

PRIVACY-PRESERVING FEDERATED LEARNING WITH DIFFERENTIAL PRIVACY AND EDGE AI FOR RESOURCE-CONSTRAINED IOT DEVICES

Shamikh Imran^{*1}, Kinza Khurshid², Zobia Shabeer³, Muhammad Naeem⁴^{*1, 2, 3, 4} Department of Computer Science, Abbottabad University of Science & Technology, Havelian, Khyber Pakhtunkhwa, PakistanDOI: <http://doi.org/10.5281/zenodo.21353633>**Keywords**

Federated Learning, Differential Privacy, Edge AI, Internet of Things, Privacy Preservation, Secure Aggregation, Edge Computing, Distributed Machine Learning.

Article History

Received: 21 November 2025

Accepted: 16 January 2026

Published: 30 January 2026

Copyright @Author

Corresponding Author: *

Shamikh Imran

Abstract

As the Internet of Things (IoT) revolution has swept the world, the data produced is vast and distributed, bringing with it opportunities for smart data analysis and concerns about privacy and security. Conventional distributed and centralized machine learning systems involve sending sensitive data to cloud servers, which can lead to communication delays and data breaches. The problem has been addressed by offering a new approach capable of collaborative model training without sharing raw data, known as Federated Learning (FL). Federated learning is still susceptible to information leakage, poisoning attacks, and resource limitations of edge devices, however. This paper introduces a Privacy-Preserving Federated Learning (PPFL) framework for resource constrained IoT applications utilizing Differential Privacy (DP), Secure Aggregation and Edge AI. The proposed framework allows for distributed model training while maintaining privacy preserving model updates and secure aggregation mechanisms. The model convergence, privacy-accuracy trade-off, communication overhead, energy consumption, resource usage, model robustness, scalability and feasibility of deployment were comprehensively evaluated through experiments. The experimental results show that the proposed framework can guarantee stable convergence with different privacy budgets and provide a good trade-off between privacy preservation and predictive performance. Under relaxed privacy settings, the framework obtains the accuracy of up to 86% in models, while still assuring strong privacy guarantees and resistance to adversarial attacks. Moreover, the communication efficient mechanisms resulted in less overhead in the network, while the evaluations in the edge deployment demonstrated the implementation of the framework with resource constrained devices. The results show that the proposed PPFL framework can be a practical solution for privacy-sensitive IoT applications such as a smart healthcare system, a smart city system, and an industrial IoT system.

1. INTRODUCTION

With the proliferation of connected devices that generate, compute and share data in tremendous amounts, the Internet of Things (IoT) has radically transformed the way this data is created. In recent years, there has been a significant increase in the use of IoT devices in various smart and intelligent systems, including smart

healthcare systems, industrial automation systems, intelligent transportation systems, environmental monitoring systems, and smart city infrastructures [1]. These devices continuously produce large amounts of heterogeneous data that can be used for training intelligent machine learning models for prediction, classification and decision making.

But the well-known machine learning techniques are typically dependent on a centralized data collection, in which data from numerous devices is forwarded to the cloud for the model to be learned. This method introduces many privacy concerns, communication overheads, security concerns and regulatory problems [2].

The valuable insights gained from data collected by the IoT becomes possible with the help of AI, which has become a prominent technology. AI has been integrated with IoT systems, leading to a wide range of smart applications like disease diagnosis, predictive maintenance, anomaly detection, traffic management and more [3]. Though these advantages, centralised AI systems require uniform interaction between edge devices and cloud-based systems, resulting in increased network traffic and data leakage or privacy breaches [4]. Moreover, data protection laws, including the General Data Protection Regulation (GDPR), are very restrictive on how personal information can be collected and shared, adding another difficulty to centralized learning systems [5].

To overcome these challenges, Federated Learning (FL) introduces a new paradigm of distributed machine learning where the local datasets are not placed in a centralized location, but rather are maintained in various local networks by distinct organizations or users [6]. In FL, the clients that participate train a model locally and only send model parameters or gradients to an aggregation server. The server aggregates them together into a global model with aggregation methods like Federated Averaging (FedAvg) [7]. This local storage allows FL to minimize privacy concerns and communication needs, and to realize large-scale distributed intelligence. Thus, the federated learning has attracted a lot of research attention for privacy sensitive application areas such as healthcare, finance, industrial IoT, and smart city applications [8].

While federated learning offers advantages over centralized training in privacy, recent research has shown that privacy cannot be fully guaranteed, as sensitive information can be deduced from model updates to be shared with

others [9]. The gradient leakage attack, model inversion attack, and membership inference attacks are some of the methods that adversarial parties can use to recover private training data from the set of parameters they communicate [10]. Besides, federated learning systems are still susceptible to poisoning attacks, Byzantine attacks, and malicious behaviors of the clients that may affect the integrity of the models and impair the performance of the system [11]. So, relying on federated learning alone will not guarantee full privacy and security for the distributed environment.

Differential Privacy (DP) is a successful technique to reduce privacy leakage in machine learning systems [12]. To prevent this, Differential Privacy adds carefully-tuned noise to updates to the model before they're sent, restricting the kind of information that can be gained about the data records themselves. The privacy guarantee is regulated by the privacy budget ϵ , which means that lower values of ϵ will guarantee more privacy protection, but may degrade the accuracy of the model [13]. Different studies have shown that DP-FL can significantly improve the level of privacy while keeping the predictive performance at an acceptable level [14].

Apart from privacy issues, another problem is communicating efficiently in federated IoT system. Edge devices are resource-constrained devices, which typically have limited computational power, memory, battery life, and network bandwidth [15]. However, frequent interactions between the client and the aggregation server can lead to higher latency and energy costs, which can limit the feasibility of federated learning in real-world scenarios. To overcome these challenges, Edge AI has emerged as a viable solution, allowing machine learning computations to be done on the edge devices instead of relying on cloud resources and communication overheads [16].

Other techniques have been proposed to enhance privacy preservation during aggregation [17] by using Secure Aggregation techniques. They guarantee that only pooled changes to the model will be accessible to the server and don't permit any information leakage while training models

cooperatively. DP, SA, and Edge AI are a complete arsenal for building privacy preserving federated learning frameworks for IoT applications [18].

As much as efforts have been made in privacy-preserving federated learning, there are still several challenges that are yet to be resolved. Existing solutions are mostly related to the protection of privacy and do not consider resource limitations, energy consumption, communication efficiency and feasibility of deployment on edge devices [19]. Moreover, privacy protection and model utility is still a key research challenge, especially for very distributed IoT systems where computing power is scarce [20].

In this work, we present a Privacy-Preserving Federated Learning (PPFL) framework using Differential Privacy, Secure Aggregation and Edge AI for resource-limited IoT devices to tackle these challenges. The proposed framework provides a secure collaborative learning framework without sharing the raw data and makes it difficult to leak data and make adversarial attacks to the model updates. The framework also analyzes the trade-off between varying privacy budgets and model convergence, predictive performance, communication overhead, energy usage, resource usage, and scalability.

The key contributions of this study are as follows:

1. A Privacy-Preserving Federated Learning framework is introduced, combining Differential Privacy, Secure Aggregation and Edge AI for resource-constrained environments in the IoT.
2. To further help protect against information leakage attacks, Differential Privacy is added to local model updates.
3. The Secure Aggregation mechanisms are used to guarantee that individual client contributions remain secure throughout the federated aggregation process.
4. Full-scale experimental assessments are performed to investigate: model convergence, privacy-utility trade-offs, communication overhead, energy consumption, security robustness, and scalability.
5. The suggested framework shows an adequate balance between the privacy protection

and effectiveness of the predictive performance and thus is appropriate for implementation in the privacy-sensitive IoT applications.

The rest of this paper is structured as follows. Section 2 summarizes the previous work in the field of federated learning, DP, and privacy-preserving IoT systems. The methodology and system architecture proposed are described in Section 3. Results and performance evaluation of the experiments is discussed in Section 4. Lastly, Section 5 summarizes the paper and suggests ideas for future research.

2. LITERATURE REVIEW

Therefore, Fed. Learn (FL), a distributed learning paradigm that maintains data privacy while supporting model training for collaboration across different devices, has become a good solution. Thus, Fed. Learn (FL) has been proposed as a distributed learning paradigm that keeps data private and allows for model training collaboratively across multiple devices. As FL is increasingly being used in IoT, tremendous research efforts have been dedicated to privacy protection, efficient communications, security, and resource-constrained deployment.

Nguyen et al. [21] provided an extensive overview of federated learning in IoT, along with key obstacles such as data heterogeneity, communication bandwidth, privacy security, and scalability. The authors highlighted that effective federated learning methods are needed that can work in environments with limited resources. Mothukuri et al. [22] highlighted that there existed several vulnerabilities in FL such as Gradient Leakage, Membership Inference, Model Poisoning, and Byzantine attacks. To make sure there is full privacy protection in federated learning, they claimed it is necessary to add extra security mechanisms, not just the usual ones. Li et al. [23] looked into federated optimization across heterogeneous networks and found that non-identically distributed, often called non-IID data, has a notable effect on model convergence and also on the overall learning performance. In their discussion they essentially underlined the need to build more comprehensive optimization

approaches, for tackling what happens in real-world federated learning settings.

Wen et al. [24] gave a kind of detailed, yet practical overview on what makes federated learning difficult and where it can be used. In that work, several research hurdles were called out that still need to be sorted so federated learning systems can actually be adopted in a wide way, like privacy preservation, communication efficiency, and those computational limits that show up during training. Aledhari et al. [25] went further and talked about the enabling technologies, together with the communications protocols used for federated learning. Their message was pretty clear: using federated learning can reduce privacy risks a lot, mainly because there is no requirement to transfer the raw data to a centralized server. Still, even if sensitive data stays local, there can be a noticeable communications overhead, which is not a small thing. To tackle that privacy angle, Liu et al. [26] proposed a federated learning based, privacy-preserving traffic flow prediction framework. The results suggested that transportation data sensitivity can indeed be protected, while keeping a reasonable level of prediction accuracy. However the study theme stayed mostly on privacy preservation, and it did not really examine resource utilization, nor did it evaluate security robustness.

Wei et al. [27] brought Differential Privacy into federated learning, and well, they also looked at that privacy versus usefulness trade off, like you know. From their experiments it shows DP can really cut down the information leakage, but if you choose a stronger privacy safeguard then, somehow, accuracy tends to drop. To explore FL for Industrial Internet of Things (IIoT) applications, Nguyen et al. [28] investigated FL in IIoT applications. The authors claimed that federated learning can enhance the data privacy, lower the communication expenses and realize distributed intelligence in industrial environments. In wireless communication networks, Chen et al. [29] investigated the use of machine learning techniques. Their findings highlighted the capabilities of intelligent edge computing in lowering latency and enhancing

network efficiency, which will enable the implementation of distributed learning approaches in IoT systems.

For big scale deployments of IoT, Mekki et al. [30] they talked about low power wide area network (LPWAN) technologies, kind of. The research emphasized the need to use energy-efficient communication paths in the case of resource-limited edge devices involved in distributed learning environments. Ali et al. [31] provided a complete review on federated learning algorithms for privacy-preserving smart healthcare systems. Their results revealed that federated learning can safeguard sensitive medical information without compromising privacy, and still allow for collaborative healthcare analytics. But there are issues of communication overhead and computation overhead still to consider. Samarakoon et al. [32] studied distributed federated learning for ultra reliable low latency vehicular communications. The research validated the edge-based FL approach by showing that it can enhance communication efficiency and decrease latency in dynamic vehicular environments.

Lu et al. [33] studied the federated learning with non-IID data distributions and emphasized how federated learning affects the convergence of models and predictive performance. The authors believe that one of the most important challenges in federated learning research is the ability to deal with non-IID data. In a federated learning setting, Muñoz-González et al. [34] introduced an adaptive model averaging Byzantine-robust federated learning approach. They showed that they were more resistant to rogue clients and attacks. The idea was, however, to add to the complexity of the computational process. Raza et al. [35] gave an overview of Low Power Wide Area Network (LPWANs) and highlighted their applications for large scale IoT deployments. They discovered that it is important to have communication efficient protocols to facilitate distributed intelligence in resource-limited settings.

While remarkable advancements have been achieved in federated learning research that is privacy-preserving, some key research challenges

persist. Current research tends to address issues related to privacy, communication efficiency, healthcare applications or attack mitigation. There has been scant work focusing on the integration of Differential Privacy, Secure Aggregation, and Edge AI in a coherent system that also ensures privacy protection, security strength, communication overhead, energy consumption, and deployment viability. This paper, sort of, presents a Privacy-Preserving Federated Learning (PPFL) framework that blends Differential Privacy and Secure Aggregation, plus Edge AI, for distributed learning in an IoT context. It aims to provide a secure, scalable, and resource-efficient approach, even when the devices are all acting separately, without so much overhead really.

3. METHODOLOGY

The current study introduces a novel Privacy-Preserving Federated Learning (PPFL) system in the context of resource constrained Internet of Things (IoT) environments. The system employs Federated Learning (FL), Differential Privacy (DP), Secure Aggregation (SA) and Edge AI concepts in the collaborative model training with no raw data exposition to the system. Initially, the central server deploys the global model to all registered edge devices. The edge devices train this global model on their private datasets and transmit their corresponding model updates to the aggregation server. To maintain privacy, DP mechanism (Gaussian noise addition) is applied on local model updates prior transmitting them towards the central server. Secure aggregation is adopted to transmit the model updates from devices to server in secure way. The server collects all local model updates, aggregates them with the Federated Averaging (FedAvg) model and updates the global model. The process is iterated over multiple rounds of communications until model convergence. The proposed framework is tested on diverse privacy budgets and edge device scenarios. Accuracy, privacy, energy efficiency, communication overhead, device resource usage, and security parameters were computed to quantify the efficacy of proposed method in private sensing applications.

3.1 Overview of the Proposed Framework

We investigate a Privacy-Preserving Federated Learning (PPFL) framework based on Differential Privacy (DP), Secure Aggregation, and Edge AI for secure and efficient learning in a few-resource IoT environment. The framework provides the ability for various edge devices to collectively train a global learning model while never sharing original data, thus protecting user privacy and protecting sensitive data from leakage.

Our design is consisting of one server and several devices. Devices update model locally and submit encrypted and privacy protected model to the server. DP mechanism applied after gradient is produced and before transmission to avoid data leakage in model gradient. Secure aggregation technology is used to secure user data from client update leakage.

3.2 System Architecture

The proposed framework comprises three primary components:

1. Edge Devices
2. Differential Privacy Module
3. Federated Aggregation Server

To start off, the parameter server sends the global model to all distributed client devices. The clients then perform a local training procedure with their own dataset on the global model. When local training is performed, an amount of Gaussian noise is added to the updates following the privacy budget (ϵ). This privacy-preserved update is then sent to the aggregation server where secure aggregation is performed to compute a new global model by aggregating updates from all clients. This procedure is iterated through a number of communication rounds and convergence is expected to take place.

3.3 Federated Learning Process

The federated learning procedure consists of the following steps:

Step 1: Global Model Initialization

The main server starts up a global neural network model, and it spreads the model parameter s to

all the participating edge devices kind of in a rapid way.

Step 2: Local Training

Each edge device trains the received model, using its local dataset for a set number of local epochs. The optimization process runs on its own, without sharing any raw data in common, so it stays separate.

Step 3: Differential Privacy Application

To protect sensitive information, Differential Privacy is kind of applied to the model update(s) before anything gets transmitted. So, Gaussian noise is added to each model parameter, more or less according to:

$$\Delta W' = \Delta W + N(0, \sigma^2)$$

where ΔW represents the local model update and $N(0, \sigma^2)$ denotes Gaussian noise with variance σ^2 .

Step 4: Secure Aggregation

The privacy preserving updates are securely transmitted to the central server. Secure Aggregation, mixes the encrypted updates from all the participating clients without revealing any individual contributions, like you can't tell who did what.

Step 5: Global Model Update

The server updates the global model using the Federated Averaging (FedAvg) algorithm:

$$W_{(t+1)} = \sum(n_k/N)W_{k(t)}$$

where W_k represents the local model parameters of client k , n_k denotes the number of local samples, and N represents the total number of training samples across all clients.

Step 6: Iterative Communication

The modified global model gets shared with all the participating clients, and this iterative process repeats till the condition for stop is fulfilled.

3.4 Differential Privacy Mechanism

Differential Privacy was put in place to keep sensitive user data from being revealed while the model was being trained, kind of like an extra shield. In this setting, the privacy guarantee is tuned by the privacy budget (ϵ), and when ϵ ends up being smaller, the privacy protection gets stronger.

Several privacy budgets were evaluated, including:

- $\epsilon = 0.1$
- $\epsilon = 0.5$
- $\epsilon = 1.0$
- $\epsilon = 10.0$
- $\epsilon = \infty$ (No Privacy)

The impact of different privacy levels on model accuracy and convergence performance was investigated experimentally.

3.5 Edge Device Simulation

To judge how deployable something is in IoT, resource constrained edge devices were kind of simulated, with realistic limits on computing, and all that stuff. Each device was described by, the following traits, roughly, in a way that felt practical:

- CPU utilization constraints
- Memory limitations
- Battery capacity restrictions
- Communication bandwidth limitations

The simulation also tracked energy consumption, communication overhead, and battery depletion during federated training.

3.6 Experimental Configuration

So, the suggested framework got tested with synthetic non-IID federated data, made through a Dirichlet distribution. The experiment parameters, are basically wrapped up in Table 1, just there.

Table 1. Experimental Configuration and Hyperparameter Settings of the Proposed Framework.

Parameter	Value
Number of Clients	10-100
Communication Rounds	20
Local Epochs	3
Batch Size	32
Optimizer	Adam

Privacy Budget (ϵ)	0.1– ∞
Aggregation Algorithm	FedAvg
Noise Mechanism	Gaussian Noise

3.7 Performance Evaluation Metrics

We used a wide collection of performance, privacy, resource, communication, and security measures to see how well the proposed Privacy-Preserving Federated Learning (PPFL) framework actually performs. The metrics we selected weren't random, they were picked because they somehow mirror the framework's ability to predict outcomes while protecting privacy, staying computationally frugal, keeping communication lean, and also being rather resistant to security threats.

3.7.1 Model Performance Metrics

We used usual classification metrics, like Accuracy, Precision, Recall and the F1-Score to check how well the proposed framework works, or well how it predicts. These measures help you do a broad evaluation of the model classification ability and how it actually learns from a distributed dataset.

3.7.2 Privacy Metrics

The Privacy Budget (ϵ) along with Information Leakage Resistance, were used to gauge how well privacy is being preserved. In practice the privacy budget gives the privacy guarantees provided by the Differential Privacy mechanism, and the information leakage resistance is used to judge the ability of the framework to conceal sensitive details in the model updates.

3.7.3 Resource Utilization Metrics

CPU Utilization (%), Memory Consumption (MB), Energy Consumption (mWh), and Battery Utilization (%) were tracked during the whole training process to see whether the framework is actually suitable in a resource constrained IoT setup. The next metrics sort of give us a sense of the compute load as well as the energy price of this approach.

3.7.4 Communication Metrics

The effectiveness of communication was measured using Communication Cost (MB) and

the Number of Active Clients participating in each Communication Round. These measures are basically used to gauge scalability of the framework, and also the communication burden from federated learning operations, which honestly can get kind of heavy as things grow.

3.7.5 Security Metrics

Attack Detection Rate, True Positive Rate (TPR), False Positive Rate (FPR), Precision and Recall are metrics employed for determining the security level of the framework. These metrics are computed to analyze how well the security mechanisms being developed would able to detect and prevent malicious activity of the federated learning system. In essence, these metrics allow for thorough and holistic assessment of the proposed security system regarding performance, efficiency, privacy guarantees, and the security of the system in the federated systems which exist distributedly.

4. RESULTS AND DISCUSSION

We tested and analyzed convergence behavior and prediction performance for the PPFL framework by perturbing different privacy budget to investigate its behavior and performance under restricted resources of IoT. The experiments were run for 20 rounds of communication and the accuracy of the global model was measured at each round. This assessment aimed to explore the effect of the mechanisms for privacy protection on the learning performance of the federated model.

4.1 Model Accuracy Convergence Analysis

The convergence behaviour of the proposed Privacy-Preserving Federated Learning (PPFL) framework for various privacy budgets (ϵ) is shown in Figure 1. As can be seen from the results, the accuracy of the models grow according to the number of communication rounds for all privacy settings as expected. When

no privacy was provided ($\epsilon = \infty$), the final accuracy was the highest, about 86%, whereas the higher the privacy constraints, the lower the accuracy. The results show the intrinsic privacy-utility trade-off in Differential Privacy. The results showed that the model achieved an accuracy of $\sim 83\%$ when $\epsilon = 10.0$, and $\sim 78\%$ when $\epsilon = 1.0$. The model reached the accuracy of about 64%

with strict privacy settings ($\epsilon = 0.1$). Although performance degraded, the proposed framework retained the predictability without compromising user privacy. These results validate that the moderate privacy budget ($\epsilon \approx 1.0$) offers a good trade-off between model usefulness and protecting privacy.




Figure 1. Global model accuracy convergence across communication rounds under different differential privacy budgets.

4.2 Privacy–Accuracy Trade-Off Analysis

As shown in Figure 2, the CIFAR-10, MNIST, and Medical datasets display similar patterns when it comes to model accuracy and privacy budget. Figure 2 shows that the relationship between privacy budget and model accuracy is similar for the CIFAR-10, MNIST, and Medical datasets. The results showed that there was a positive relationship between the classification performance and the privacy budget. For CIFAR-10, accuracy increased from approximately 51% at $\epsilon = 0.01$ to 83% at $\epsilon = 100$. The similar trends

were also noted for MNIST and Medical datasets with an accuracy of about 95% and 87% respectively.

The highlighted practical privacy region ($0.5 \leq \epsilon \leq 2.0$) shows a good privacy guarantee/prediction performance balance. All the datasets in this range were able to maintain competitive accuracy while also maintaining good privacy guarantees. The results confirm the appropriateness of Differential Privacy for various privacy-sensitive IoT applications in the healthcare and smart cities sectors.




Figure 2. Privacy-accuracy trade-off across multiple datasets under varying differential privacy budgets.

4.3 Energy Consumption Analysis

Figure 3, checks the energy needs for the proposed framework across various privacy settings. It looks like energy consumption climbs as the number of participating edge devices increases as shown in Figure 3(a). The high-privacy configuration ($\epsilon = 0.1$) used the most energy because of the cost of privacy-preserving operations. Advances in energy consumption were seen with 50 edge devices, which consumed

about 540 mWh, compared to 470 mWh for non-private use. The energy distribution among components of the system is shown in Figure 3(b). In both cases, local training was the largest contributor of energy use. But Differential Privacy added extra 20% energy overhead as a result of noise generation and privacy-preserving computation. The results suggest that privacy protection can also come at a cost of energy, but it is still affordable for practical IoT deployments.




Figure 3. Energy consumption analysis showing (a) energy scaling with the number of edge devices and (b) component-wise energy distribution.

4.4 Communication Overhead Evaluation

Efficiency of communication is key challenge in Federated Learning systems. As shown in Figure 4, the communication costs vary among different FL algorithms. The communication cost without model compression varied from 40 MB to 52 MB per communication round. The new plan would

need 52 megabytes for each round, thanks to privacy-preserving updates. Communication costs were greatly reduced by applying model compression techniques. The 50% compression brought down communication requirements to about 26 MB per round, and the 75% compression brought them down to almost 13

MB. The results show the effectiveness of communication efficient mechanisms in

mitigating overhead caused by Differential Privacy.




Figure 4. Communication overhead comparison among federated learning algorithms under different compression ratios.

4.5 Resource Utilization Analysis

The resource utilization properties of different edge devices are shown in Figure 5. The CPU utilization results indicate that the CPU load of ESP32 had the highest average CPU load (~75%) while that of Google Coral had the lowest CPU load (~45%). The amount of memory required is different for each device, with the Jetson Nano needing the most memory. We found the

consumption of batteries by the Jetson Nano and the ESP32 to be the fastest and longest respectively, from the battery consumption analysis. The resource efficiency scores show that Google Coral has the highest efficiency, followed by Jetson Nano and Raspberry Pi. The results demonstrate that the choice of hardware can have a profound impact on the feasibility of implementing privacy-preserving Federated Learning models at the edge of the network.




Figure 5. Resource utilization across edge devices including CPU usage, memory consumption, battery behavior, and overall resource efficiency.

4.6 Security Analysis

In Figure 6, the robustness of the proposed framework was assessed in various adversarial attacks. In the absence of defenses, the accuracy of the models was heavily affected by Data Poisoning, Model Poisoning, Byzantine, and Gradient Leakage attacks. The proposed combined defense approach performed best in all the attack cases. For instance, in the Gradient Leakage attacks, the baseline model's accuracy was only 50% while the proposed framework retained around 82% accuracy. The same

improvements were seen for Byzantine and Backdoor attacks. The results of the attack detection further highlight the effectiveness of the proposed security mechanisms. Adopting the combined approach led to:

- True Positive Rate = 92%
- Precision = 90%
- Recall = 88%
- F1-Score = 89%
- False Positive Rate = 5%

The results show that the approach of combining Differential Privacy and Secure Aggregation significantly enhances resilience of the system.




Figure 6. Security performance under multiple attack scenarios and attack detection effectiveness.

4.7 Comparative Analysis with Baseline Methods

Figure 7 offers a detailed comparison of the proposed framework and FedAvg baseline. The proposed model performed better than the baseline model in all the parameters that were used for evaluating. The security performance of the privacy scores improved from 30% to 85%, and the privacy scores went up from 20%. The

effectiveness of the communication rose from 50% to 75% and the energy efficiency rose from 60% to 70%. While the accuracy improvements were relatively small (72% to 76%), the proposed framework provided significantly more privacy and security assurances. The balanced performance of the proposed framework is clearly reflected in the radar chart in terms of various objectives.




Figure 7. Multi-objective comparison between the proposed privacy-preserving federated learning framework and the FedAvg baseline.

4.8 Scalability Analysis

For scalability with the number of clients, the authors studied their proposed framework as

shown in figure 8. Results became more accurate as the number of clients in the federated process grew. With IID data distributions, accuracy rose

from ~73% with 10 clients to 82% with 100 clients. The non-IID datasets had lower performance, but had a similar trend. When the number of clients was expanded from 10 to 100, the convergence time changed from 12 minutes to around 29 minutes. The communication cost

also rose in a linear fashion to about 310MB for 100 clients joining in.

Despite these growths, the proposed framework demonstrated the scalability stability and satisfactory performance on large scale deployment condition.




Figure 8. Scalability analysis showing accuracy, convergence time, and communication cost with increasing numbers of clients.

4.9 Hyperparameter Sensitivity Analysis

The effects of the important hyperparameters on the system performance are analyzed in Figure 9. The learning rate had the largest impact on the accuracy of the model (85%) and the noise multiplier had the greatest impact on the privacy preserved (90%). Local epochs had a major impact on energy usage (80%) and latency (75%).

Likewise, the bounds in the clip moderately affected the communication efficiency and privacy performance. The outcomes indicate that the hyperparameters have to be carefully tuned in order to obtain an optimal balance between privacy, accuracy, energy efficiency, and communication overhead.




Figure 9. Hyperparameter sensitivity analysis showing the impact of system parameters on performance metrics.

4.10 Real-World Edge Deployment Evaluation

The feasibility of the proposed framework in real world edge devices is evaluated in figure 10. Google Coral had the quickest inference time (20ms) and the best model accuracy (87.1%). Jetson Nano also had impressive performance in terms of inference latency at 24ms and accuracy at 85%. The ESP32-S3 was the smallest memory footprint (57 MB), which is great for very resource intensive applications. Arduino Portenta

used the least memory and had the lowest predictive performance. The outcomes show that Google Coral and Jetson Nano are the most suitable platforms for deploying privacy-preserving Federated Learning models, given their performance, latency, and energy efficiency.

Figure 10. Proposed framework deployment performance analysis in the real world on various edge computing platforms.




Figure 10. Real-world deployment performance of the proposed framework across multiple edge computing platforms.

4.11 Discussion of Findings

The experimental results show that the proposed Privacy-Preserving Federated Learning framework can effectively balance the privacy, accuracy, security and resource efficiency. Differential Privacy is an effective solution because it enables the protection of sensitive user information and achieves competitive model performance. Communication compression is used to achieve network overhead reduction and secure aggregation is used to increase adversarial resistance. The proposed framework is overall suitable for the IoT application areas where privacy is a key concern such as healthcare monitoring, smart city, and industrial IoT applications, where distributed learning is an

essential component that must be secure and efficient.

5. CONCLUSION

This paper introduced a Privacy-Preserving Federated Learning (PPFL) framework, which combines Differential Privacy, Secure Aggregation, and Edge AI to facilitate secure and efficient distributed learning in resource-constrained IoT settings. The concept of the proposed framework was to address the following limitations and issues in the context of privacy leakage, communication overhead, security vulnerability at the edge devices, and computation limitation at the edge devices. The effectiveness of the proposed framework in

maintaining a stable convergence of the models for different privacy settings and protecting sensitive information was demonstrated in the experimental evaluation. The results confirmed the idea of privacy-utility trade-off and demonstrated that the higher the privacy (typically the more severe) the worse the prediction performance. The framework was effective however, it was well balanced and provided reasonable trade-off between privacy and model accuracy for reasonable privacy budgets. Moreover, the proposed Secure Aggregation greatly improved against adversarial attacks and the communication-efficient mechanisms minimized the overhead cost of federated learning operations. Resource utilization and energy consumption analysis showed that the framework will be suitable to be applied on resource constrained edge devices. Simultaneously, the experiments on scalability have demonstrated that the proposed approach was able to support the number of client participants in a satisfactory way. The test of the framework in the real world was also verified to be applicable to practical IoT applications. To conclude, the framework proposed in the suggested PPFL provides a broad answer for distributed intelligence from privacy perspective in IoT environment. Further research would be beneficial in the areas of adaptive privacy mechanisms, lightweight cryptographic techniques, blockchain-based federated learning, and deployments in the real world with heterogeneous IoT networks and large-scale industrial applications.

REFERENCES

- Li, S., Xu, L. D., & Zhao, S. (2015). The internet of things: a survey. *Information systems frontiers*, 17(2), 243-259.
- Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., & Ayyash, M. (2015). Internet of things: A survey on enabling technologies, protocols, and applications. *IEEE communications surveys & tutorials*, 17(4), 2347-2376.
- Andrade, R. O., Yoo, S. G., Tello-Oquendo, L., & Ortiz-Garcés, I. (2020). A comprehensive study of the IoT cybersecurity in smart cities. *Ieee Access*, 8, 228922-228941.
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE internet of things journal*, 3(5), 637-646.
- Protection, D. (2018). General data protection regulation. *Intersoft Consulting*, Accessed in October, 24(1).
- McMahan, B., Moore, E., Ramage, D., Hampson, S., & y Arcas, B. A. (2017, April). Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics* (pp. 1273-1282). Pmlr.
- Konečný, J., McMahan, H. B., Felix, X. Y., Suresh, A. T., Bacon, D., & Richtárik, P. (2018). Federated learning: Strategies for improving communication efficiency.
- Yang, Q., Liu, Y., Chen, T., & Tong, Y. (2019). Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2), 1-19.
- Ryffel, T., Trask, A., Dahl, M., Wagner, B., Mancuso, J., Rueckert, D., & Passerat-Palmbach, J. (2018). A generic framework for privacy preserving deep learning. *arXiv preprint arXiv:1811.04017*.
- Wei, W., & Liu, L. (2021). Gradient leakage attack resilient deep learning. *IEEE Transactions on Information Forensics and Security*, 17, 303-316.
- Bagdasaryan, E., Veit, A., Hua, Y., Estrin, D., & Shmatikov, V. (2020, June). How to backdoor federated learning. In *International conference on artificial intelligence and statistics* (pp. 2938-2948). PMLR.
- Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and trends® in theoretical computer science*, 9(3-4), 211-487.

- El Ouadrhiri, A., & Abdelhadi, A. (2022). Differential privacy for deep and federated learning: A survey. *IEEE access*, 10, 22359-22380.
- Truex, S., Baracaldo, N., Anwar, A., Steinke, T., Ludwig, H., Zhang, R., & Zhou, Y. (2019, November). A hybrid approach to privacy-preserving federated learning. In *Proceedings of the 12th ACM workshop on artificial intelligence and security* (pp. 1-11).
- Wang, X., Han, Y., Wang, C., Zhao, Q., Chen, X., & Chen, M. (2019). In-edge ai: Intelligentizing mobile edge computing, caching and communication by federated learning. *Ieee Network*, 33(5), 156-165.
- Mao, Y., You, C., Zhang, J., Huang, K., & Letaief, K. B. (2017). A survey on mobile edge computing: The communication perspective. *IEEE communications surveys & tutorials*, 19(4), 2322-2358.
- Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., ... & Seth, K. (2017, October). Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security* (pp. 1175-1191).
- Kairouz, P., & McMahan, H. B. (2021). Advances and open problems in federated learning. *Foundations and trends in machine learning*, 14(1-2), 1-210.
- Li, H., Ota, K., & Dong, M. (2018). Learning IoT in edge: Deep learning for the Internet of Things with edge computing. *IEEE network*, 32(1), 96-101.
- Wu, Q., He, K., & Chen, X. (2020). Personalized federated learning for intelligent IoT applications: A cloud-edge based framework. *IEEE Open Journal of the Computer Society*, 1, 35-44.
- Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., & Poor, H. V. (2021). Federated learning for internet of things: A comprehensive survey. *IEEE communications surveys & tutorials*, 23(3), 1622-1658.
- Mothukuri, V., Parizi, R. M., Pouriyeh, S., Huang, Y., Dehghantanha, A., & Srivastava, G. (2021). A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115, 619-640.
- Li, T., Sahu, A. K., Zaheer, M., Sanjabi, M., Talwalkar, A., & Smith, V. (2020). Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2, 429-450.
- Wen, J., Zhang, Z., Lan, Y., Cui, Z., Cai, J., & Zhang, W. (2023). A survey on federated learning: challenges and applications. *International journal of machine learning and cybernetics*, 14(2), 513-535.
- Aledhari, M., Razzak, R., Parizi, R. M., & Saeed, F. (2020). Federated learning: A survey on enabling technologies, protocols, and applications. *IEEE access*, 8, 140699-140725.
- Liu, Y., James, J. Q., Kang, J., Niyato, D., & Zhang, S. (2020). Privacy-preserving traffic flow prediction: A federated learning approach. *IEEE Internet of Things Journal*, 7(8), 7751-7763.
- Wei, K., Li, J., Ding, M., Ma, C., Yang, H. H., Farokhi, F., ... & Poor, H. V. (2020). Federated learning with differential privacy: Algorithms and performance analysis. *IEEE transactions on information forensics and security*, 15, 3454-3469.
- Nguyen, D. C., Ding, M., Pathirana, P. N., Seneviratne, A., Li, J., Niyato, D., & Poor, H. V. (2021). Federated learning for industrial internet of things in future industries. *IEEE Wireless Communications*, 28(6), 192-199.

- Chen, M., Challita, U., Saad, W., Yin, C., & Debbah, M. (2017). Machine learning for wireless networks with artificial intelligence: A tutorial on neural networks. *arXiv preprint arXiv:1710.02913*, 9.
- Mekki, K., Bajic, E., Chaxel, F., & Meyer, F. (2019). A comparative study of LPWAN technologies for large-scale IoT deployment. *ICT express*, 5(1), 1-7.
- Ali, M., Naeem, F., Tariq, M., & Kaddoum, G. (2022). Federated learning for privacy preservation in smart healthcare systems: A comprehensive survey. *IEEE journal of biomedical and health informatics*, 27(2), 778-789.
- Samarakoon, S., Bennis, M., Saad, W., & Debbah, M. (2019). Distributed federated learning for ultra-reliable low-latency vehicular communications. *IEEE Transactions on Communications*, 68(2), 1146-1159.
- Lu, Z., Pan, H., Dai, Y., Si, X., & Zhang, Y. (2024). Federated learning with non-iid data: A survey. *IEEE Internet of Things Journal*, 11(11), 19188-19209.
- Muñoz-González, L., Co, K. T., & Lupu, E. C. (2019). Byzantine-robust federated machine learning through adaptive model averaging. *arXiv preprint arXiv:1909.05125*.
- Raza, U., Kulkarni, P., & Sooriyabandara, M. (2017). Low power wide area networks: An overview. *ieee communications surveys & tutorials*, 19(2), 855-873.

