High

COMPARATIVE ANALYSIS IN TABULAR FORM

4.1 Comparative Analysis of Federated Learning Architectures

Parameters	Centralized FL	Decentralized FL	Hierarchical FL
Architecture	Central aggregation server	Distributed peer-to-peer nodes	Multi-level aggregation
Data Privacy	High with safeguards	High with safeguards	High with safeguards
Detection Performance	High	Relatively High	Very High potential
Communication Efficiency	Moderate	Variable	High potential
Scalability	Moderate	High	Very High
Fault Tolerance	Limited by central server	High potential	Moderate to High
Computational Complexity	Moderate	High	Moderate to High
Implementation	Relatively Easy	Complex	Moderately Complex
Industrial Suitability	Small and medium networks	Distributed networks	Large-scale IIoT networks

Analysis: Hierarchical Federated Learning offers significant potential for large-scale industrial environments because it combines local model training with regional and global aggregation. Centralized FL provides simpler coordination, whereas decentralized FL reduces dependence on a single aggregation server.

4.2 Comparative Analysis of Intrusion Detection Algorithms

Algorithm	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
Decision Tree	91.60	90.80	89.90	90.35
Random Forest	94.70	94.10	93.60	93.85
SVM	92.80	92.10	91.40	91.75
CNN	95.60	95.20	94.70	94.95
LSTM	96.10	95.70	95.20	95.45
CNN-LSTM	97.40	97.10	96.80	96.95

Analysis: CNN-LSTM demonstrates the highest illustrative classification performance. Its combination of spatial feature extraction and temporal sequence learning makes it suitable for identifying complex and evolving industrial cyber threats.

4.3 Comparative Analysis of Privacy-Preserving Techniques

Privacy Technique	Data Confidentiality	Security Level	Computational Overhead	Major Limitation
Differential Privacy	High	High	Moderate	Accuracy degradation
Secure Aggregation	High	High	Moderate to High	Aggregation complexity
Homomorphic Encryption	Very High	High	High	Computational cost
Blockchain Integration	Supports integrity and auditability	Application-dependent	High	Storage and latency
TLS Encryption	Protects data in transit	High	Low to Moderate	Does not protect exposed endpoint data

Analysis: Combining differential privacy, secure aggregation, and encrypted communication can provide multiple layers of protection. Blockchain can additionally support auditability and model-update verification, but it should not be considered a replacement for privacy mechanisms.

4.4 Comparative Analysis of Data Distribution Strategies

Parameters	IID Data	Mild Non-IID Data	Moderate Non-IID Data	Severe Non-IID Data
Data Distribution	Similar across clients	Slightly different	Significantly different	Highly heterogeneous
Model Convergence	Stable	Relatively Stable	Moderate	Difficult
Detection Accuracy (%)	97.40	96.20	94.80	91.90
Aggregation Performance	High	Good	Reduced	Challenging
Training Stability	High	High	Moderate	Low
Recommended Strategy	FedAvg	FedAvg with tuning	FedProx	Personalized FL / SCAFFOLD

Analysis: Non-IID data distribution represents an important challenge in industrial federated learning. Adaptive optimization and personalized learning strategies can help improve global model convergence when industrial clients possess different network traffic characteristics.

4.5 Comparative Analysis of Detection Errors

Algorithm	False Positive Rate (%)	False Negative Rate (%)	Detection Reliability
Decision Tree	5.20	6.10	Moderate
Random Forest	3.80	4.20	High
CNN	3.10	3.50	High
LSTM	2.90	3.20	Very High
CNN-LSTM	2.40	2.80	Very High

Analysis: The CNN-LSTM model demonstrates lower illustrative false positive and false negative rates than the other evaluated algorithms. This is important for industrial environments because false alarms can disrupt operations, while undetected attacks may compromise critical infrastructure.

LIMITATIONS & DRAWBACKS

Although Federated Learning (FL) provides significant opportunities for privacy-preserving intrusion detection in Industrial Internet of Things (IIoT) networks, its practical implementation involves several technical, operational, and security-related limitations. The distributed nature of federated learning introduces challenges associated with communication efficiency, heterogeneous data distributions, computational resources, model security, privacy protection, and system scalability.

The integration of machine learning, deep learning, and privacy-enhancing technologies can improve industrial cybersecurity; however, these approaches do not completely eliminate the risks associated with cyberattacks or information leakage. The major limitations and drawbacks are discussed below.

5.1 Non-Independent and Identically Distributed (Non-IID) Data

One of the major challenges in federated learning is the heterogeneous distribution of industrial network data. Different industrial devices, manufacturing units, and organizations generate network traffic according to their specific operational characteristics.

For example, traffic generated by a smart manufacturing system may differ significantly from traffic generated by an energy management network or industrial control system.

These variations can negatively affect global model convergence and reduce intrusion detection accuracy.

Limitation: Conventional federated optimization algorithms may perform poorly when local datasets have significantly different statistical distributions.

5.2 Communication Overhead

Federated Learning requires repeated communication between participating industrial devices and aggregation servers. Although raw datasets remain locally stored, model parameters and gradients must be exchanged during multiple training rounds.

Large deep learning models, particularly CNN-LSTM architectures, may require substantial communication resources.

Drawbacks:

- Increased network bandwidth requirements.
- Communication delays during model aggregation.
- Additional energy consumption.
- Dependence on reliable network connectivity.
- Reduced suitability for bandwidth-constrained industrial devices.

5.3 Computational Resource Limitations

Industrial IoT environments frequently contain resource-constrained devices with limited processing power, memory, and battery capacity.

Advanced deep learning models require significant computational resources for training and inference. Consequently, implementing complex federated learning algorithms on low-capacity sensors and embedded devices may be challenging.

Limitation: Not all industrial devices possess sufficient computational capabilities to participate effectively in distributed model training.

5.4 Privacy Leakage Through Model Updates

Although Federated Learning avoids directly sharing raw datasets, it does not guarantee complete data privacy. Model gradients and parameter updates may contain information about the original training data.

Attackers may exploit these updates through techniques such as:

- Membership inference attacks.
- Model inversion attacks.
- Gradient reconstruction attacks.
- Property inference attacks.

Limitation: Additional mechanisms, including differential privacy, secure aggregation, and encryption, are required to reduce information leakage risks. These mechanisms may introduce additional computational overhead and affect model performance.

5.5 Vulnerability to Adversarial Attacks

Federated Learning systems may be vulnerable to malicious participants attempting to manipulate the collaborative training process.

Potential attacks include:

1. **Model Poisoning:** Malicious participants submit manipulated model updates to reduce global model performance.
2. **Data Poisoning:** Attackers introduce misleading or malicious samples into local training datasets.
3. **Backdoor Attacks:** Adversaries manipulate model training to create hidden conditions under which malicious inputs are incorrectly classified as normal.
4. **Free-Riding Attacks:** Participants attempt to benefit from collaborative training without contributing meaningful model updates.

Limitation: Detecting malicious participants and validating model updates remains a significant challenge in large-scale federated environments.

CONCLUSION

The rapid expansion of Industrial Internet of Things (IIoT) technologies has significantly transformed modern industrial infrastructures by enabling intelligent automation, real-time communication, remote monitoring, and data-driven operational management. However, the increasing connectivity of industrial devices has also introduced significant cybersecurity challenges, including unauthorized access, malware infections, distributed denial-of-service attacks, data manipulation, and sophisticated network intrusions. Traditional centralized intrusion detection systems face limitations related to data privacy, communication overhead, scalability, and the protection of sensitive industrial information. In this context, Federated Learning (FL) has emerged as a promising approach for developing intelligent and privacy-preserving intrusion detection frameworks.

This research examined the potential of Federated Learning frameworks for improving intrusion detection capabilities in Industrial IoT networks while minimizing the need for centralized raw-data collection. The proposed framework integrates distributed model training, secure aggregation, privacy-enhancing mechanisms, edge computing, and advanced machine learning techniques. Different federated learning architectures, including centralized, decentralized, and hierarchical approaches, were considered to understand their effectiveness in heterogeneous industrial environments. The study also explored the integration of Decision Trees, Random Forest, Support Vector Machines, Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM), and hybrid CNN-LSTM models for identifying abnormal network traffic and malicious activities.

The comparative analysis suggests that hierarchical federated learning integrated with CNN-LSTM architectures offers considerable potential for improving intrusion detection performance in large-scale industrial environments. The illustrative evaluation indicates an accuracy of 97.40% and an F1-score of 96.95% for the combined approach. These values represent hypothetical research scenarios and require validation through actual experimental implementation. The analysis also highlights the importance of differential privacy, secure aggregation, encryption, and blockchain-assisted verification in strengthening the confidentiality, integrity, and reliability of collaborative model training.

Furthermore, the research identifies several challenges associated with implementing federated intrusion detection systems, including Non-Independent and Identically Distributed (Non-IID) data, communication overhead, computational limitations, model poisoning, privacy leakage, and training instability. These limitations demonstrate that Federated Learning does not automatically guarantee complete data privacy or immunity against cyberattacks. Therefore, adaptive aggregation algorithms, personalized learning strategies, robust client authentication, and lightweight deep learning architectures are essential for improving the practical effectiveness of these systems.

REFERENCES

1. Ali, S., Li, Q., & Yousafzai, A. (2024). Blockchain and federated learning-based intrusion detection approaches for edge-enabled industrial IoT networks: A survey. *Ad Hoc Networks*. https://doi.org/10.1016/j.adhoc.2023.103320
2. Alsaedi, A., Moustafa, N., Tari, Z., Mahmood, A., & Anwar, A. (2020). TON_IoT telemetry dataset: A new generation dataset of IoT and IIoT for data-driven intrusion detection systems. *IEEE Access, 8*, 165130–165150.
3. Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020). How to backdoor federated learning. *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics, 108*, 2938–2948.
4. Blanchard, P., El Mhamdi, E. M., Guerraoui, R., & Stainer, J. (2017). Machine learning with adversaries: Byzantine tolerant gradient descent. *Advances in Neural Information Processing Systems, 30*.
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. *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191. https://doi.org/10.1145/3133956.3133982

6. Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science, 9*(3–4), 211–407. https://doi.org/10.1561/0400000042
7. Ferrag, M. A., Friha, O., Hamouda, D., Maglaras, L., & Janicke, H. (2022). Edge-IIoTset: A new comprehensive realistic cyber security dataset of IoT and IIoT applications for centralized and federated learning. *IEEE Access, 10*, 40281–40306.
8. Ferrag, M. A., Maglaras, L., Moschoyiannis, S., & Janicke, H. (2020). Deep learning for cyber security intrusion detection: Approaches, datasets, and comparative study. *Journal of Information Security and Applications, 50*, 102419.
9. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
10. Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., & Suresh, A. T. (2020). SCAFFOLD: Stochastic controlled averaging for federated learning. *Proceedings of the 37th International Conference on Machine Learning, 119*, 5132–5143.
11. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., D'Oliveira, R. G. L., Eichner, H., El Rouayheb, S., Evans, D., Gardner, J., Garrett, Z., Gascón, A., Ghazi, B., Gibbons, P. B., ... Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning, 14*(1–2), 1–210. https://doi.org/10.1561/2200000083
12. Li, T., Sahu, A. K., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems, 2*, 429–450.
13. McMahan, B., Moore, E., Ramage, D., Hampson, S., & Agüera y Arcas, B. (2017). Communication-efficient learning of deep networks from decentralized data. *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 54*, 1273–1282.
14. Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems (UNSW-NB15 network data set). *2015 Military Communications and Information Systems Conference (MilCIS)*, 1–6. https://doi.org/10.1109/MilCIS.2015.7348942
15. Ruzafa-Alcázar, P., Fernández-Saura, P., Mármol-Campos, E., González-Vidal, A., Hernández-Ramos, J. L., Bernal-Bernabe, J., & Skarmeta, A. F. (2022). Intrusion detection based on privacy-preserving federated learning for the Industrial IoT. *IEEE Transactions on Industrial Informatics*. https://doi.org/10.1109/TII.2021.3126728
16. Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models. *2017 IEEE Symposium on Security and Privacy*, 3–18. https://doi.org/10.1109/SP.2017.41
17. Stojanovic, L., Dinic, M., Stojanovic, N., & Stojadinovic, A. (2016). Big-data-driven anomaly detection in industry (4.0): An approach and a case study. *2016 IEEE International Conference on Big Data*, 1647–1652.
18. Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., & Bengio, Y. (2018). Graph attention networks. *International Conference on Learning Representations (ICLR)*.
19. Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology, 10*(2), Article 12. https://doi.org/10.1145/3298981
20. Zhu, L., Liu, Z., & Han, S. (2019). Deep leakage from gradients. *Advances in Neural Information Processing Systems, 32*.