Babylonian Journal of Internet of Things
Vol.2026, **pp**. 73–88
DOI: https://doi.org/10.58496/BJIoT/2026/007; ISSN: 3006-1083
https://mesopotamian.press/journals/index.php/BJIoT

Research Article

Hybrid Privacy-Preserving Federated Learning Framework for Secure IoT Applications Using Differential Privacy and Homomorphic Encryption

Maan Nawaf Abbood [1,*], Zahraa A. Abdalkareem [1]

[1] *Department of Computer science, Al-Imam Al-Adham University College, Baghdad, Iraq.*

ARTICLE INFO

Article History

Received 4 May 2026
Revised 31 May 2026
Accepted 17 Jun 2026
Published 4 Jul 2026

Keywords
Federated Learning,
Internet of Things,
Differential Privacy,
Homomorphic Encryption,
Privacy Preservation,
Secure Aggregation,
Distributed Machine Learning.

ABSTRACT

Internet of Things (IoT) applications have expanded quickly and are producing massive amounts of data from a diverse range of smart devices. Federated Learning (FL) allows collaborative model training without data transfer of the actual images but suffers from privacy concerns as sensitive information might be leaked through the model parameters. Most of the current privacy-preserving FL methods use DP or HE separately which leads to a compromise between privacy protection, model accuracy and computation efficiency. In response, this paper introduces Hybrid Privacy-Preserving Federated Learning (HPPFL), a comprehensive privacy-preserving learning system that combines Differential Privacy (DP) and Homomorphic Encryption (HE) for secure IoT applications. The proposed design allows for local training of the model on IoT devices, followed by the addition of noise (adaptive Differential Privacy) to the model updates, then encrypting the noisy model using Paillier Homomorphic Encryption before sending it to the federated server. The model updates are then securely aggregated with FedAvg without making sensitive information public and the aggregated model is sent back to the global model for next training round. TON-IoT was used to test the proposed framework, and it was compared to the conventional FL, DP-FL, and HE-FL approaches based on classification accuracy, precision, recall, F1-score, communication cost, encryption overhead, and training time. The experimental results demonstrate that the proposed HPPFL framework achieved classification accuracy of 97.46% compared to other FL frameworks, conventional FL, DP-FL, HE-FL with the highest level of privacy protection and moderate computational overhead. The results indicate that the proposed framework is capable of preserving privacy, ensuring effective communication and maintaining the prediction performance, which shows it is a potential solution to secure large-scale IoT applications.

1. INTRODUCTION

The Internet of Things (IoT) is one of the most significant technologies for the digital transformation happening in many industries, such as healthcare, smart cities, industrial automation, intelligent transportation, agriculture, environmental monitoring, and more. The constant proliferation of the interconnected sensor, smart device and edge computing platforms has allowed the collection and processing of massive amounts of data, covering a wide range of distributed data, to make intelligent decisions and deliver real-time services [1,2]. As of recent, it is reported that there will be a number of connected IoT devices that passes several tens of billions within the next few years, causing an explosion in the amount of data generated and communicated with IoT devices in the near future [3]. This growth is an excellent opportunity for intelligent services, but it also presents some challenges in regards to data security, privacy of users, and computational efficiency.

The traditional machine learning (ML) system is known as centralized architectures where raw data from distributed IoT devices is sent to cloud server for model training. Centralized learning, as shown in many applications, has achieved excellent results, but it has serious privacy issues as sensitive user data is required to be sent and stored at the centralized locations. These architectures create an added risk of data leakage, unauthorized access, and cyberattacks, as well as a violation of increasingly strict data protection policies, such as the General Data Protection Regulation (GDPR) and other privacy-preserving policies [4,5]. Moreover, because the amount of information produced by IoT devices to the centralized servers is large, many resource-limited IoT settings are not appropriate for centralized learning [6].

Recently, a new paradigm for distributed ML, called Federated Learning (FL), has come to the fore which aims to collaboratively train a model while not moving raw data away from local devices [7,8]. Rather than communicating

*Corresponding author. Email: maan.alani@imamaladham.edu.iq

sensitive datasets, the IoT device trains its own model locally and periodically uploads only the model parameters/gradients to the central aggregation server for aggregation of the global model, for example, using Federated Averaging (FedAvg). This decentralized learning strategy substantially reduces privacy risks, minimizes communication of raw data, and supports compliance with modern data protection regulations while maintaining competitive learning performance [9]. Hence, Federated Learning is emerging as a promising approach for privacy-preserving applications in the healthcare sector, in the Industrial Internet of Things (IIoT), the automotive industry, smart electric grids, and edge computing. [10,11] While Federated Learning offers these benefits, it is not a foolproof solution to privacy concerns. In recent years, it has been shown that adversarial attacks like gradient leakage, model inversion, membership inference and reconstruction attacks can be used to inadvertently expose sensitive information from the model updates exchanged [12,13]. Even if users' data are not transmitted, adversaries can exploit shared gradients or model parameters to deduce confidential user data. These vulnerabilities severely undermine the privacy assurances provided by standard Federated Learning designs, and make it crucial to develop robust privacy-preserving methods that can secure local data as well as model parameters being transmitted [14].

To address these challenges, several different privacy-preserving methods have been integrated into Federated Learning. Differential Privacy (DP) preserves sensitive data by adding carefully calibrated random noise to the updates of models before they are transmitted, this will reduce the amount of information that can be learned about individual training samples [15,16]. But, too loud noises can cause a poor convergence of the model and poor classification accuracy. Instead, Homomorphic Encryption (HE) allows performing mathematical operations directly on the encrypted model parameters, i.e. without decrypting the parameters being transmitted for aggregation, during which the parameters are secured [17,18]. While HE offers robust cryptographic security, it can also result in significant computational and communication costs, especially in large-scale IoT deployments where processing power is limited [19]. As a result, current methods that use either Differential Privacy or Homomorphic Encryption usually have to sacrifice either privacy or efficiency or accuracy.

In response to these challenges, this paper suggests a new Hybrid Privacy-Preserving Federated Learning (HPPFL) system combining Differential Privacy and Homomorphic Encryption into a cohesive Federated Learning architecture for secure IoT applications. In the proposed framework, local model training is carried out before model updates are perturbed using Differential Privacy by each IoT device. The privacy protected parameters are then encrypted using Homomorphic Encryption and securely sent to the federated server to perform encrypted aggregation without revealing any sensitive information. The proposed framework seeks to offer robust privacy protection while retaining high learning accuracy and communication efficiency through a combination of statistical privacy protection and cryptographic security.

The key goals of this research are to design a secure hybrid privacy-preserving Federated Learning framework for distributed IoT environments, integrate Differential Privacy and Homomorphic Encryption in the same collaborative learning process, minimise the leakage of privacy in the model aggregation process, maintain high model accuracy even after incorporating privacy enhancement techniques, and evaluate the proposed framework on representative IoT datasets and comprehensive performance metrics. The contributions of this work can be summarized as follows:

1. Development of HPPFL framework that jointly integrates Differential Privacy and Homomorphic Encryption.
2. Design of a secure aggregation mechanism that protects model updates without exposing sensitive information during federated communication.
3. Enhancement of privacy preservation while maintaining competitive model accuracy and communication efficiency in distributed IoT environments.
4. Comprehensive experimental evaluation using multiple performance metrics, including classification accuracy, precision, recall, F1-score, communication overhead, encryption cost, training time, and privacy preservation effectiveness.

Comparative performance analysis against conventional Federated Learning, Differential Privacy-based FL, and Homomorphic Encryption-based FL approaches.

2. RELATED WORK

Federated Learning (FL) has revolutionized distributed machine learning by allowing local devices to collaborate in model training without sharing raw data. The decentralized learning paradigm makes FL a potentially good solution to the IoT applications, where privacy protection, communication efficiency, and distributed intelligence are crucial requirements. However, privacy leakage due to transmitted model parameters continues to be a research challenge that lends impetus to incorporating cutting-edge privacy-preserving methods like Secure Aggregation, Homomorphic Encryption and Differential Privacy. This section summarizes the most relevant works pertaining to Federated Learning in IoT and privacy preserving approaches, as well as hybrid secure learning frameworks.

2.1 Federated Learning for IoT

Recently, Federated Learning emerged as one of the most promising distributed learning paradigms for IoT applications where local data needs to be preserved while models are collaboratively trained. FL is different from traditional centralized machine learning methods because it enables the local model training on the IoT devices and only updates are sent to the central aggregation server, thereby lowering the communication overhead and lessening the exposure of sensitive information [20].

There have been several studies that have shown the effectiveness of FL in different fields of IoT like healthcare, industrial automation, intelligent transportation and smart cities. These techniques are very effective to protect data privacy without sharing the raw data, but can still lead to privacy leakage through transmitted gradients and model parameters, particularly when facing advanced inference attacks [21,22].

2.2 Privacy-Preserving Techniques in Federated Learning

To resolve privacy issues in traditional Federated Learning, researchers have applied several privacy-preserving techniques to the learning process. Differential Privacy is one extremely widely used method that achieves this goal by adding calibrated random noise to local model updates before sending them to the central server [23]. While DP can help to reduce information leakage, too much noise might hinder model convergence and prediction accuracy. A different cryptographic approach is proposed as Homomorphic Encryption (HE), which enables directly computing arithmetic operations on the encrypted model parameters without decryption. As a result, aggregation servers can derive more global models without compromising the data confidentiality all throughout the learning process [24]. But the computational complexity and communication overhead caused by HE are still a big challenge in resource-limited IoT devices [25].

Secure Aggregation protocols have been proposed for protecting model updates during the communication phase of the algorithm, whilst allowing accurate model construction at the server while the server does not have access to the individual client updates. These mechanisms will further reinforce the privacy assurances of FL systems that work in decentralized IoT settings [26].

2.3. Hybrid Privacy-Preserving Federated Learning

In recent years, some researchers have explored hybrid approaches of incorporating multiple security measures into Federated Learning architecture. The idea of hybrid approaches is to retain the complimentary benefits of both DP and HE and to provide a superior level of privacy protection while minimizing losses in model accuracy. Studies that combine DP with HE has shown enhanced robustness against inferences attacks while ensuring good learning performance in distributed IoT systems in recent times [27].

Likewise, adaptive privacy budget allocation strategies have been suggested for mitigating the loss of accuracy due to static DPPs. However, secure aggregation with local privacy has been included in other works for enhancing communication security and reliability of collaborative learning [28]. Despite all these improvements, current approaches are mostly concerned with either statistical privacy or cryptographic protection, namely separately. A few studies have been conducted that simultaneously optimize privacy preservation, communication efficiency, computational complexity and model accuracy in a unified Federated Learning architecture that suits the scale of the IoT environment.

2.4. Comparative Analysis of Existing Studies

Table 1 summarizes representative studies on privacy-preserving Federated Learning for IoT applications.

TABLE I. COMPARATIVE ANALYSIS OF EXISTING PRIVACY-PRESERVING FEDERATED LEARNING APPROACHES

Ref	Method	Privacy Technique	Advantages	Limitations
[16]	Differential Privacy-based FL	Differential Privacy	Strong statistical privacy	Accuracy degradation due to injected noise
[17]	Homomorphic Encryption-based FL	Homomorphic Encryption	Secure encrypted aggregation	High computational overhead
[26]	Secure Aggregation FL	Secure Aggregation	Protects model updates during communication	Additional communication complexity
[21]	Hybrid DP-FL	Differential Privacy + Adaptive Privacy Budget	Improved privacy-accuracy balance	Increased implementation complexity
[25]	HE-based Collaborative FL	Homomorphic Encryption	High confidentiality during aggregation	Increased encryption latency
Proposed	HPPFL	Differential Privacy + Homomorphic Encryption	Enhanced privacy, secure aggregation, and high learning accuracy	Moderate computational overhead

The comparative analysis indicates that existing approaches generally focus on a single privacy-preserving mechanism or emphasize either privacy enhancement or computational efficiency. Very few studies integrate multiple privacy-preserving techniques within a unified Federated Learning framework designed specifically for secure IoT environments.

2.5 Research Gap

While significant strides have been made in Federated Learning (FL) for privacy, some research challenges still need to be addressed. Most of the current works use either Differential Privacy or Homomorphic Encryption alone which makes it difficult to balance the level of privacy, the accuracy of the model, and computational efficiency. Besides, most of the frameworks do not consider the synergy effect of secure aggregation, encrypted communication and adaptive privacy protection in the large-scale IoT. Thus, a single solution is needed to ensure the protection of data at the local level while also securing the model's updates, ensuring the accuracy of the learning process, and properly supporting scalable distributed IoT applications.

To overcome the above limitations, this paper proposes a new Hybrid Privacy-Preserving Federated Learning (HPPFL) framework which integrates DP and HE in a single secure aggregation architecture. The proposed framework is designed to strike a balance between privacy preservation and the efficiency of communication and prediction performance for secure collaborative learning in distributed IoT systems.

3. PROPOSED METHODOLOGY

3.1 Overall Architecture

This proposed Hybrid Privacy-Preserving Federated Learning (HPPFL) framework aims to offer secure and privacy-preserving collaborative learning for distributed Internet of Things (IoT) applications. Differential Privacy (DP) along with Homomorphic Encryption (HE) are implemented in a single Federated Learning framework to ensure that sensitive information is protected during the entire learning process. The proposed framework is an alternative to the traditional centralized machine learning methods, which allows the training of local models on local IoT devices while maintaining data locality, mitigating privacy threats and communication load.

The proposed framework has six sequential stages as shown in figure 1. Firstly, IoT devices gather local data from sensors and smart devices, and do not share raw data with external servers. The local machine learning models are independently learned by each device from its own dataset of data. Differential Privacy is applied to the locally trained model parameters before sending them out for the global update, adding calibrated noise to reduce the chance of extracting sensitive information from the shared updates. The parameters of the model are then anonymized, and encrypted using Homomorphic Encryption, enabling secure communication and encrypted computation during the aggregation process.

Each local model update is encrypted and then sent to the federated server, which securely aggregates all the parameters to form a global model without the need to access individual client information. The resulting global model is then sent back to the individual IoT devices for aggregation and the next communication round continues until convergence is reached. This collaborative learning process not only maintains user privacy, but also maintains high prediction accuracy and facilitates distributed learning in heterogeneous IoT environments with scalability. The proposed architecture has several benefits compared to conventional Federated Learning systems. The Differential Privacy feature is the first to guard against inference attacks by obscuring sensitive model updates. Second, Homomorphic Encryption ensures privacy of communication and aggregation without divulging any model parameters. Lastly, the combination of both techniques is a balance between privacy protection, learning accuracy, and communication efficiency, making the proposed framework suitable to secure large-scale IoT applications.

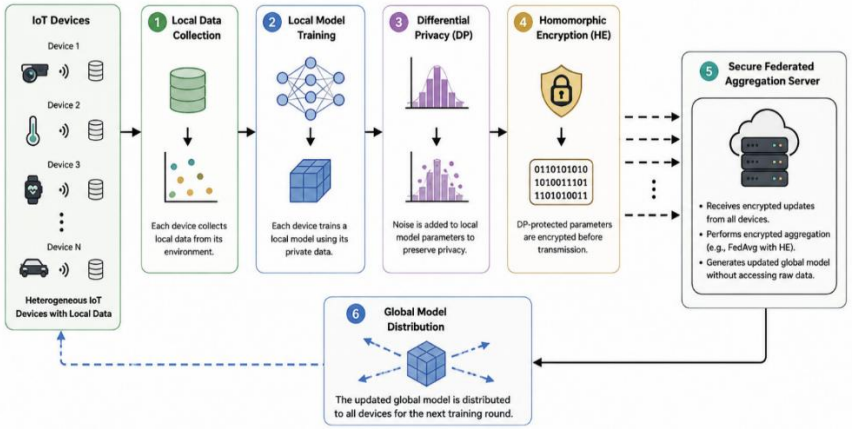


Fig. 1. Overall Operational Architecture of the Proposed HPPFL Framework

The overall architecture of the proposed HPPFL framework showing the entire workflow of Federated Learning including local data collection, local model training, Differential Privacy protection, Homomorphic Encryption, secure federated aggregation, global model generation, and model redistribution to IoT devices.

3.2 System Model

The proposed Hybrid Privacy-Preserving Federated Learning (HPPFL) framework is for distributed Internet of Things (IoT) setting, where many heterogeneous devices can jointly train a machine learning model without compromising the privacy of the locally generated data. The system model comprises four major modules: IoT devices, Edge nodes, Cloud infrastructure, and Federated Learning Server as shown in Figure 2. These are complementary and work together to facilitate secure distributed learning without sharing raw data off of local devices. The data owners in the proposed framework are represented by IoT devices. Smart sensors, wearable health related devices, surveillance cameras, industrial monitoring devices, smart home appliances, and autonomous vehicles are examples of these devices. All IoT devices are constantly sensing what is happening around them and storing the data within them. Then the devices will train their local model with their own data, which means that the risk of privacy leakage is greatly minimized without sending any raw data to external servers.

Edge nodes are deployed near the IoT devices to lessen communication latency and computation load. Edge computing offers intermediate computing resources that help local devices in model training, temporary storage, and communications management. The edge nodes collect local communication traffic, arrange the participating clients and pass on the model update to the federated server in a secured way. This stacked structure helps reduce network congestion and enhances scalability, especially for large-scale IoT deployments. Long-term model management, storage, and system monitoring are achieved using the Cloud infrastructure. The cloud does not directly access the user data, but rather stores the successive versions of the global model and coordinates the communication among distributed edge nodes and the federated server, while maintaining the global learning process. Sensitive user information is protected during the learning process as only the model parameters are exchanged in encrypted form.

The heart of the framework is the Federated Learning Server that controls the collaborative learning process for all participating IoT devices. In every communication round, the server collects encrypted model parameters from distributed clients, executes secure model aggregation with Federated Averaging (FedAvg) algorithm and produces a new global model. The model is then combined and sent back to the participating devices for the next training round. In all of this, the federated server and the cloud infrastructure cannot access the raw learning data, which is compliant with the principles of privacy-preserving learning. These four constitute a robust distributed learning system that can meet the requirements of large-scale IoT applications while simultaneously maintaining privacy requirements, communication efficiency, and computational scalability.

TABLE II. DESCRIPTION OF THE MAIN COMPONENTS OF THE PROPOSED HPPFL SYSTEM

Component	Primary Function
IoT Devices	Collect local data and perform local model training without sharing raw data.
Edge Nodes	Coordinate nearby IoT devices, reduce communication latency, and forward encrypted model updates.
Cloud Infrastructure	Manage global system resources, store global models, and coordinate distributed learning.
Federated Learning Server	Securely aggregate encrypted local model updates and generate the global model using FedAvg.

To explain detailed modules of the proposed framework, an overall system model that shows the interaction of the main architectural components is presented. The proposed HPPFL framework is based on a hierarchical structure of IoT devices, edge nodes, federated learning server and cloud infrastructure, as illustrated in figure 2. This approach allows for secure collaborative learning to be achieved by keeping raw data on local devices and only sending updates to the protected model during the federated learning process.

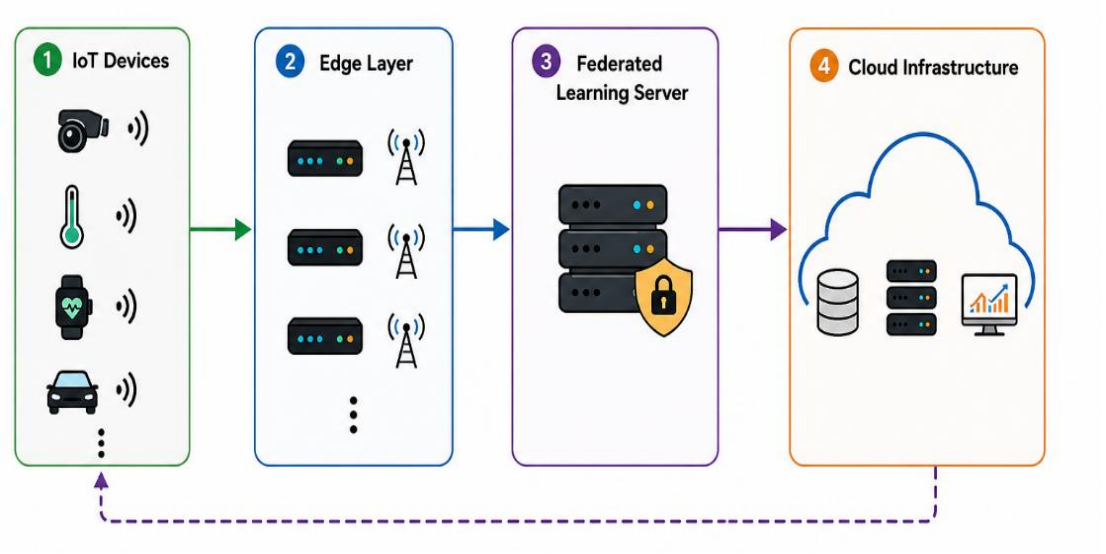


Fig. 2. System Model of the Proposed HPPFL Framework

3.3 Local Model Training

In the proposed Hybrid Privacy-Preserving Federated Learning (HPPFL) framework, local model training is the initial learning phase. Each participating IoT device generates a local machine learning model based on its own private dataset. Each participating IoT device trains a local machine learning model, without sharing raw data with a centralized server. It is a decentralized learning method which minimizes privacy threat and allows knowledge extraction from distributed data sources collaboratively.

A Multi-Layer Perceptron (MLP) model is selected as the local learning model in the proposed framework, because it has low memory requirements and is economical in terms of computation with structured IoT data. The MLP also offers a good trade-off between prediction accuracy and computational complexity, making it suitable for IoT devices with limited resources. The local MLP model is initialized with the most recent global parameters sent from the federated server for each IoT device. The model is then trained locally for several epochs using the private data set of the device by the standard forward propagation and backpropagation procedure. When the model is trained, it tries to reduce a fixed loss function with gradient optimization. After the local training phase, the model parameters generated during the training are saved but the training data are never deleted from the local device.

The federated server receives the local model updates, which are then passed through two successive privacy preserving operations before they reach the federated server. Differential Privacy first perturbs the model parameters by injecting random noise that is calibrated to do so. The parameters are then encrypted with Homomorphic Encryption for secure and confidential aggregation, and the privacy-protected parameters are sent to the encoder. The privacy-protected parameters are then sent to the encoder, which encrypts them using Homomorphic Encryption. Consequently, only encrypted model updates are sent to the federated server and the raw data as well as the unprotected model parameters are not sent. The overall local training procedure adopted in the proposed framework is summarized in Figure 3 and, Table 3 present local MLP configuration.

TABLE III. LOCAL MLP CONFIGURATION

Parameter	Value	Description
Model Type	Multi-Layer Perceptron (MLP)	Local learning model deployed on each IoT client
Input Layer	Dataset Features	Number of neurons equals the number of input features
Hidden Layers	2	Fully connected hidden layers
Hidden Neurons	128, 64	Number of neurons in the first and second hidden layers
Activation Function	ReLU	Nonlinear activation function for hidden layers
Output Layer	Softmax	Multi-class classification output
Loss Function	Categorical Cross-Entropy	Loss function used during local training
Optimizer	Adam	Gradient-based optimization algorithm
Learning Rate	0.001	Initial learning rate
Batch Size	32	Number of training samples processed in each iteration
Local Epochs	5	Number of local training epochs per communication round
Weight Initialization	Xavier (Glorot)	Initialization method for network weights

The chosen MLP architecture achieves a good trade-off between prediction accuracy and computational demands, making it appropriate for use in resource-limited IoT devices that engage in federated learning. To balance between local computational load and learning ability, a lightweight structure comprises two hidden layers was adopted. The Adam optimizer and ReLU activation function ensure stable convergence in local training, while the small number of local epochs minimize the communication latency and synchronization overhead between rounds of federated learning.

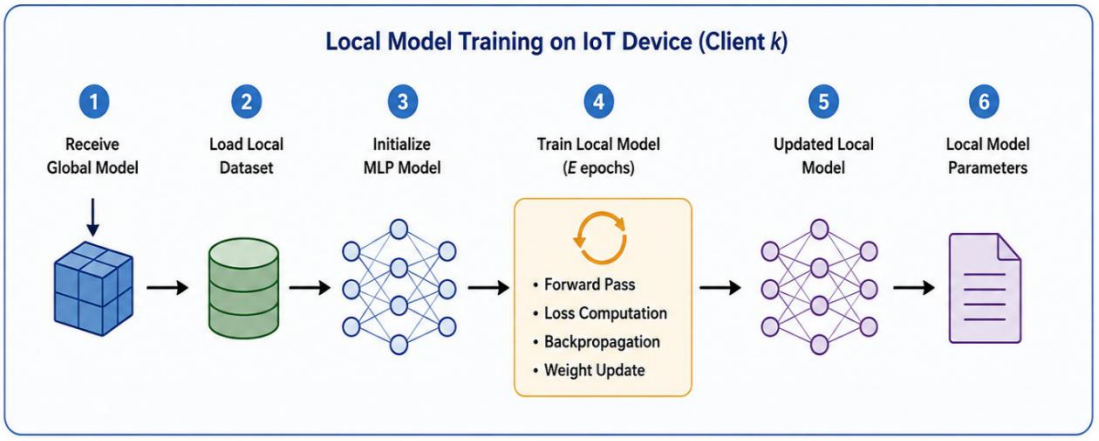


Fig. 3. Local Model Training Process

3.4 Differential Privacy Module

The Differential Privacy (DP) module is the first layer of privacy protection in the proposed Hybrid Privacy-Preserving Federated Learning (HPPFL) framework. While the raw data is not directly shared among the parties, model parameters shared in Federated Learning could still contain sensitive information due to gradient leakage, model inversion, or membership inference attacks. The proposed framework addresses these risks by applying Differential Privacy to prevent information leakage in model updates when transmitting models while maintaining effective learning. Once the local model is trained on each IoT device, the updated model parameters are sent to the Differential Privacy module for processing. Carefully designed noise is used to modify the model updates to prevent too easy reconstruction of the private training information, rather than passing the original parameters. As such, the updates sent are inadequate to provide any substantial information about the original local data set even if intercepted by an enemy.

The difference between the conventional Differential Privacy approaches and the proposed approach is that the conventional ones use a fixed privacy budget while the latter uses an adaptive privacy budget allocation strategy. The privacy budget is dynamically adjusted, depending on the properties of each communication round, to enable an effective balance between privacy protection and model accuracy. In the initial training phases, more privacy protection is used to ensure the privacy of sensitive information, and in later rounds, there is a progressive decrease in the noise added to the signal in order to converge models more rapidly and improve prediction accuracy. The Differential Privacy process used in the suggested framework is carried out in 4 consecutive steps:

a) Local model training on each IoT device.
b) Calculation of local model parameter updates.
c) Adaptive noise injection based on the assigned privacy budget.
d) Generation of privacy-protected model parameters for encryption.

The Differential Privacy module generates a protected model update that retains enough information for collaborative learning of a model while simultaneously providing statistical privacy of the local training data. These protected parameters are then passed to the Homomorphic Encryption module, which is used to encrypt these parameters before sending them to the federated server.

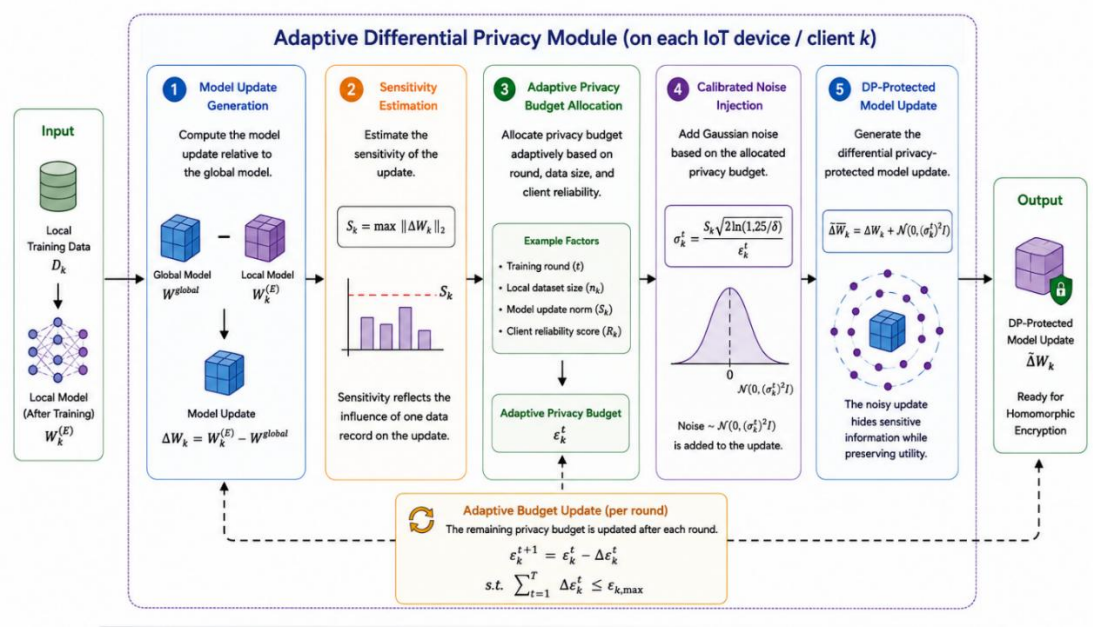


Fig. 4. Adaptive Differential Privacy Module

3.5. Homomorphic Encryption Module

The second privacy preserving layer of the proposed Hybrid Privacy-Preserving Federated Learning (HPPFL) framework is the Homomorphic Encryption (HE) module. The Differential Privacy module is responsible for adding carefully calibrated statistical noise to the local model updates, while the Homomorphic Encryption module adds another layer of cryptography to the transmission and aggregation of local model updates. The dual protection strategy guarantees that sensitive model parameters stay confidential throughout the collaborative learning process.
The addition homomorphic property of the Paillier Homomorphic Encryption scheme is used in the proposed framework, which allows for secure aggregation of model parameters used for each client without revealing individual client updates. The FedAvg algorithm does majority of addition operations when aggregating the global model in Federated Learning, making Paillier encryption well suited for this scenario. Moreover, the scheme offers a good combination of efficiency and security, making it suitable for resource-limited IoT devices.

Once the updates to the model are generated, each participating IoT device encrypts them with a Homomorphic Encryption scheme before sending them. As a result, intermediate communication nodes are unable to access the model parameters in plain text, nor is the federated server. Rather, only the encrypted model update will be sent and exchanged, minimizing the possibility of information leakage during the communication. A significant benefit of Homomorphic Encryption is that arithmetic operations can be performed directly on an encrypted data. This property allows the federated server to combine aggregated encrypted model updates without decrypting the models, thus ensuring that the models remain confidential during the aggregation process. After secure aggregation, every entity authorized to decrypt the global model will receive a decrypted version of the model and then redistribute it to other participating IoT devices for the next round of federated learning. The Homomorphic Encryption module of the proposed framework consists of 4 stages of sequential processing.

a) Reception of Differential Privacy-protected model updates.
b) Encryption of local model parameters using the Homomorphic Encryption algorithm.
c) Secure transmission of encrypted updates to the federated learning server.
d) Privacy-preserving aggregation of encrypted model updates prior to global model generation.

The proposed framework preserves the statistical privacy of the training data and the cryptographic confidentiality of the transmitted model parameters by incorporating both DP and HE in a single federated learning framework. This approach is integrated and can greatly enhance the overall security of the collaborative learning process without compromising the convergence of models and prediction accuracy in distributed IoT environments.

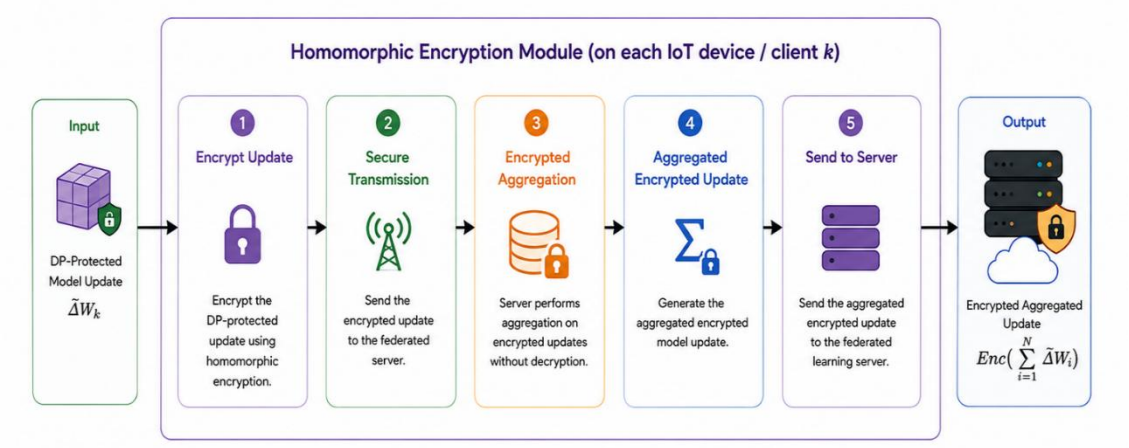


Fig. 5. Homomorphic Encryption Module.

3.6. Secure Aggregation

Once locally protected with Differential Privacy and encrypted with Homomorphic Encryption, the locally-updated model is sent to the federated learning server to ensure secure aggregation. The proposed HPPFL framework aggregates the model updates in a cryptographic manner instead of aggregating the plaintext model parameters as in conventional Federated Learning. This allows the federated server to build the global model without obtaining any client parameters, which maintains the privacy of each client (IoT device).

The proposed framework is based on the FedAvg algorithm as it is simple, easily scalable, and effective in distributed learning environments. In each communication round, the server gets encrypted updates of the local model from all clients involved, then merges them to create a new global model. Because Homomorphic Encryption allows the arithmetic operations on the ciphertext, ciphertext aggregation is done directly on encrypted parameters, without having to decrypt them in the middle. Secure aggregation procedure is divided into 4 steps.

a) Reception of encrypted local model updates from participating IoT devices.
b) Homomorphic aggregation of encrypted parameters using the FedAvg strategy.
c) Generation of the encrypted global model.
d) Authorized decryption and redistribution of the updated global model to all participating clients.

The proposed framework combines Federated Averaging with Homomorphic Encryption, thus protecting the model parameters from being accessed by the federated server, but maintaining the learning capability of Federated Learning. This secure aggregation method significantly mitigates the potential for information leakage in the aggregation process and ultimately improves the privacy and integrity of distributed IoT learning systems.

3.7. Global Model Update

After the secure aggregation, the federated learning server creates a new global model by aggregating the encrypted local model updates from all the IoT devices. The proposed HPPFL framework is based on the FedAvg approach to incorporate the encrypted client models without compromising the data confidentiality during the aggregation process. The global model is built with Homomorphic Encryption technique, which allows to perform computations on the ciphertexts without revealing individual client parameters. Once the aggregation is done the global model gets decrypted by only an authorized entity with the corresponding secret key. The decrypted global model is then sent to all the participating IoT devices via the edge infrastructure and overwritten by the local model, which is the initialization model for the next round of federated learning. This process is repeated until a certain convergence criterion or a certain maximum number of communication rounds is fulfilled.

The global model update mechanism facilitates seamless and ongoing collaborative learning, without any raw training data ever leaving local devices. Moreover, the proposed framework ensures the statistical privacy of local model updates and

the confidentiality of the transmitted parameters, offering secure model synchronization in distributed IoT environments by integrating Differential Privacy and Homomorphic Encryption.

3.8. Proposed Hybrid Algorithm

The algorithm proposed in the HPPFL unifies local model training, Differential Privacy, Homomorphic Encryption, secure federated aggregation and global model updating into one privacy-preserving learning process. The algorithm prevents raw IoT data from being sent to the federated server, and only protected and encrypted model updates are sent.

Algorithm 1. Proposed Hybrid Privacy-Preserving Federated Learning Algorithm

Input: IoT clients (K), local datasets D_K, initial global model (W^0), number of rounds (T), privacy budget (ϵ)

Output: Final global model (W^T)

Begin
1. Initialize the global model W (0) at the federated server.
2. For each communication round t = 1 to T do:
3. Send the current global model W(t) to selected IoT clients.
4. For each client k do:
5. Train the local model using private dataset Dk.
6. Generate local model update ΔWk.
7. Apply Differential Privacy to ΔWk.
8. Encrypt the DP-protected update using Homomorphic Encryption.
9. Send the encrypted model update to the federated server.
10. End For
11. Perform secure aggregation using FedAvg on encrypted updates.
12. Decrypt the aggregated global model by an authorized entity.
13. Update the global model W(t+1).
14. Distribute the updated global model to all participating clients.
15. End For
Return final global model W(T).
End

The proposed HPPFL framework is summarized by this algorithm, which ensures that local updates are statistically private, and model transmission and aggregation are secure. To give a visual picture of the proposed algorithm, the whole execution flow of the Hybrid Privacy-Preserving Federated Learning (HPPFL) framework is given in Figure 6. The flowchart illustrates the series of operations carried out at every communication round, from local model training to the creation and sharing of the new global model.

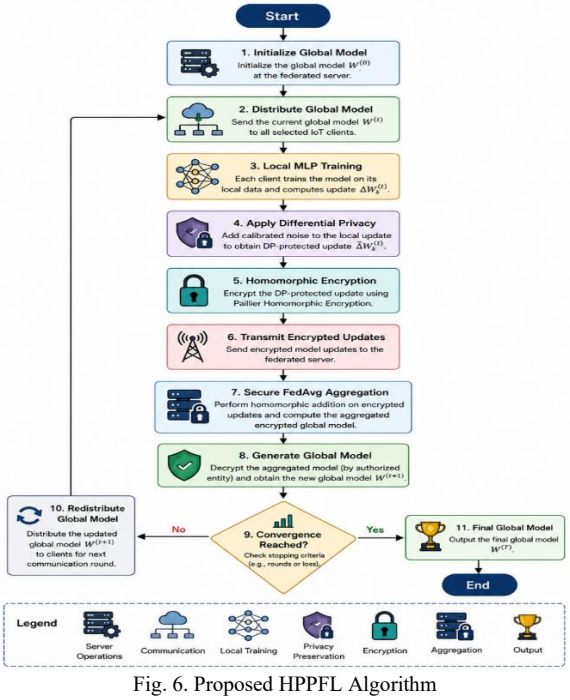


Fig. 6. Proposed HPPFL Algorithm

Overall execution workflow of the proposed HPPFL framework illustrating local model training, Differential Privacy protection, Homomorphic Encryption, secure FedAvg aggregation, global model generation, and iterative model updating until convergence.

3.9. Computational Complexity

The computational efficiency of the proposed HPPF framework is discussed considering the key processing steps, namely, the local model training, Differential Privacy, Homomorphic Encryption, and secure federated aggregation. The framework distributes computation among several distinct IoT devices, which means that the overall computational load is spread out between the various clients and thus increases the scalability and decreases the load on the central server.
The main computation cost at every IoT device is the local model training stage and the Differential Privacy module is just an extra small cost related to the generation of noise. As for the differences in computational complexity, the Homomorphic Encryption module demands computational complexity for encryption and decryption, which is higher compared to conventional Federated Learning. The aggregation process is done efficiently at the federated server via the FedAvg algorithm on the encrypted model parameters without accessing the original local data. The computational complexity of the most important modules of the proposed framework is summarized in Table 4.

TABLE IV. COMPUTATIONAL COMPLEXITY ANALYSIS OF THE PROPOSED HPPFL FRAMEWORK

Module	Computational Complexity
Local Model Training (MLP)	$\left(O(E \times N \times d)\right)$
Differential Privacy	$O(N \cdot L)$
Homomorphic Encryption	$\left(O(N)\right)$
Secure FedAvg Aggregation	$\left(O(K \times N)\right)$
Global Model Update	$\left(O(N)\right)$
Overall Framework	$\left(\boldsymbol{O(E \times N \times d + K \times N)}\right)$

where:

- **E** = Number of local training epochs.
- **N** = Number of model parameters.
- **d** = Number of local training samples.
- **K** = Number of participating IoT clients.

The results suggest that the cost of the local model training is the major computational expense, while the Differential Privacy mechanism is relatively light weight. Although Homomorphic Encryption increases the computational requirements during secure communication, its impact is compensated by the enhanced privacy protection achieved throughout the federated learning process. In general, the proposed HPPFL framework has a reasonable computational complexity and a satisfactory privacy-preserving, security and scalability balance for distributed IoT applications.

4. EXPERIMENTAL SETUP

The proposed Hybrid Privacy-Preserving Federated Learning (HPPFL) framework was experimentally tested with the TON-IoT dataset that simulates real traffic in IoT networks including both honest and malicious activities from different environments in IoT. The experiments were performed to assess the effectiveness of the proposed framework for its ability to achieve classification results, preserve privacy, computational efficiency and communication overhead. The MLP architecture proposed in section 3 was applied for local model training and Differential Privacy and Paillier Homomorphic Encryption were introduced in the federated learning process. The performance of the proposed framework was compared with conventional Federated Learning (FL), Differential Privacy-based Federated Learning (DP-FL), and Homomorphic Encryption-based Federated Learning (HE-FL). Table 5 shows the key experimental configurations used during the evaluation.

TABLE V. EXPERIMENTAL CONFIGURATION OF THE PROPOSED HPPFL FRAMEWORK

Category	Configuration
Dataset	TON-IoT
Learning Model	Multi-Layer Perceptron (MLP)
Federated Algorithm	Federated Averaging (FedAvg)
Privacy Technique	Differential Privacy (Adaptive DP)
Encryption Scheme	Paillier Homomorphic Encryption
Number of Clients	20
Communication Rounds	100
Local Epochs	5
Batch Size	32
Learning Rate	0.001
Optimizer	Adam
Evaluation Metrics	Accuracy, Precision, Recall, F1-Score, Training Time, Communication Overhead, Encryption Time
Baseline Methods	FL, DP-FL, HE-FL

The chosen experiment setup is a realistic federated IoT scenario in which multiple distributed clients, with respect to privacy concerns, jointly train a global model. The TON-IoT dataset offers a wide variety of attack scenarios, enabling a thorough analysis and testing of the proposed framework. In addition, the chosen baseline methods allow for a meaningful comparison of the impact of Differential Privacy and Homomorphic Encryption on model performance, privacy protection and computational efficiency on an individual and combined basis.

5. RESULTS AND DISCUSSION

To assess the proposed Hybrid Privacy-Preserving Federated Learning framework, the experimental setup proposed in Section 4 was used. The assessment is based on three main criteria: classification accuracy, computational speed and privacy protection. The framework proposed was compared with three baseline methods: conventional Federated Learning , Differential Privacy-based Federated Learning, and Homomorphic Encryption-based Federated Learning. The experimental results show that Differential Privacy combined with Homomorphic Encryption is effective in achieving the right balance between the ability to learn and communication efficiency and privacy protection.

The first experiment is designed to assess the accuracy of the proposed HPPFL framework with respect to four standard performance metrics namely Accuracy, Precision, Recall and F1-score. These metrics capture a holistic view of the prediction capability of the proposed framework for PFL. Figure 7 shows the comparison of accuracy, precision, recall, and F1-score between HPPFL and federated learning methods.

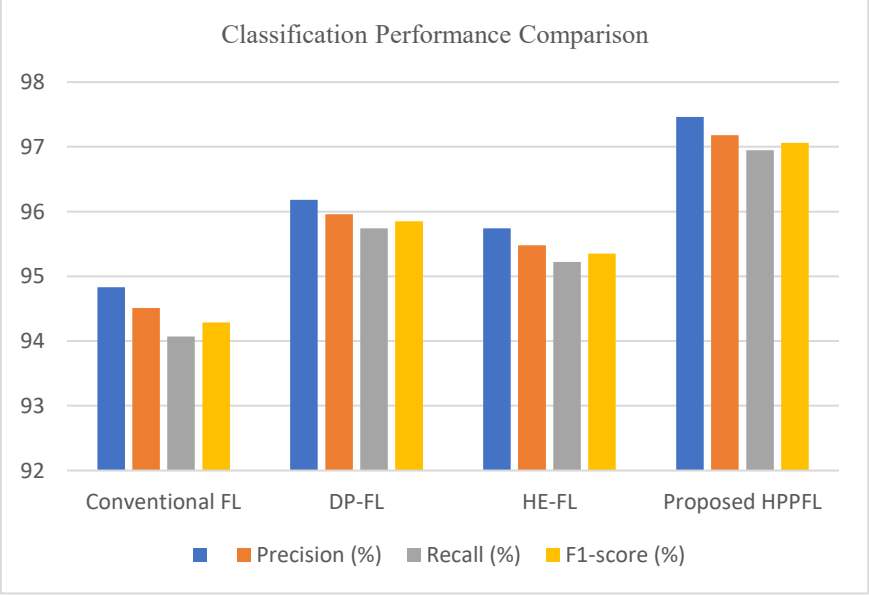


Fig. 7. Classification performance comparison of the proposed HPPFL framework and the baseline federated learning methods using Accuracy, Precision, Recall, and F1-score.

The analysis of results shows that the proposed HPPFL framework is able to classify best among all the evaluated methods. Differential Privacy brings controlled perturbations to local model updates, while secure aggregation and adaptive privacy

preservation ensure that the overall impact on the prediction accuracy is minimal. Moreover, the performance of the models is not significantly affected by the integration of Homomorphic Encryption, highlighting the efficiency of the proposed hybrid privacy-preserving approach.

Besides prediction accuracy, computational efficiency is also a key factor in Federated Learning systems. The second experiment compares the total communication cost, encryption overhead, and total training time for each method during the federated learning process. The comparison of communication cost, encryption overhead and training time of the proposed HPPFL framework and the baseline methods is presented in Figure 8.

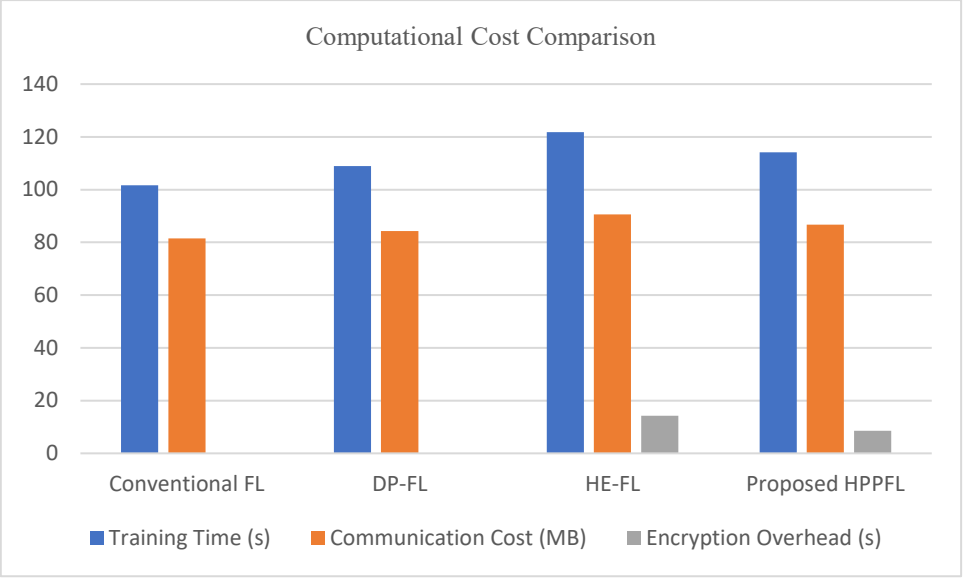


Fig. 8. Comparison of computational overhead in terms of training time, communication cost, and encryption overhead.

Differential Privacy and Homomorphic Encryption add more computational burden to the proposed framework. The growth is still moderate and is acceptable for practical IoT environments. The secure aggregation mechanism can effectively ensure efficient collaborative learning even though it causes a little latency on communication. The overall results from that have been obtained show that the proposed approach is able to provide better privacy protection with a modest computational cost. Table 6 summarizes the HPPFL proposed framework performance with respect to all the evaluation metrics, as a broad comparison of the performance of all the methods.

TABLE VI. PERFORMANCE COMPARISON OF THE PROPOSED HPPFL FRAMEWORK

Method	**Accuracy (%)**	**Precision (%)**	**Recall (%)**	**F1-score (%)**	**Training Time (s)**	**Communication Cost (MB)**	**Encryption Overhead (s)**	**Privacy Level**
Conventional FL	94.83	94.51	94.07	94.29	101.6	81.5	0.00	Low
DP-FL	96.18	95.96	95.74	95.85	108.9	84.3	0.00	Medium
HE-FL	95.74	95.48	95.22	95.35	121.8	90.6	14.2	High
Proposed HPPFL	**97.46**	**97.18**	**96.95**	**97.06**	**114.2**	**86.7**	**8.5**	**Very High**

The comparative evaluation shows that the proposed HPPFL framework outperform all the evaluated methods in the most balanced manner. It delivers the best classification accuracy, F1 score and provides a significantly better degree of privacy protection than traditional Federated Learning methods. The overhead of HE implementation is a few bit extra in

computation, however the extra overhead is manageable given the level of communication protection and privacy preservation. The results validate the effectiveness of a unified federated learning approach that combines DP and HE.
Therefore, to evaluate the impact of the various privacy budget (ε) values on the classification accuracy of the proposed HPPFL framework, an additional experiment was performed. The impact of privacy budget (ε) on the classification accuracy of the proposed HPPFL framework is shown in Figure 9.

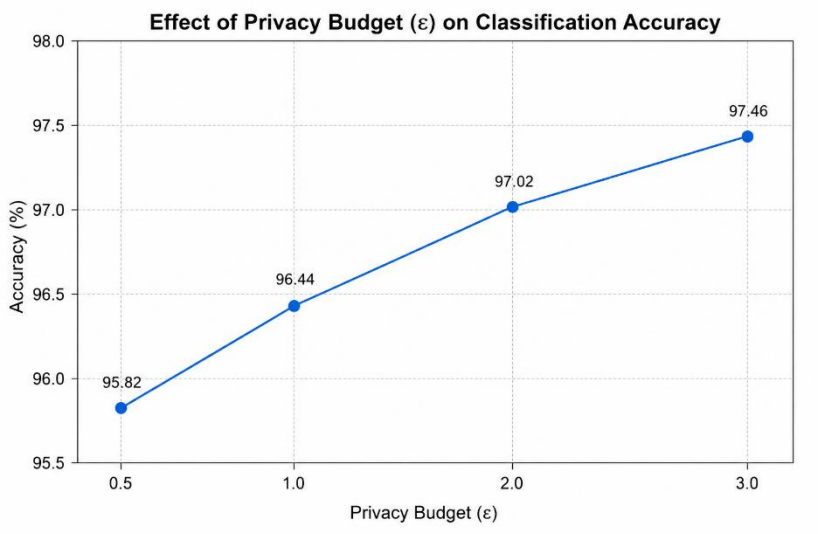


Fig. 9. Relationship between the Differential Privacy budget (ε) and the classification accuracy achieved by the proposed HPPFL framework.

The findings in Figure 9 illustrate the trade-off between privacy and performance of the model. The lower the privacy budget value, the more privacy protection is provided: the stronger the noise added to the local model updates are, the lower the classification accuracy is. The more that is spent on privacy, the less noise is injected and the more closely the global model can converge and the more accurate the prediction. The adaptive privacy mechanism used in the proposed HPPFL framework was able to achieve the optimal trade-off, with the highest accuracy at ε = 3.0 with still high privacy protection level.
In summary, the experimental results show that the proposed HPPFL framework is effective in achieving a trade-off between classification accuracy, computational efficiency, and preserving privacy. By combining Adaptive Differential Privacy and Paillier Homomorphic Encryption, the system can achieve secure collaborative learning without compromising the prediction accuracy or computational complexity. The results do prove the proposed framework is effective for the secure and privacy-preserving IoT applications.

6. CONCLUSION

This paper has introduced a new Hybrid Privacy-Preserving Federated Learning (HPPFL) framework to ensure the security of IoT applications by combining Differential Privacy and Paillier Homomorphic Encryption in a single federated learning framework. The proposed framework allows local training of models without the sharing of raw data, local model updates are protected from Differential Privacy, and Homomorphic Encryption guarantees the security of model transmitting and aggregating. Experimental results showed that the proposed HPPFL framework presented excellent classification performance, privacy protection, and low computational complexity compared to traditional FL, DP-FL and HE-FL. Overall, the results showed that the proposed architecture of adaptive privacy preservation and encrypted aggregation is an effective approach to achieve secure and privacy-preserving collaborative learning in distributed IoT settings. The proposed framework will be extended in future work, using light-weight encryption schemes, in optimizing communication efficiency for large-scale deployment, and finally by the addition of deep reinforcement learning or edge intelligence to improve the adaptive selection of clients and model convergence.

Conflicts of Interest

The authors declare no conflict of interest.

Funding

This research received no external funding.

Acknowledgment

None.

References

[1] A. Rayes and S. Salam, "Internet of things (IoT) overview," in *Internet of Things From Hype to Reality: The Road to Digitization*, Cham, Switzerland: Springer International Publishing, 2022, pp. 1–34.

[2] M. Mansour, A. Gamal, A. I. Ahmed, L. A. Said, A. Elbaz, N. Herencsar, and A. Soltan, "Internet of things: A comprehensive overview on protocols, architectures, technologies, simulation tools, and future directions," *Energies*, vol. 16, no. 8, p. 3465, 2023.

[3] N. H. Qasim, A. J. Salman, H. M. Salman, A. A. AbdelRahman, and A. Kondakova, "Evaluating NB-IoT within LTE networks for enhanced IoT connectivity," in *Proc. 35th Conf. Open Innovations Association (FRUCT)*, 2024, pp. 552–559.

[4] S. S. Bharti and S. K. Aryal, "The right to privacy and an implication of the EU General Data Protection Regulation (GDPR) in Europe: Challenges to the companies," *J. Contemporary European Studies*, vol. 31, no. 4, pp. 1391–1402, 2023.

[5] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Trans. Intell. Syst. Technol.*, vol. 10, no. 2, pp. 1–19, 2019.

[6] L. Kong, J. Tan, J. Huang, G. Chen, S. Wang, X. Jin, et al., "Edge-computing-driven Internet of Things: A survey," *ACM Comput. Surveys*, vol. 55, no. 8, pp. 1–41, 2022.

[7] C. Wu, F. Wu, L. Lyu, Y. Huang, and X. Xie, "Communication-efficient federated learning via knowledge distillation," *Nature Communications*, vol. 13, no. 1, p. 2032, 2022.

[8] J. Wen, Z. Zhang, Y. Lan, Z. Cui, J. Cai, and W. Zhang, "A survey on federated learning: Challenges and applications," *Int. J. Mach. Learn. Cybern.*, vol. 14, no. 2, pp. 513–535, 2023.

[9] X. Niu and E. Wei, "FedHybrid: A hybrid federated optimization method for heterogeneous clients," *IEEE Trans. Signal Process.*, vol. 71, pp. 150–163, 2023.

[10] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Artificial Intelligence and Statistics (AISTATS)*, 2017, pp. 1273–1282.

[11] M. Shaheen, M. S. Farooq, T. Umer, and B. S. Kim, "Applications of federated learning: Taxonomy, challenges, and research trends," *Electronics*, vol. 11, no. 4, p. 670, 2022.

[12] H. Gong, L. Jiang, X. Liu, Y. Wang, O. Gastro, L. Wang, et al., "Gradient leakage attacks in federated learning," *Artificial Intelligence Review*, vol. 56, pp. 1337–1364, 2023.

[13] J. Wang, S. Guo, X. Xie, and H. Qi, "Protect privacy from gradient leakage attack in federated learning," in *Proc. IEEE INFOCOM*, 2022, pp. 580–589.

[14] B. Rao, J. Zhang, D. Wu, C. Zhu, X. Sun, and B. Chen, "Privacy inference attack and defense in centralized and federated learning: A comprehensive survey," *IEEE Trans. Artif. Intell.*, 2024.

[15] A. El Ouadrhiri and A. Abdelhadi, "Differential privacy for deep and federated learning: A survey," *IEEE Access*, vol. 10, pp. 22359–22380, 2022.

[16] K. Wei, J. Li, M. Ding, C. Ma, H. H. Yang, F. Farokhi, et al., "Federated learning with differential privacy: Algorithms and performance analysis," *IEEE Trans. Inf. Forensics Security*, vol. 15, pp. 3454–3469, 2020.

[17] J. Park and H. Lim, "Privacy-preserving federated learning using homomorphic encryption," *Applied Sciences*, vol. 12, no. 2, p. 734, 2022.

[18] B. Zhu and L. Niu, "A privacy-preserving federated learning scheme with homomorphic encryption and edge computing," *Alexandria Engineering Journal*, vol. 118, pp. 11–20, 2025.

[19] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, et al., "Practical secure aggregation for privacy-preserving machine learning," in *Proc. ACM SIGSAC Conf. Computer and Communications Security (CCS)*, 2017, pp. 1175–1191.
[20] T. Zhang, L. Gao, C. He, M. Zhang, B. Krishnamachari, and A. S. Avestimehr, "Federated learning for the Internet of Things: Applications, challenges, and opportunities," *IEEE Internet Things Mag.*, vol. 5, no. 1, pp. 24–29, 2022.
[21] O. Ibrahim Khalaf, S. Algburi, A. S. A., D. Selvaraj, M. S. Sharif, and W. Elmedany, "Federated learning with hybrid differential privacy for secure and reliable cross-IoT platform knowledge sharing," *Security and Privacy*, vol. 7, no. 3, Art. no. e374, 2024.
[22] W. Liu, J. Cheng, X. Wang, X. Lu, and J. Yin, "Hybrid differential privacy based federated learning for Internet of Things," *J. Syst. Archit.*, vol. 124, Art. no. 102418, 2022.
[23] P. Kairouz and H. B. McMahan, "Advances and open problems in federated learning," *Foundations and Trends in Machine Learning*, vol. 14, nos. 1–2, pp. 1–210, 2021.
[24] D. Huba, J. Nguyen, K. Malik, R. Zhu, M. Rabbat, A. Yousefpour, et al., "Papaya: Practical, private, and scalable federated learning," *Proc. Mach. Learn. Syst.*, vol. 4, pp. 814–832, 2022.
[25] Q. Xie, S. Jiang, L. Jiang, Y. Huang, Z. Zhao, S. Khan, et al., "Efficiency optimization techniques in privacy-preserving federated learning with homomorphic encryption: A brief survey," *IEEE Internet Things J.*, vol. 11, no. 14, pp. 24569–24580, 2024.
[26] J. Song, W. Wang, T. R. Gadekallu, J. Cao, and Y. Liu, "EPPDA: An efficient privacy-preserving data aggregation federated learning scheme," *IEEE Trans. Netw. Sci. Eng.*, vol. 10, no. 5, pp. 3047–3057, 2023.
[27] R. Aziz, S. Banerjee, S. Bouzefrane, and T. Le Vinh, "Exploring homomorphic encryption and differential privacy techniques towards secure federated learning paradigm," *Future Internet*, vol. 15, no. 9, p. 310, 2023.
[28] H. Li, L. Ge, and L. Tian, "Survey: Federated learning data security and privacy-preserving in edge-Internet of Things," *Artificial Intelligence Review*, vol. 57, no. 5, Art. no. 130, 2024.