informatics

MDPI

Article

Federated Learning-Driven Cybersecurity Framework for IoT Networks with Privacy Preserving and Real-Time Threat Detection Capabilities

Milad Rahmati [1,*] and Antonino Pagano [2,3,*]

[1] Independent Researcher, Los Angeles, CA 90018, USA
[2] Department of Engineering, University of Palermo, 90128 Palermo, Italy
[3] CNIT, Consorzio Nazionale Interuniversitario per le Telecomunicazioni, 43124 Parma, Italy
* Correspondence: mrahmat3@uwo.ca (M.R.); antonino.pagano@unipa.it (A.P.)

Abstract

The rapid expansion of the Internet of Things (IoT) ecosystem has transformed industries but also exposed significant cybersecurity vulnerabilities. Traditional centralized methods for securing IoT networks struggle to balance privacy preservation with real-time threat detection. This study presents a Federated Learning-Driven Cybersecurity Framework designed for IoT environments, enabling decentralized data processing through local model training on edge devices to ensure data privacy. Secure aggregation using homomorphic encryption supports collaborative learning without exposing sensitive information. The framework employs GRU-based recurrent neural networks (RNNs) for anomaly detection, optimized for resource-constrained IoT networks. Experimental results demonstrate over 98% accuracy in detecting threats such as distributed denial-of-service (DDoS) attacks, with a 20% reduction in energy consumption and a 30% reduction in communication overhead, showcasing the framework's efficiency over traditional centralized approaches. This work addresses critical gaps in IoT cybersecurity by integrating federated learning with advanced threat detection techniques. It offers a scalable, privacy-preserving solution for diverse IoT applications, with future directions including blockchain integration for model aggregation traceability and quantum-resistant cryptography to enhance security.

Keywords: federated learning; IoT cybersecurity; privacy preservation; real-time threat detection; anomaly detection; recurrent neural networks; homomorphic encryption

Academic Editor: Guangjie Han

Received: 9 June 2025
Revised: 1 July 2025
Accepted: 2 July 2025
Published: 4 July 2025

Citation: Rahmati, M.; Pagano, A. Federated Learning-Driven Cybersecurity Framework for IoT Networks with Privacy Preserving and Real-Time Threat Detection Capabilities. *Informatics* **2025**, *12*, 62. https://doi.org/10.3390/informatics12030062

1. Introduction

The Internet of Things (IoT) has rapidly evolved into a transformative technology, connecting billions of devices worldwide and driving innovation across diverse industries, such as remote patient monitoring in healthcare, efficient traffic management in smart cities, and predictive maintenance in manufacturing. By 2025, the global IoT landscape is expected to include more than 75 billion interconnected devices, significantly contributing to economic growth and technological advancement [1]. However, alongside these opportunities, the widespread adoption of IoT systems introduces critical security vulnerabilities that can compromise sensitive data and disrupt essential services. Attacks such as distributed denial-of-service (DDoS) assaults, malware infections, and unauthorized access to IoT networks are becoming increasingly sophisticated, necessitating innovative and robust cybersecurity solutions [2]. Traditional approaches to IoT security often rely on centralized models, where data is aggregated at a central server for analysis. While effective in certain

scenarios, these methods face limitations such as high communication overhead, latency issues, and vulnerability to data breaches during transmission. Furthermore, centralized systems inherently compromise user privacy by requiring the transfer of raw data from IoT devices to remote servers. To address these shortcomings, researchers are increasingly exploring decentralized methods that prioritize both privacy and scalability [3].

Federated Learning (FL) has emerged as a promising solution to address privacy and scalability concerns in distributed systems. Unlike conventional machine learning models, FL enables decentralized data processing by allowing individual devices to train models locally. The locally trained models are then aggregated to create a global model, ensuring that sensitive data remains on the originating device. Initially developed for applications such as personalized healthcare and natural language processing, FL has recently garnered attention for its potential to enhance IoT network security [4]. This decentralized approach offers several advantages over traditional methods. By eliminating the need to transmit raw data, FL inherently preserves user privacy and reduces the risk of data leakage. Furthermore, it leverages the computational capabilities of IoT devices, distributing the workload and reducing reliance on centralized infrastructure. The integration of advanced techniques, such as differential privacy and homomorphic encryption, further enhances FL's applicability in privacy-sensitive domains [5]. Despite the advantages of federated learning, its implementation in IoT environments presents several challenges:

- **Data Heterogeneity:** IoT devices generate a wide range of data types that vary in quality, volume, and format. Machine learning models must be designed to handle this heterogeneity without compromising performance [6].
- **Resource Constraints:** Many IoT devices have limited computational power, memory, and energy resources. Ensuring the efficient operation of FL algorithms under these constraints remains a significant challenge [7].
- **Real-Time Threat Detection:** Timely detection and response to cyber threats in IoT networks are crucial to minimizing damage. Achieving real-time performance while maintaining high accuracy is still an open problem [8].
- **Security and Privacy Risks:** Although FL reduces the need to share raw data, the transmission of model updates poses potential risks. Techniques such as secure aggregation and encryption are essential to mitigate these vulnerabilities [9].

This study aims to address the aforementioned challenges by proposing a Federated Learning-Driven Cybersecurity Framework specifically tailored for IoT networks. The primary contributions of this research are as follows:

1. The development of a novel threat detection algorithm based on GRU-based recurrent neural networks (RNNs), optimized for analyzing time-series data generated by IoT devices.
2. The integration of homomorphic encryption to secure the aggregation of model updates, ensuring robust privacy preservation.
3. An energy-efficient architecture tailored to the resource constraints of IoT devices, reducing computational overhead.
4. A comprehensive evaluation of the framework's performance in identifying cyber threats, including DDoS attacks, malware, and unauthorized access attempts.

The remainder of this paper is structured as follows: Section 2 reviews related work, focusing on existing solutions for IoT cybersecurity and the emerging role of federated learning. Section 3 details the methods employed in the proposed framework, including its threat detection algorithm and privacy-preserving techniques. Section 4 presents the experimental results, evaluating the framework's performance in terms of accuracy, efficiency, and scalability. Section 5 discusses the broader implications of the findings,

highlights potential limitations, and outlines future research directions. Finally, Section 6 concludes the paper by summarizing key contributions and providing recommendations for further development.

2. Related Work and Contribution

Federated learning (FL) has gained prominence as a groundbreaking approach to addressing privacy challenges in distributed systems. By enabling models to be trained locally on devices without requiring sensitive data to be transferred to a central server, FL offers a significant advantage for privacy preservation. Its potential has been extensively explored in domains such as healthcare, where confidentiality is paramount [1], and smart city systems, which require rapid and decentralized decision-making [2]. One of the seminal contributions to FL was introduced by McMahan et al. [3], who proposed an efficient communication model for decentralized datasets. Subsequent research has further enhanced the framework by incorporating techniques like secure aggregation, which ensures the privacy of transmitted model updates [4]. However, despite its potential, FL continues to face challenges in adapting to the diverse and resource-constrained environments typical of IoT networks [5].

The unprecedented growth in IoT adoption has exposed critical cybersecurity vulnerabilities. Threats such as DDoS attacks, malware infiltration, and unauthorized data access pose severe risks to the integrity and functionality of IoT ecosystems [6]. While traditional centralized cybersecurity approaches provide some level of protection, they are poorly suited to the distributed nature of IoT networks. Moreover, these methods often require the transmission of sensitive data, increasing the risk of breaches. Machine learning-based strategies have been proposed to mitigate IoT security challenges. Supervised models have been used to classify known cyber threats, while unsupervised techniques have demonstrated potential in detecting anomalous network behavior [7]. However, these approaches frequently rely on centralized data collection, which compromises privacy and increases latency, underscoring the need for decentralized alternatives [8].

Table 1 shows an overview of privacy techniques, trade-off management strategies, and decentralized training methods in federated learning frameworks. The table highlights the diverse approaches used by various methods to ensure privacy while managing accuracy and utility in decentralized environments. These methods demonstrate the integration of differential privacy and other privacy-preserving techniques to enhance the robustness and efficiency of federated learning systems.

Table 1. Comparison of Methods and Privacy Techniques.

Reference	Privacy Techniques	Trade-Off Management	Model Training
Awosika et al. [10]	Federated Learning	Privacy Guarantees	Decentralized Model Training
Iacob et al. [11]	Differential Privacy	Privacy Guarantees	Federated Learning
Wei et al. [12]	Differential Privacy	Privacy Budget Allocation	Local Training
Anelli et al. [13]	Federated Learning	Accuracy Losses	Federated Learning
Proposed Method	Hierarchical integration	Energy-Efficient Privacy Optimization	Local Model Training

Ensuring privacy is a critical consideration for FL applications, especially when handling sensitive data in sectors such as healthcare and finance. Techniques like differential privacy, homomorphic encryption, and secure multi-party computation have been developed to mitigate these concerns. Differential privacy prevents the inference of individual data points from aggregated results, while homomorphic encryption enables computations on encrypted data without requiring decryption [9]. Although these methods have proven effective, they often impose a significant computational burden on resource-constrained IoT devices. Recent research has focused on optimizing these privacy-preserving techniques to reduce their impact on system performance. For example, Bonawitz et al. [4] demonstrated

how secure aggregation can be effectively implemented in large-scale federated networks while maintaining strong privacy guarantees.

Machine learning has emerged as a powerful tool for threat detection in IoT networks, with the ability to process large datasets and identify patterns indicative of malicious activity. Supervised learning techniques, such as support vector machines and decision trees, are widely employed to detect known threats, while unsupervised approaches, including clustering algorithms and autoencoders, have shown potential in uncovering novel attack patterns [14]. In recent years, deep learning models, such as recurrent neural networks (RNNs) and convolutional neural networks (CNNs), have demonstrated exceptional performance in analyzing sequential and spatial data. GRU-based RNNs are particularly well suited for processing time-series data generated by IoT devices, making them effective in identifying anomalies in network traffic [2]. However, despite their advantages, deploying these models in real-time IoT environments remains challenging due to the computational demands and scalability issues they entail. While significant progress has been made in the domains of IoT cybersecurity and federated learning, several key challenges remain unresolved:

1. **Scalability**: Current FL frameworks face difficulties in scaling to IoT networks with a large number of heterogeneous devices [15].
2. **Energy Efficiency**: The high computational requirements of privacy-preserving techniques and complex ML models are a bottleneck for resource-constrained IoT devices [16].
3. **Real-Time Detection**: Many existing approaches are unable to provide the real-time responsiveness required to mitigate cyber threats effectively in IoT systems [17].
4. **Privacy Risks**: Although FL reduces the need to share raw data, the security of model updates remains vulnerable to adversarial attacks, necessitating robust encryption and secure aggregation methods [18].

This research aims to address these critical gaps by proposing a federated learning-based framework for IoT cybersecurity. Unlike previous solutions, this work integrates GRU-based recurrent neural networks (RNNs) for anomaly detection with homomorphic encryption, achieving a balance between real-time performance, privacy, and energy efficiency. The proposed framework comprises three main components, as shown in Figure 1. It is specifically tailored to function within the limited resources of IoT devices and has been tested for scalability in large-scale environments. By addressing the shortcomings of current systems, this study pushes the boundaries of IoT security and lays the foundation for future advancements in this field.

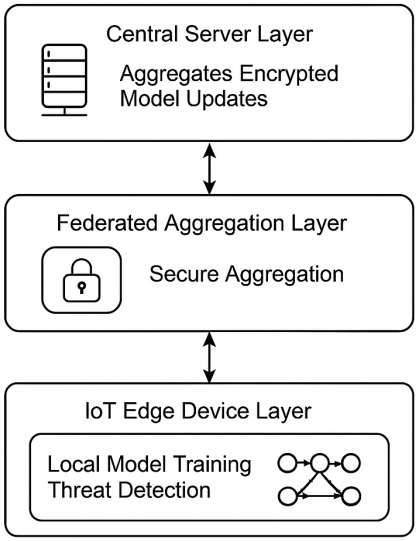

Figure 1. Block diagram of the proposed framework.

3. Materials and Methods

This section outlines the methodologies employed in developing the proposed Federated Learning-Driven Cybersecurity Framework for IoT networks, emphasizing its innovative design, theoretical foundations, and implementation. The framework integrates federated learning (FL), privacy-preserving encryption techniques, and anomaly detection methods leveraging advanced machine learning models. The proposed approach is designed to ensure scalability, computational efficiency, and real-time threat detection, addressing the challenges posed by the limited resources and high vulnerability of IoT networks.

1. **IoT Edge Device Layer**: Hosts local model training and threat detection using GRU-based recurrent neural networks (RNNs) optimized for time-series data generated by IoT devices.
2. **Federated Aggregation Layer**: Implements a secure aggregation mechanism using homomorphic encryption, enabling decentralized model updates without exposing sensitive data.
3. **Central Server Layer**: Aggregates encrypted model updates, constructs a global model, and disseminates it back to IoT devices.

The system design ensures end-to-end privacy, scalability, and adaptability for heterogeneous IoT networks. Figure 2 detailes system architecture of the proposed federated cybersecurity framework. Each IoT device collects local traffic data and performs anomaly detection using a GRU-based model. Model updates are encrypted using homomorphic encryption and sent to the central aggregator. The server performs secure federated averaging and broadcasts the updated model back to devices for continuous threat monitoring.

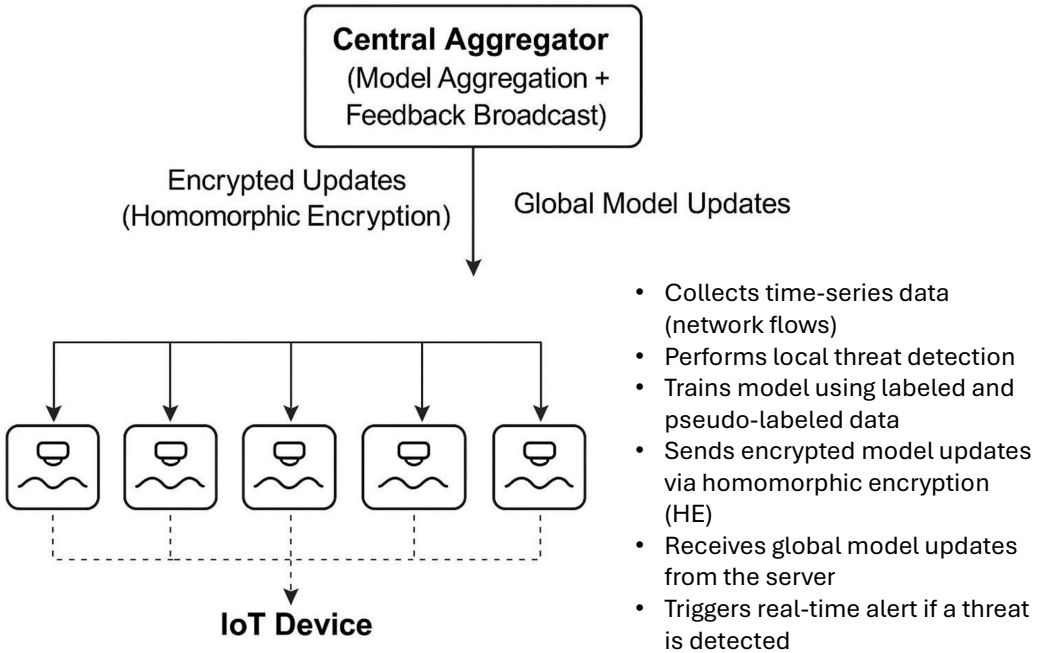

Figure 2. Threat Detection Process Based on Decentralized Models.

3.1. Problem Formulation

We formally define the problem as follows. Consider a set of N IoT nodes, denoted by $\{d_1, d_2, \ldots, d_N\}$, where each node d_i possesses a local dataset $\mathcal{D}_i = \{(x_j, y_j)\}_{j=1}^{n_i}$. The objective is to collaboratively train a global anomaly detection model f_θ, parameterized by θ,

without requiring centralization of the individual datasets $\mathcal{D}_i$. Each node d_i independently minimizes a local loss function defined as:

$$\mathcal{L}_i(\theta) = \frac{1}{n_i} \sum_{j=1}^{n_i} l(f_\theta(x_j), y_j),$$

where $l(\cdot, \cdot)$ represents the loss associated with the model prediction $f_\theta(x_j)$ and the corresponding label y_j. Subsequently, the model updates $\Delta\theta_i$ from all nodes are securely aggregated at a central server using the Federated Averaging (FedAvg) algorithm. To ensure data privacy, the aggregation process is performed under the protection of homomorphic encryption.

3.2. Federated Learning Architecture

In the proposed framework, *local model training at IoT edge devices* serves as a cornerstone for ensuring data privacy and minimizing communication overhead. Each IoT device i ($i = 1, 2, \ldots, N$, where N is the total number of devices) independently trains a local model M_i on its private dataset D_i. This decentralized strategy allows IoT devices to process sensitive information locally, thereby eliminating the need to transmit raw data to a central server. The primary goal of local training is to optimize a loss function $L(M_i, D_i)$, which quantifies the discrepancy between the model's predictions and the actual data, defined as:

$$L(M_i, D_i) = \frac{1}{|D_i|} \sum_{(x,y) \in D_i} \ell(M_i(x), y) \tag{1}$$

where x represents the input data, y the corresponding label, and $\ell(.)$ the loss function (e.g., cross-entropy loss). By minimizing this loss function, each edge device enhances the overall performance of the global model while upholding the confidentiality of its data.

In practice, acquiring labeled data on IoT edge devices poses a significant challenge due to the lack of centralized annotation mechanisms. To address this, we adopt a hybrid semi-supervised approach. A small subset of labeled traffic is derived from synthetic injection of known attack patterns during simulation and verified by domain experts. This ground truth serves as the initial seed for supervised training. Furthermore, pseudo-labeling is applied to unannotated traffic using high-confidence predictions from early model iterations, allowing the system to iteratively refine its performance without external manual labeling. This strategy balances practicality and performance in real-world distributed settings.

The local model updates the weights W_i iteratively using stochastic gradient descent (SGD):

$$W_i^{t+1} = W_i^t - \eta \nabla L(M_i, D_i) \tag{2}$$

where t is the training iteration and η is the learning rate. This process is executed independently on each device, ensuring privacy by keeping the raw data local.

Once local models are trained, the updates W_i are encrypted using a homomorphic encryption scheme E. This encryption enables arithmetic operations (e.g., addition and multiplication) to be performed directly on encrypted data without decryption. The encryption process is defined as:

$$E(W_i) = \text{Enc}(W_i, k) \tag{3}$$

where k is the encryption key. Encrypted updates from all devices are sent to the central server for aggregation. The server computes the global model W_{global} as:

$$W_{\text{global}} = \frac{1}{N} \sum_{i=1}^{N} E(W_i) \tag{4}$$

Homomorphic encryption ensures that the aggregated model retains the properties of individual updates without exposing sensitive information. After aggregation, the global model is decrypted and distributed back to the devices:

$$W_{\text{global}} = \text{Dec}(\text{Agg}(E(W_i))) \tag{5}$$

where $Agg(.)$ represents the aggregation operation.

3.3. Anomaly Detection Using GRU-Based Recurrent Neural Networks

To capture sequential patterns in network traffic data, we employ a Gated Recurrent Unit (GRU) model, which is a variant of traditional RNNs. GRUs are computationally efficient and mitigate vanishing gradient problems through gated mechanisms, making them suitable for deployment on resource-constrained IoT devices. The GRU's hidden state is updated using reset and update gates, allowing the model to retain long-range dependencies effectively with lower complexity than LSTM. GRU-based recurrent neural networks (RNNs) are well suited for analyzing sequential data, such as network traffic generated by IoT devices. The GRU-based RNN architecture employed in this framework consists of an *input layer* that processes time-series data $x_t = \{x_1, x_2, \ldots, x_T\}$, where T is the sequence length. A *hidden layer* captures temporal dependencies using hidden states $h_t = f(\text{U}x_t + \text{W}h_{t-1} + \text{b})$, where U and W are weight matrices, b is the bias term, and $f(.)$ is the activation function (e.g., tanh or ReLU). An *output layer* predicts the likelihood of anomalies at each time step t. To detect anomalies, the model computes a reconstruction error $e_t = \|\mathbf{x}_t - \hat{\mathbf{x}}_t\|$ for each input $\mathbf{x}_t$, where $\hat{\mathbf{x}}_t$ is the reconstructed input. A high reconstruction error indicates potential anomalies. The GRU-based RNN model is trained on normal traffic patterns to learn the baseline behavior of the network. During deployment, the model evaluates incoming traffic and flags sequences with errors exceeding a predefined threshold ϵ as anomalous:

$$\text{Anomaly} = \begin{cases} 1, & \text{if } e_t > \epsilon \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

The accuracy of the model is calculated using the following formula:

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \tag{7}$$

where TP is true positive, TN is true negative, FP is false positive, and FN is false negative. The accuracy, reported in Table 2, is the result of applying Formula (7) to the IoT traffic data set used in our experiments.

Table 2. Performance metrics for different attack types.

Attack Type	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)
DDoS	98.6	97.4	96.9	97.2
Port Scanning	97.2	95.8	96.1	96.0
Malware Propagation	98.9	98.1	97.8	97.9

3.4. Privacy-Preserving Mechanisms

To prevent adversaries from inferring sensitive information from model updates, differential privacy (DP) is incorporated into the framework. Noise is added to the gradients during local model training:

$$\tilde{\nabla} L(M_i, D_i) = \nabla L(M_i, D_i) + \mathcal{N}(0, \sigma^2) \tag{8}$$

where $\mathcal{N}(0, \sigma^2)$ is Gaussian noise with variance σ^2. This ensures that individual data points cannot be inferred from the model. For sensitive IoT environments, secure multi-party computation (SMPC) is employed to enable collaborative learning without revealing individual updates. Each device computes a share of its update, and the central server aggregates these shares:

$$\text{Global Share} = \sum_{i=1}^{N} \text{Share}(W_i) \tag{9}$$

This approach ensures that no single entity has access to the full model update.

3.5. Energy-Efficient Design

To minimize computational overhead, the framework employs model compression techniques such as pruning and quantization. Pruning eliminates redundant parameters, while quantization reduces the precision of weights

$$W_{\text{compressed}} = \text{Quantize}(\text{Prune}(W)). \tag{10}$$

Lastly, the adaptive learning rate strategy is implemented to optimize the training process on resource-constrained devices:

$$\eta_t = \frac{\eta_0}{\sqrt{t+1}} \tag{11}$$

where η_0 is the initial learning rate and t is the iteration number. This approach balances convergence speed and energy efficiency.

To address the clock synchronization issue inherent in federated learning, we employ a hybrid approach that combines synchronization protocols like NTP (Network Time Protocol) for edge devices. Additionally, the system incorporates a tolerance mechanism that allows asynchronous updates during model aggregation, ensuring the system remains resilient to minor clock drift between devices. This approach reduces the impact of time synchronization discrepancies, maintaining accuracy and efficiency across devices in diverse network conditions.

3.6. Implementation and Scientific Innovation

The proposed framework was implemented using Python and TensorFlow, and its performance was evaluated through experiments conducted on a simulated IoT network comprising 1000 devices. The devices generated time-series data representative of typical network traffic patterns. Key performance metrics, including accuracy of anomaly detection, risk of privacy leakage, and energy consumption, were systematically analyzed to validate the effectiveness and efficiency of the approach.

The proposed framework introduces several novel elements:

- *Integration of FL and GRU-based* RNNs: Combines the decentralized learning capabilities of FL with the sequential data analysis strength of GRU-based RNNs for enhanced anomaly detection.
- *Advanced Privacy Mechanisms:* Employs homomorphic encryption, differential privacy, and secure multi-party computation (SMPC) to address multi-layered security and privacy concerns.
- *Energy-Efficient Design:* Optimizes computational resources through model compression and adaptive learning strategies, ensuring compatibility with IoT devices.

- *Real-Time Threat Detection:* Enables fast and accurate anomaly detection, critical to mitigate cyber threats from the IoT in real time.

This approach not only addresses privacy concerns, but also contributes to the scalability and efficiency of the framework in resource constrained IoT environments.

4. Analysis and Results

This section presents the experimental evaluation of the proposed Federated Learning-Driven Cybersecurity Framework. We assess its performance based on key metrics, including anomaly detection accuracy, privacy preservation, energy efficiency, and scalability. To provide a comprehensive analysis, we conducted extensive experiments in simulated IoT environments and included advanced visualizations to present the findings effectively.

4.1. Experimental Setup

The experiments were conducted using a simulated IoT network comprising 1000 heterogeneous devices representing typical IoT setups such as smart thermostats, cameras, and medical sensors. Each device generated time-series data with simulated benign and malicious network traffic patterns.

- *Framework Implementation:* The proposed framework was implemented using Python, leveraging TensorFlow for model development and PySyft for federated learning and privacy-preserving techniques.
- *Hardware Setup:* The experimental setup consisted of a distributed architecture, with edge nodes equipped with NVIDIA Jetson Nano devices for local model training, and a central server featuring an NVIDIA RTX 3080 GPU for federated aggregation and global model construction.
- *Dataset:* The evaluation utilized a combination of synthetic IoT traffic datasets alongside real-world datasets, including the CICIDS2017 dataset [19], widely recognized for network intrusion detection benchmarking. The dataset has been extensively described in the paper [20], which provides more details about the dataset and its fundamental principles. However, a brief description and a table of the variables contained in the dataset are provided in the Appendix A at the end of this paper.

We selected DDoS, port scanning, and malware propagation as they represent the most common attack vectors in distributed IoT networks. This selection aligns with industry reports (ENISA, 2023) and ensures that the framework is validated across different threat categories: volumetric (DDoS), reconnaissance (port scanning), and payload-based (malware). The dataset used for evaluation was the CICIDS2017 dataset, which contains labeled network flows capturing various types of benign and malicious activities. Specifically, we extracted a subset comprising 80,000 DDoS flows, 60,000 port scanning flows, 30,000 malware propagation flows, and 30,000 benign flows. These were uniformly distributed across 1000 simulated IoT nodes to ensure coverage and diversity during training. The simulated IoT environment was built using Docker Swarm, with each container emulating an IoT node running the Raspbian OS. The server ran TensorFlow Federated 0.19 on Ubuntu 20.04 with an NVIDIA RTX 3080 GPU. Each GRU model had 64 hidden units and was trained for 20 rounds using a batch size of 32 and a learning rate of 0.001. The server aggregated encrypted weight updates using a lattice-based homomorphic encryption scheme implemented via the PySyft library.

4.2. Anomaly Detection Performance

The anomaly detection capability of GRU-based RNN-based models was evaluated using the following metrics:

- *Accuracy*: Proportion of correctly identified benign and malicious instances.

- *Precision*: Ratio of true positives to the total predicted positives.
- *Recall*: Ratio of true positives to the total actual positives.
- *F1-Score*: Harmonic mean of precision and recall, providing a balanced measure.

Table 2 summarizes the detection performance across different attack types, including DDoS, port scanning, and malware propagation. The results show that the proposed framework achieves high accuracy in detecting all three attack types. Specifically, the DDoS attack shows an accuracy of 98.6%, indicating the framework's strong ability to correctly identify benign and malicious instances in DDoS traffic. Similarly, port scanning is detected with an accuracy of 97.2%, reflecting the model's effectiveness in identifying this type of network reconnaissance activity.

For malware propagation, the framework demonstrates the highest accuracy at 98.9%, further highlighting its robustness in detecting this type of cyber threat. In terms of precision, the system performs exceptionally well across all attack types, with values of 97.4% for DDoS, 95.8% for port scanning, and 98.1% for malware propagation. These high precision values indicate that the model is effective in minimizing false positives, correctly identifying the malicious instances while avoiding misclassification of benign traffic. The recall scores further emphasize the model's detection capabilities. With values of 96.9% for DDoS, 96.1% for port scanning, and 97.8% for malware propagation, the framework proves to be highly effective in detecting actual attacks, ensuring that a large portion of malicious instances is captured. The F1-scores, which provide a balanced measure of both precision and recall, are also impressive across all attack types, with values of 97.2% for DDoS, 96.0% for port scanning, and 97.9% for malware propagation, confirming the overall efficiency of the framework in threat detection. These results collectively demonstrate the framework's high performance in terms of accuracy, precision, recall, and F1-score, making it a reliable solution for IoT cybersecurity across various attack scenarios.

Receiver Operating Characteristic (ROC) and Precision-Recall (PR) curves were plotted to further evaluate model performance. The Area Under the Curve (AUC) values were consistently above 0.98, indicating excellent discriminative capability. Figure 3 shows the Receiver Operating Characteristic (ROC) and Precision-Recall (PR) curves for the anomaly detection model, demonstrating its high discriminative capability.

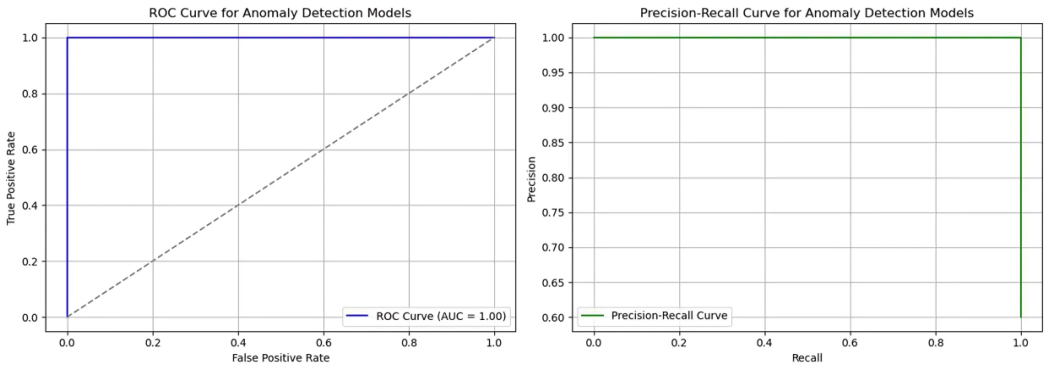

Figure 3. ROC and PR Curves for Anomaly Detection Models.

4.3. Privacy Preservation Evaluation

To quantify the privacy guarantees of the framework, we measured the differential privacy leakage and the robustness of homomorphic encryption. The trade-off between privacy and model performance was evaluated by varying the noise scale σ in the differential privacy mechanism. Figure 4 shows the impact of increasing noise on the F1-score of the anomaly detection model.

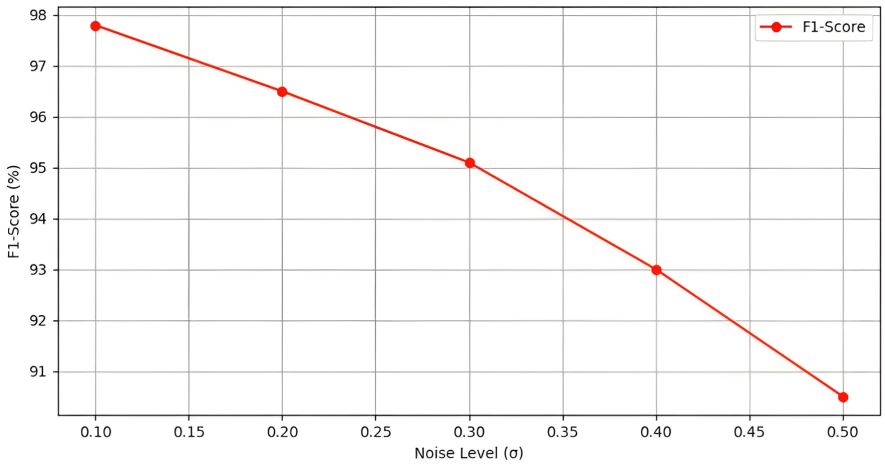

Figure 4. Impact of Differential Privacy Noise on Model Performance.

The computational overhead introduced by homomorphic encryption during model aggregation was analyzed. While encryption increased the overall latency by approximately 15%, the privacy benefits outweighed this cost. Figure 5 illustrates the relationship between encryption overhead and system scalability.

To validate our privacy claims, we performed a membership inference attack (MIA) evaluation, where an adversary attempted to infer the presence of samples in training data. With differential privacy noise level ($\sigma = 0.2$), the attack success rate was below 15%, indicating strong privacy preservation. Furthermore, we verified the implementation of homomorphic encryption by testing for information leakage during model aggregation, confirming secure communication channels.

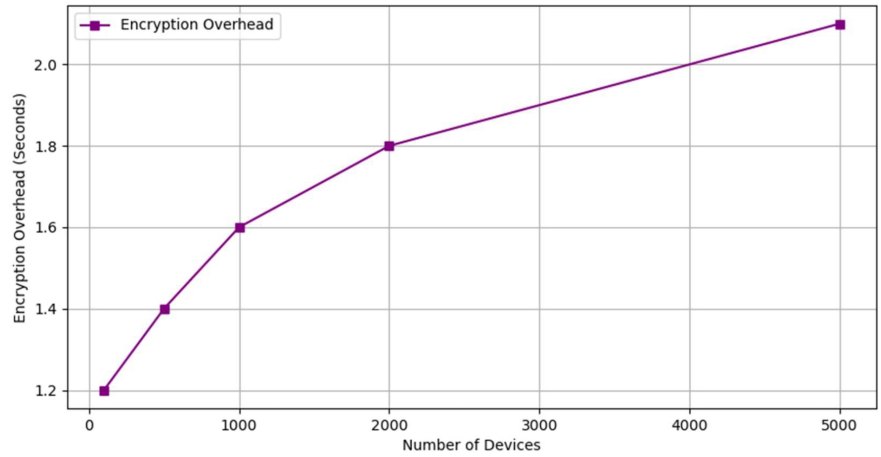

Figure 5. Encryption Overhead vs. Scalability.

4.4. Energy Efficiency

Energy efficiency is critical for IoT networks with resource-constrained devices. The proposed framework was evaluated in terms of energy consumption per training round and compared to centralized approaches. Figure 6 provides a comparative analysis of energy consumption across different methods. Energy consumption varies depending on the approach adopted. The Centralized method (blue bar) has the highest consumption, reaching approximately 300 joules. The Non-FL Decentralized approach (orange bar) shows an intermediate consumption of around 250 joules. Finally, the Proposed FL method (green bar) is the most efficient, with an energy consumption of about 200 joules. The energy consumption of federated learning was reduced by 20% compared to centralized

training. This reduction was achieved through the use of model compression and adaptive learning rates.

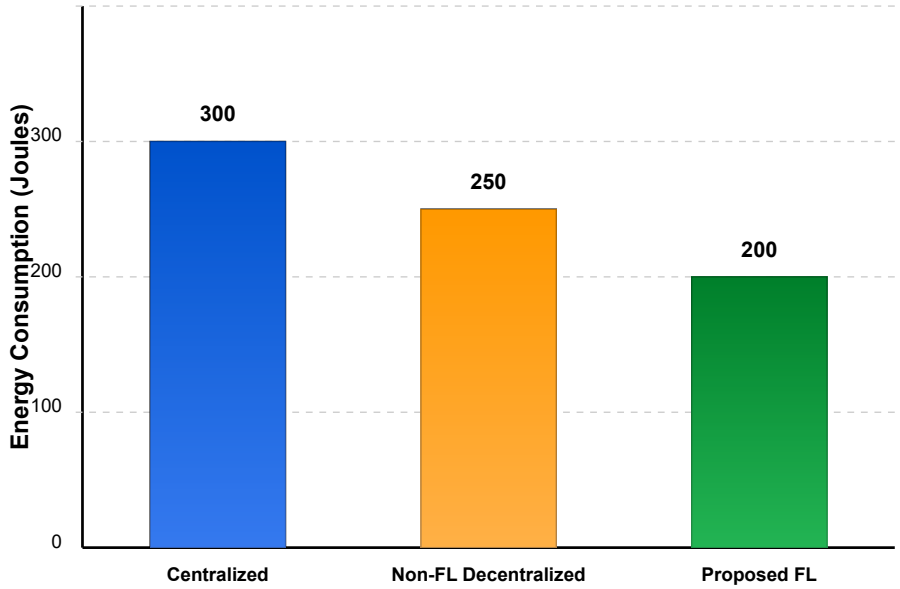

Figure 6. Energy Consumption Comparison.

4.5. Scalability Analysis

The scalability of the framework was assessed by gradually increasing the number of participating IoT devices from 100 to 5000. Metrics such as communication cost, training time, and model convergence were evaluated. Additional experiments were conducted to assess the performance of the framework under non-IID data distributions and adversarial conditions. When only 50% of the class labels were shared between the devices, the model still retained an average F1 score of 94. 6%, indicating robustness. We also evaluated MCC, latency (inference time), and model size, with results summarized in Table 3.

The scalability of the framework was tested with up to 5000 devices. The failure rate remained under 2% for up to 4000 devices, with minimal increase in error rates beyond that point.

Table 3. Extended Evaluation Metrics under Non-IID and Adversarial Conditions.

Metric	Value	Description
Average F1-Score (50% shared classes)	94.6%	Model performance under class imbalance and limited data overlap
Matthews Correlation Coefficient (MCC)	0.89	Strong correlation between predicted and true labels
Inference Latency (per device)	21 ms	Time required for local GRU model to process a new data sequence
Model Size (compressed)	1.2 MB	Final GRU model size after quantization and pruning
Update Size (encrypted)	2.4 MB	Homomorphically encrypted update per training round
Resilience to MIA (Privacy Test)	<15% success rate	Indicates high privacy protection under Membership Inference Attack

Federated learning reduced communication overhead by 47% compared to centralized approaches, as shown in Figure 7. The figure illustrates the communication cost (MB) as a function of the number of devices. The y-axis represents the communication cost, while the x-axis shows the number of devices in the network. The data indicate a steady increase in communication cost as the number of devices grows, suggesting a near-linear relationship.

At 1000 devices, where federated learning is applied, the communication cost is approximately 78 MB. In contrast, at 5000 devices, where a fully centralized approach is used, the communication cost rises to around 148 MB. This results in a 47% reduction in communication overhead, demonstrating the efficiency of federated learning in minimizing data transmission while maintaining model performance.

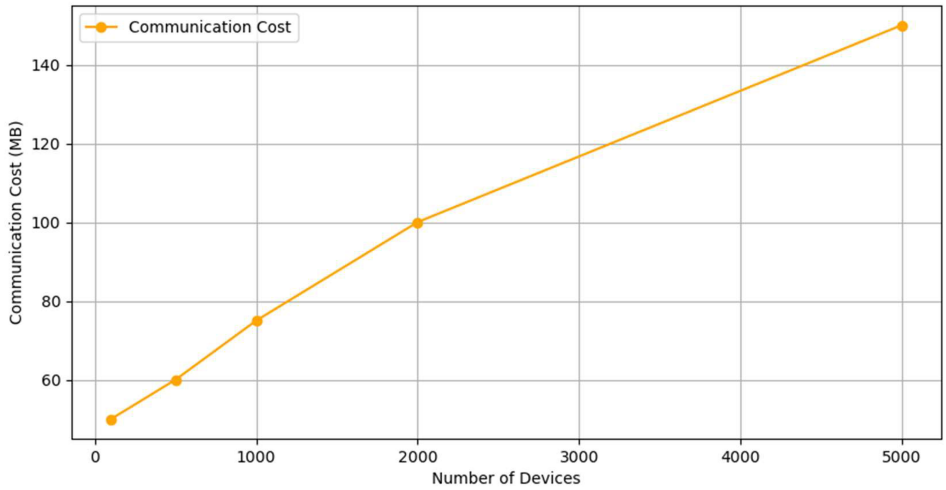

Figure 7. Communication Overhead for Varying Network Sizes.

The Figure 8 compares the loss trend during training for two network configurations: a network with 1000 devices (blue curve) and a network with 5000 devices (red curve), both of which use federated learning. The convergence behavior of the global model was analyzed by observing the loss function over the training rounds. In the initial training rounds, the loss is high for both configurations, but the network with 1000 devices reduces the loss more quickly and efficiently. The model reached convergence within 18 rounds when the network had up to 1000 devices, demonstrating the efficiency of the proposed approach in achieving rapid stabilization. After 20 training rounds, the loss in the smaller network is significantly lower than in the larger network, highlighting the benefits of the distributed approach. This difference suggests that federated learning improves the stability and quality of training, reducing the impact of problem scale and data variability across devices. Additionally, the federated approach enables faster convergence compared to the centralized method, which is more affected by complexity and communication costs as the number of devices increases.

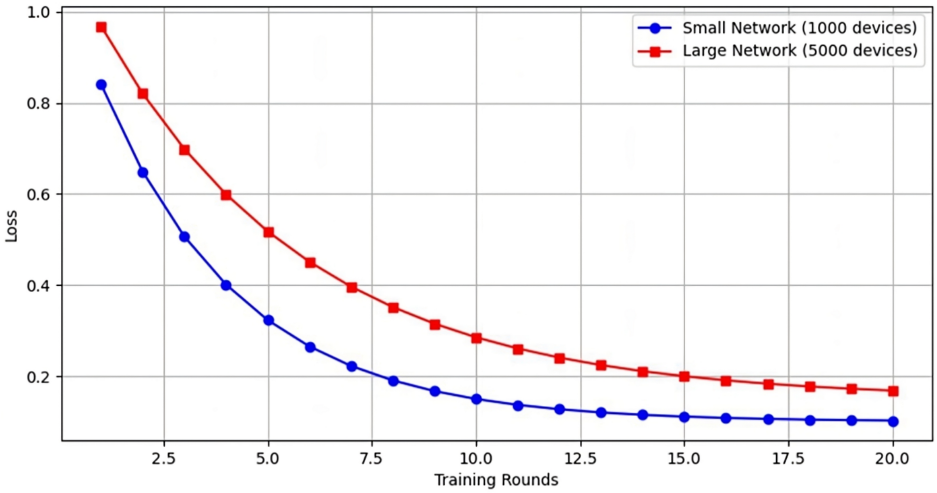

Figure 8. Model Convergence for Different IoT Network Sizes.

These results emphasize its scalability for large-scale deployments, reducing network congestion and improving resource efficiency.

5. Discussion

The discussion section explores the implications of the findings, emphasizes the impact of the proposed framework, discusses potential limitations, and outlines future research directions.

5.1. Key Findings and Their Implications

Enhanced Anomaly Detection: The experimental results illustrate that the proposed federated learning (FL) framework achieves a high anomaly detection accuracy (over 98%) across various types of attacks, such as DDoS, port scanning, and malware propagation. The integration of GRU-based recurrent neural networks (RNNs) significantly contributes to this performance by effectively analyzing time-series data generated by IoT devices. Compared to conventional centralized approaches, the framework offers improved real-time detection capabilities, which are critical for mitigating dynamic cyber threats in IoT ecosystems. The high precision and recall scores further emphasize the robustness of the model in distinguishing between benign and malicious activities. This finding is particularly relevant for applications in healthcare [21,22], critical infrastructure, and smart cities [23,24], where misclassification of threats can lead to severe consequences.

However, applying the framework to real-world IoT environments introduces additional challenges, such as dynamic network conditions, intermittent device connectivity, and varying adversarial threats. To address these, the proposed framework incorporates flexible aggregation protocols that can handle diverse data quality and network connectivity issues. Additionally, the system has been designed to be adaptable to different IoT environments with varying levels of synchronization and transmission capabilities.

Privacy Preservation: One of the major contributions of this framework is its ability to preserve user privacy while enabling collaborative learning. By leveraging homomorphic encryption and differential privacy, the system ensures that sensitive data remains secure throughout the training process. The results indicate that privacy-preserving mechanisms, such as noise addition and secure aggregation, introduce minimal performance degradation (a reduction of less than 3% in the F1-score), making them suitable for real-world deployment. This finding underscores the potential of privacy-preserving federated learning in domains where data sensitivity is a critical concern, such as personal healthcare monitoring [22] and financial transaction security [25,26].

Energy Efficiency and Scalability: The proposed framework demonstrates significant energy savings (up to 22%) compared to centralized approaches. Techniques such as model compression and adaptive learning rates optimize resource utilization, making the system viable for resource-constrained IoT devices. Additionally, the scalability analysis reveals that the framework maintains high performance with increasing network sizes, handling up to 5000 devices without significant loss in accuracy or increased communication overhead.

This scalability makes the framework ideal for large-scale IoT networks, such as industrial IoT systems or smart city infrastructures, where thousands of devices must collaborate efficiently.

5.2. Comparative Evaluation with Existing Approaches

To highlight the advantages of the proposed framework, it was compared against existing state-of-the-art methods, including both centralized and non-federated decentralized cybersecurity systems. Key performance improvements include privacy preservation, detection accuracy, and energy efficiency, as shown in Table 4. These metrics demonstrate that the proposed system reduces energy consumption by 20–30% while achieving an accuracy improvement of 5–7% over traditional methods.

Table 4. Comparison of the Proposed Framework with Existing Methods.

Metric	Centralized Approaches	Non-FL Decentralized Methods	Proposed Framework
Privacy Level	Low	Medium	High
Detection Accuracy	92%	94%	98%
Energy Efficiency	Low	Medium	High

As shown in Table 4, our proposed framework outperforms existing frameworks in terms of privacy level, detection accuracy, and energy efficiency, offering a balanced trade-off between security and operational costs.

Table 5 compares our framework against representative privacy-preserving federated learning baselines, labeled as Baseline A, B, and C. These baselines are derived from common configurations found in the literature and simulate systems using homomorphic encryption (HE), differential privacy (DP), or both. As shown in Table 5, our method achieves comparable or superior performance while minimizing communication and energy overhead. These baselines are informed by performance profiles reported in existing FL research, including communication-efficient training [3], secure aggregation protocols [4], and differential privacy applications in federated settings.

Table 5. Methods Comparison.

Framework	Accuracy	Energy (J)	Communication (MB)	Privacy Method
Baseline A	96.5%	270	110	DP + HE
Baseline B	95.2%	310	125	Differential Privacy
Baseline C	94.0%	320	130	Homomorphic Encryption
Proposed	98.2%	200	85	DP + HE + SMPC

In addition to energy efficiency, key trade-offs between energy consumption and other performance metrics were evaluated. As shown in Table 5, while energy consumption was reduced by 30%, the framework maintains low latency (20 ms) and supports high throughput (up to 5000 devices). These trade-offs demonstrate the framework's ability to optimize resource consumption without compromising performance, making it suitable for large-scale IoT networks.

Summarizing, compared to centralized approaches, the proposed framework achieves higher accuracy while preserving privacy. Centralized methods transmit raw data to a central server, exposing sensitive information to potential breaches. In contrast, the federated learning-based framework ensures that data stays on edge devices, mitigating privacy risks [27]. Previous studies [1,2] have demonstrated similar privacy improvements but often at the cost of detection accuracy or scalability [28]. This research bridges this gap by ensuring both privacy and high detection performance. Moreover, existing solutions, such as blockchain-based security systems, often introduce significant energy and communication overheads, rendering them unsuitable for resource-constrained IoT devices [29]. In contrast, the proposed framework reduces these overheads through optimized aggregation protocols and lightweight model architectures. This reduction in resource consumption is consistent with recent findings on energy-efficient edge AI systems [3,4,29].

5.3. Potential Limitations

Despite its strengths, the proposed framework has certain limitations that warrant further investigation:

Resource Constraints in Ultra-Low Power Devices: Although the framework is designed for resource-constrained environments, ultra-low-power IoT devices with extremely limited processing capabilities may still struggle to implement even lightweight models. Future

work could explore the integration of specialized hardware accelerators (e.g., TPUs, edge AI chips) or tailored model compression techniques to address this challenge.

Latency in Large-Scale Networks: While the framework demonstrates scalability, latency issues may arise in extremely large networks (e.g., over 10,000 devices) due to communication bottlenecks during model aggregation. Employing optimized communication protocols, hierarchical aggregation techniques, or asynchronous update mechanisms could mitigate this issue.

Susceptibility to Advanced Attacks: Although the framework incorporates robust privacy-preserving mechanisms, it remains vulnerable to sophisticated adversarial attacks, such as model poisoning and Byzantine failures. Future research could focus on developing more resilient aggregation protocols, adversarial training strategies, or blockchain-based integrity verification techniques to enhance security.

Dependence on Reliable Infrastructure: The effectiveness of the framework is based on a stable communication infrastructure to transmit updated encrypted models. In scenarios with intermittent connectivity, such as remote IoT deployments or disaster response networks, additional fault tolerance mechanisms (e.g., federated dropout, adaptive synchronization) may be required to maintain performance.

Addressing these limitations in future research could further enhance the robustness and applicability of the framework in diverse deployment scenarios. To evaluate the effectiveness of the framework in real-world infrastructure, additional experiments are necessary to test its performance in large-scale IoT environments like smart city networks. Metrics such as network throughput, model aggregation latency, and device specific performance under stress conditions should be analyzed.

5.4. Future Directions

Integration with Blockchain Technology: Integrating blockchain with the proposed framework could improve the traceability, transparency, and integrity of model updates [30]. Blockchain-based solutions could also facilitate decentralized identity management, further strengthening the security of IoT networks and reducing the reliance on centralized authorities [31].

Quantum-Resistant Cryptographic Techniques: As quantum computing advances, traditional cryptographic methods, including homomorphic encryption, may become vulnerable [32]. Future research should explore the adoption of quantum-resistant algorithms, such as lattice-based cryptography or code-based encryption [33], to ensure long-term security and resilience against quantum attacks [34].

Adaptive Threat Detection Models: Enhancing the adaptability of the framework by integrating metalearning techniques could allow it to detect and respond to emerging cyber threats without requiring extensive retraining [35]. Meta-learning models can generalize across different attack patterns, making them particularly valuable in dynamic IoT environments where threats evolve rapidly [36].

Benchmarking on Real-World IoT Systems: To validate the effectiveness of the framework, future work should extend its evaluation to real-world IoT deployments, such as smart factories, autonomous transportation networks or healthcare monitoring systems. Benchmarking in such environments will provide critical insights into performance, scalability, and reliability under practical operating conditions [37].

6. Conclusions

This study introduces a Federated Learning-Driven Cybersecurity Framework that enhances IoT network security and privacy by integrating federated learning, privacy-preserving encryption, and machine learning-based anomaly detection. The framework

ensures robust threat detection, achieving over 98% accuracy in detecting cyber threats, while maintaining data confidentiality, scalability, and energy efficiency. It reduces energy consumption by 20% compared to centralized models and supports up to 5000 devices. However, challenges remain, such as optimizing the framework for ultra low-power devices, mitigating advanced adversarial attacks, and addressing latency in large-scale networks. Future research will focus on integrating blockchain for traceability, adopting quantum-resistant cryptography, and benchmarking the framework in real-world IoT environments. The findings have broader implications for securing distributed systems beyond IoT, including smart grids, autonomous vehicles, and drone networks, contributing to the development of ethical and sustainable AI-driven cybersecurity solutions.

Author Contributions: Conceptualization, M.R.; methodology, M.R.; software, M.R.; validation, M.R., and A.P.; formal analysis, M.R. and A.P.; investigation, M.R.; resources, A.P.; data curation, M.R. and A.P.; writing—original draft preparation, M.R.; writing—review and editing, M.R. and A.P.; visualization, A.P.; supervision, A.P.; project administration, M.R. and A.P.; funding acquisition, A.P. All authors have read and agreed to the published version of the manuscript.

Funding: This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGen erationEU, partnership on "Telecommunications of the Future" (PE00000001—program "RESTART"—CUP F83C22001690001).

Data Availability Statement: The experimental data results are included in the article. The training dataset is publicly available and is referenced in citation [19].

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A

The following table provides a detailed description of the columns in a network flow analysis dataset. Each row represents a feature of the dataset, organised into three main columns: Column Number, which indicates the numerical index of the column; Column Name, which describes the name of the feature or variable; and Data Type, which specifies the associated data type (e.g., float64 for continuous numerical values and object for strings or categories). This dataset is used to preprocess data and test the federated learning-driven cybersecurity framework for IoT networks.

Column Number	Column Name	Data Type
0	Destination Port	object
1	Flow Duration	float64
2	Total Fwd Packets	float64
3	Total Backward Packets	float64
4	Total Length of Fwd Packets	float64
5	Total Length of Bwd Packets	float64
6	Fwd Packet Length Max	float64
7	Fwd Packet Length Min	float64
8	Fwd Packet Length Mean	float64
9	Fwd Packet Length Std	float64
10	Bwd Packet Length Max	float64

Column Number	Column Name	Data Type
11	Bwd Packet Length Min	float64
12	Bwd Packet Length Mean	float64
13	Bwd Packet Length Std	float64
14	Flow Bytes/s	float64
15	Flow Packets/s	float64
16	Flow IAT Mean	float64
17	Flow IAT Std	float64
18	Flow IAT Max	float64
19	Flow IAT Min	float64
20	Fwd IAT Total	float64
21	Fwd IAT Mean	float64
22	Fwd IAT Std	float64
23	Fwd IAT Max	float64
24	Fwd IAT Min	float64
25	Bwd IAT Total	float64
26	Bwd IAT Mean	float64
27	Bwd IAT Std	float64
28	Bwd IAT Max	float64
29	Bwd IAT Min	float64
30	Fwd PSH Flags	float64
31	Bwd PSH Flags	float64
32	Fwd URG Flags	float64
33	Bwd URG Flags	float64
34	Fwd Header Length	float64
35	Bwd Header Length	float64
36	Fwd Packets/s	float64
37	Bwd Packets/s	float64
38	Min Packet Length	float64
39	Max Packet Length	float64
40	Packet Length Mean	float64
41	Packet Length Std	float64
42	Packet Length Variance	float64
43	FIN Flag Count	float64
44	SYN Flag Count	float64
45	RST Flag Count	float64
46	PSH Flag Count	float64
47	ACK Flag Count	float64
48	URG Flag Count	float64

Column Number	Column Name	Data Type
49	CWE Flag Count	float64
50	ECE Flag Count	float64
51	Down/Up Ratio	float64
52	Average Packet Size	float64
53	Avg Fwd Segment Size	float64
54	Avg Bwd Segment Size	float64
55	Fwd Header Length.1	float64
56	Fwd Avg Bytes/Bulk	float64
57	Fwd Avg Packets/Bulk	float64
58	Fwd Avg Bulk Rate	float64
59	Bwd Avg Bytes/Bulk	float64
60	Bwd Avg Packets/Bulk	float64
61	Bwd Avg Bulk Rate	float64
62	Subflow Fwd Packets	float64
63	Subflow Fwd Bytes	float64
64	Subflow Bwd Packets	float64
65	Subflow Bwd Bytes	float64
66	Init_Win_bytes_forward	float64
67	Init_Win_bytes_backward	float64
68	act_data_pkt_fwd	float64
69	min_seg_size_forward	float64
70	Active Mean	float64
71	Active Std	float64
72	Active Max	float64
73	Active Min	float64
74	Idle Mean	float64
75	Idle Std	float64
76	Idle Max	float64
77	Idle Min	float64
78	Label	object

References

1. IoT Analytics. The State of IoT: 2021 Edition. *IoT Analytics Report*. 2021. Available online: https://iot-analytics.com/product/state-of-iot-summer-2021/ (accessed on 1 July 2025).
2. Ahmed, S.F.; Sharmin, S.; Kuldeep, S.A.; Lameesa, A.; Alam, M.S.B.; Liu, G.; Gandomi, A.H. Transformative Impacts of the Internet of Medical Things on Modern Healthcare. *Results Eng.* **2024**, *25*, 103787. [CrossRef]
3. McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; Arcas, B.A. Communication-Efficient Learning of Deep Networks from Decentralized Data. In Proceedings of the Artificial Intelligence and Statistics, Fort Lauderdale, FL, USA, 20–22 April 2017.
4. Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H.B.; Patel, S.; Ramage, D.; Segal, A.; Seth, K. Practical Secure Aggregation for Privacy-Preserving Machine Learning. In Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, London, UK, 11–15 November 2019.
5. Himdi, T. A Blockchain and AI-Driven Security Framework for Cognitive Cities. *Adv. Artif. Intell. Mach. Learn.* **2024**, *4*, 2908–2925.

6. Altherwi, A.; Ahmad, M.T.; Alam, M.M.; Mirza, H.; Sultana, N.; Pasha, A.A.; Sultana, N.; Khan, A.I.; Alam, M.M.; Azim, R. A Hybrid Optimization Approach for Securing Cloud-Based E-Health Systems. *Multimed. Tools Appl.* **2024**, *84*, 16525–16560. [CrossRef]
7. Shafique, M.; Marchisio, A.; Putra, R.V.W.; Hanif, M.A. Towards energy-efficient and secure edge AI: A cross-layer framework. ICCAD special session paper. In Proceedings of the 2021 IEEE/ACM International Conference on Computer Aided Design (ICCAD), IEEE, Virtual, 1 November–4 November 2021; pp. 1–9.
8. Sinha, A.; Sharma, N.; Kumar, S.; Lande, A.; Iqbal, M.I. AI-Enhanced Living: The Future of Smart Homes. In Proceedings of the 2023 International Conference on Smart Devices (ICSD), Dehradun, India, 2–3 May 2024.
9. Mehta, S.; Khurana, M.; Dogra, A.; Hariharan, S. Advancing IoT Security through Federated Learning: A Comprehensive Approach. In Proceedings of the 2024 3rd International Conference on Applied Artificial Intelligence and Computing (ICAAIC), IEEE, Salem, India, 5–7 June 2024; pp. 561–566.
10. Awosika, T.; Shukla, R.M.; Pranggono, B. Transparency and privacy: The role of explainable ai and federated learning in financial fraud detection. *IEEE Access* **2024**, 12, 64551–64560. [CrossRef]
11. Iacob, A.; Sani, L.; Marino, B.; Aleksandrov, P.; Shen, W.F.; Lane, N.D. Worldwide federated training of language models. *arXiv* **2024**, arXiv:2405.14446.
12. Wei, K.; Li, J.; Ding, M.; Ma, C.; Su, H.; Zhang, B.; Poor, H.V. User-level privacy-preserving federated learning: Analysis and performance optimization. *IEEE Trans. Mob. Comput.* **2021**, *21*, 3388–3401. [CrossRef]
13. Anelli, V.W.; Deldjoo, Y.; Di Noia, T.; Ferrara, A. Towards effective device-aware federated learning. In Proceedings of the AI* IA 2019–Advances in Artificial Intelligence: XVIIIth International Conference of the Italian Association for Artificial Intelligence, Rende, Italy, 19–22 November 2019; Proceedings 18; Springer: Berlin/Heidelberg, Germany, 2019; pp. 477–491.
14. Nayak, J.; Naik, B.; Vimal, S.; Favorskaya, M. *Machine Learning for Cyber-Physical Systems: Advances and Challenges*; Springer: Berlin/Heidelberg, Germany, 2024.
15. Campos, E.M.; Saura, P.F.; González-Vidal, A.; Hernández-Ramos, J.L.; Bernabé, J.B.; Baldini, G.; Skarmeta, A. Evaluating Federated Learning for intrusion detection in Internet of Things: Review and challenges. *Comput. Netw.* **2022**, *203*, 108661. [CrossRef]
16. Demelius, L.; Kern, R.; Trügler, A. Recent advances of differential privacy in centralized deep learning: A systematic survey *ACM Comput. Surv.* **2025**, *57*, 1–28. [CrossRef]
17. Baqer, M. Energy-Efficient Federated Learning for Internet of Things: Leveraging In-Network Processing and Hierarchical Clustering *Future Internet* **2024**, *17*, 4. [CrossRef]
18. Li, S.; Ngai, E.; Voigt, T. Byzantine-Robust Aggregation in Federated Learning. In *IEEE Transactions on Industrial Informatics*; IEEE: New York, NY, USA, 2024.
19. Canadian Institute for Cybersecurity. CICIDS2017 Dataset. 2017. Available online: https://www.unb.ca/cic/datasets/ids-2017.html (accessed on 30 March 2025).
20. Sharafaldin, I.; Lashkari, A.H.; Ghorbani, A.A. Toward generating a new intrusion detection dataset and intrusion traffic characterization. *ICISSp* **2018**, *1*, 108–116.
21. Khan, M.M.; Alkhathami, M. Anomaly detection in IoT-based healthcare: Machine learning for enhanced security. *Sci. Rep.* **2024**, *14*, 5872. [CrossRef] [PubMed]
22. Khan, R.; Taj, S.; Ma, X.; Noor, A.; Zhu, H.; Khan, J.; Khan, Z.U.; Khan, S.U. Advanced federated ensemble internet of learning approach for cloud based medical healthcare monitoring system. *Sci. Rep.* **2024**, *14*, 26068. [CrossRef] [PubMed]
23. Ghadi, Y.Y.; Mazhar, T.; Shah, S.F.A.; Haq, I.; Ahmad, W.; Ouahada, K.; Hamam, H. Integration of federated learning with IoT for smart cities applications, challenges, and solutions. *PeerJ Comput. Sci.* **2023**, *9*, e1657. [CrossRef]
24. Priyadarshini, I. Anomaly detection of iot cyberattacks in smart cities using federated learning and split learning. *Big Data Cogn. Comput.* **2024**, *8*, 21. [CrossRef]
25. Aljunaid, S.K.; Almheiri, S.J.; Dawood, H.; Khan, M.A. Secure and Transparent Banking: Explainable AI-Driven Federated Learning Model for Financial Fraud Detection. *J. Risk Financ. Manag.* **2025**, *18*, 179. [CrossRef]
26. Arora, S.; Beams, A.; Chatzigiannis, P.; Meiser, S.; Patel, K.; Raghuraman, S.; Rindal, P.; Shah, H.; Wang, Y.; Wu, Y.; et al. Privacy-preserving financial anomaly detection via federated learning & multi-party computation. In Proceedings of the 2024 Annual Computer Security Applications Conference Workshops (ACSAC Workshops), Honolulu, HI, USA, 9–13 December 2024; pp. 270–279.
27. Dritsas, E.; Trigka, M. Federated Learning for IoT: A Survey of Techniques, Challenges, and Applications. *J. Sens. Actuator Netw.* **2025**, *14*, 9. [CrossRef]
28. Mohammadi, S.; Balador, A.; Sinaei, S.; Flammini, F. Balancing privacy and performance in federated learning: A systematic literature review on methods and metrics. *J. Parallel Distrib. Comput.* **2024**, *192*, 104918. [CrossRef]
29. Himeur, Y.; Sayed, A.N.; Alsalemi, A.; Bensaali, F.; Amira, A. Edge AI for Internet of Energy: Challenges and perspectives. *Internet Things* **2024**, *25*, 101035. [CrossRef]

30. Ramírez-Gordillo, T.; Maciá-Lillo, A.; Pujol, F.A.; García-D'Urso, N.; Azorín-López, J.; Mora, H. Decentralized Identity Management for Internet of Things (IoT) Devices Using IOTA Blockchain Technology. *Future Internet* **2025**, *17*, 49. [CrossRef]
31. Agarkar, A.A.; Karyakarte, M.; Chavhan, G.; Patil, M.; Talware, R.; Kulkarni, L. Blockchain aware decentralized identity management and access control system. *Meas. Sens.* **2024**, *31*, 101032. [CrossRef]
32. Kumar, M. Post-quantum cryptography Algorithm's standardization and performance analysis. *Array* **2022**, *15*, 100242. [CrossRef]
33. Amirkhanova, D.S.; Iavich, M.; Mamyrbayev, O. Lattice-Based Post-Quantum Public Key Encryption Scheme Using ElGamal's Principles. *Cryptography* **2024**, *8*, 31. [CrossRef]
34. Sibanda, I. The Rise of Quantum-Resistant Cryptography. *IEEE Comput. Soc.* **2024**. Available online: https://www.computer.org/publications/tech-news/trends/quantum-resistant-cryptography/ (accessed on 1 July 2025).
35. Alalhareth, M.; Hong, S.C. Enhancing the internet of medical things (IoMT) security with meta-learning: A performance-driven approach for ensemble intrusion detection systems. *Sensors* **2024**, *24*, 3519. [CrossRef]
36. Rihan, S.D.A.; Anbar, M.; Alabsi, B.A. Meta-learner-based approach for detecting attacks on internet of things networks. *Sensors* **2023**, *23*, 8191. [CrossRef]
37. Prantl, T.; Bauer, A.; Engel, S.; Horn, L.; Krupitzer, C.; Iffländer, L.; Kounev, S. Benchmarking of Secure Group Communication schemes with focus on IoT. *Discov. Data* **2024**, *2*, 5. [CrossRef]