

PROFL: A Privacy-Preserving Federated Learning Method with Stringent Defense Against Poisoning Attacks

Yisheng Zhong, Li-Ping Wang*

Key Laboratory of Cyberspace Security Defense, Institute of Information Engineering, Chinese Academy of Sciences
School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China
{zhongyisheng,wangliping}@iie.ac.cn

Abstract—Federated Learning (FL) faces two major issues: privacy leakage and poisoning attacks, which may seriously undermine the reliability and security of the system. Overcoming them simultaneously poses a great challenge. This is because privacy protection policies prohibit access to users' local gradients to avoid privacy leakage, while Byzantine-robust methods necessitate access to these gradients to defend against poisoning attacks. To address these problems, we propose a novel privacy-preserving Byzantine-robust FL framework PROFL. PROFL is based on the two-trapdoor additional homomorphic encryption algorithm and blinding techniques to ensure the data privacy of the entire FL process. During the defense process, PROFL first utilize secure Multi-Krum algorithm to remove malicious gradients at the user level. Then, according to the Pauta criterion, we innovatively propose a statistic-based privacy-preserving defense algorithm to eliminate outlier interference at the feature level and resist impersonation poisoning attacks with stronger concealment. Detailed theoretical analysis proves the security and efficiency of the proposed method. We conducted extensive experiments on two benchmark datasets, and PROFL improved accuracy by 39% to 75% across different attack settings compared to similar privacy-preserving robust methods, demonstrating its significant advantage in robustness.

Index Terms—Federated Learning, Poisoning Attack, Privacy-Preserving, Defense Strategy.

I. INTRODUCTION

Federated Learning (FL) allows data owners to enhance the efficiency and accuracy of global models without exposing their private data. FL has broad potential applications in fields such as healthcare [1] and finance [2]. However, FL also faces problems including privacy leakage and poisoning attacks, which pose a threat to its security and reliability.

Privacy leakage that results in data security issues is the first major threat faced by FL. As the gradients contain some private information, uploading them directly to the central server without undergoing special treatment can result in leakage issues [3]. Poisoning attack leading to a decrease in model reliability is the another major threat to FL. Byzantine adversaries employ poisoning attacks to manipulate malicious participants and intentionally inject harmful information into the uploaded gradient. This manipulation leads to gradient deviation, ultimately undermining model performance [4] and altering model decisions [5]. Therefore, developing a FL

mechanism that simultaneously preserves privacy and resists poisoning attacks is crucial to making FL a trustworthy and reliable technique.

However, there are still two important obstacles to achieving such a FL system. One is that privacy-preserving FL aims to guarantee the indistinguishability of the data, while defensive approaches necessitate access to local gradients and use similarity measures to distinguish malicious gradients from benign gradients, which seems paradoxical. The other is that there are now cunning Byzantine adversaries who make poisoning gradient attacks more covert. For example, some camouflage methods [6] are capable of analyzing the range of parameter changes that the defender cannot detect and then evade defense mechanisms. Recent approaches achieve a balance between privacy and robustness. For instance, PEFL [7] utilizes the Pearson correlation coefficient to identify outliers, employing a dual-server model for secure computation where trusted servers hold keys to collaborate with additive homomorphic encryption algorithms. The ShieldFL [8] adopts Two-trapdoor homomorphic encryption technology to relax the trust requirement for servers, using cosine similarity to identify malicious gradients. However, their robustness and privacy are not ideal.

To address the above issues, we proposes a composite privacy-preserving Byzantine-robust algorithm called PROFL. It performs two-trapdoor additional homomorphic algorithm (AHE) [9] and blinding techniques on two non-colluding servers to prevent privacy leakage during the defense process. In addition, PROFL first uses the multi-krum algorithm [4] based on the Euclidean distance as the similarity measure to detect and eliminate poisoning gradients at the user level. Then, PROFL uses statistical methods based on the Pauta criterion to eliminate outliers at the feature level, thereby resisting concealed impersonation poisoning attack. Such a composite defense can defend against malicious gradient attacks from macro and micro perspectives simultaneously. The main contributions are summarized as follows.

- We propose a composite robust mechanism to counter both general and concealed poisoning attacks. Across multiple attack scenarios, compared to similar ap-

proaches, our method demonstrates an improvement of 39% to 75% in performance. This showcases the broader and more rigorous advantages of our composite defense.

- Compared to similar methods, both the security and efficiency of our privacy-preserving strategy are higher. PROFL effectively mitigates a broader range of collusion attacks and, unlike multi-key methods, significantly reduces the computational overhead associated with key pair generation and joint decryption.
- Theoretical analysis demonstrates that PROFL performs well in managing communication and time overheads.

The rest of this paper is organized as follows. In §II, we introduce the problem statement. In §III, we present the PROFL in detail. Subsequently, we perform the theoretical analysis in §IV and evaluate the performance in §V. Finally, the paper is concluded in §VI.

II. PROBLEM STATEMENT

A. System Model

Our system consists of a key center, a pair of servers, and multiple users. Their relationship is as shown in the Fig 1.

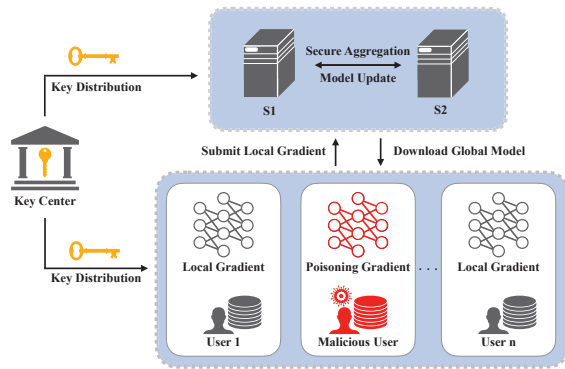


Fig. 1. System Model.

Key Center(KC): KC is a trusted agency responsible for the generation and distribution of keys. First, it generates two public-secret key pairs and splits one of the secret keys into two shares. After generation, it distributes public and secret keys and key shares to the designated participants.

Server(S_1 and S_2): The servers S_1 and S_2 cooperate and restrict each other to complete complex operations over the encrypted data and resist poisoning attacks. S_1 is the primary server, which is responsible for communicating with users and executing robust algorithms. S_2 is responsible for assisting S_1 to complete the secure computation.

User(U): Users use the latest global model to locally train the model on private data. Then, the gradient is encrypted and uploaded to S_1 . After S_1 and S_2 cooperate to defend and update the global model, users download the latest global model to replace their local models.

B. Threat Model

We assume KC is completely trustworthy, so the possible threats come from servers and users. Firstly, the two servers

are semi-honest. They both follow the pre-defined protocol, but they attempt to decrypt the users' gradients, and they cannot collude with each other [10]. Similarly, users are semi-honest, but they try to collude with S_1 to decrypt other users' gradients.

Secondly, the robustness of the model is threatened by malicious users. We assume that malicious users can have knowledge of benign users' gradient distributions and other information in advance. Malicious users submit poisoning gradients to deviate the global model from the training target. Moreover, considering that defense techniques often utilize the principle of similarity to resist poisoning attacks, malicious adversaries may employ countermeasures. These adversaries cunningly design malicious gradients to be as close as possible to benign gradients in terms of distance or similarity in order to evade defenses. They also aim to make significant changes in a few dimensions of these gradients to achieve their poisoning objectives. Specifically, a malicious user will choose p parameters from its malicious gradient \mathbf{g}_m and modify them to arbitrarily large values of t such that:

$$\text{Max}(\text{Dis}(\mathbf{g}_i, \mathbf{g}_j) | (i, j \in (1, n))) > \text{Dis}(\mathbf{g}_m = (\dots, t_1, \dots, t_p, \dots), \mathbf{g}_b)$$

where \mathbf{g}_b is a benign gradients and Dis is a distance function. To ensure the practicability and realizability of the defense strategy, we assume that the upper limit of malicious users is $\lfloor \frac{n}{2} \rfloor$ of all users.

III. FRAMEWORK

PROFL employs the Two-trapdoor AHE algorithm to preserve privacy, which is a type of distributed public-key cryptosystem. It is a variant of the Paillier Cryptosystem [11] and possesses additive homomorphic properties. It mainly consists of the components KeyGen, Enc, FullDec, KeySplit, PartDec1 and PartDec2. Besides, the two-trapdoor AHE algorithm satisfies the additive homomorphism, and we define their operator symbols as \oplus and \odot respectively. Therefore, given two ciphertexts $[[x_1]]$ and $[[x_2]]$ and an integer $A \in \mathbb{Z}_N$, we can perform such a secure calculation in the encrypted state:

$$[[x_1]] \oplus [[x_2]] = [[x_1]] \cdot [[x_2]] = [[x_1 + x_2]] \quad (1)$$

$$[[x_1]] \odot A = [[x_1]]^A = [[x_1 \cdot A]] \quad (2)$$

A. Initialization

Key Generation and Distribution: Firstly, KC generates two sets of public-secret keys using $(pk_m, sk_m) \leftarrow \text{KeyGen}(1^{k_1})$ and $(pk_g, sk_g) \leftarrow \text{KeyGen}(1^{k_2})$, then delivers the public key pk_m to S_1 and broadcasts the secret key sk_m to all users, which are respectively used for encrypting and decrypting the new global model. Next, KC broadcasts the public key pk_g to all users for encrypting the local gradients submitted to the server. KC delivers the public key pk_g to S_1 to encrypt the intermediate variables for secure computation. KC then splits the secret key sk_g into two secret shares sk_{g_1} and sk_{g_2} using $(sk_{g_1}, sk_{g_2}) \leftarrow \text{KeySplit}(sk_g)$, and distributes

them to servers S_1 and S_2 , respectively, for jointly decrypting the gradients submitted by the users.

Model Initialization and Distribution: Firstly, the server S_1 randomly initializes a global model parameter \mathbf{w}_0 , and then uses the public key pk_m to encrypt it as $[[\mathbf{w}_0]]$ using $[[\mathbf{w}_0]] \leftarrow \text{Enc}(pk_m, \mathbf{w}_0)$. Then S_1 broadcasts the encrypted model $[[\mathbf{w}_0]]$ to each user. Each user U_i decrypt the global model to obtain \mathbf{w}_0 with the secret key sk_m using $\mathbf{w}_0 \leftarrow \text{FullDec}(sk_m, [[\mathbf{w}_0]])$.

B. Local Training

During the gradient generation phase, Benign users utilize local data to conduct mini-batch local training on the latest model, and then accumulate the obtained gradients. In the same phase, malicious users will not submit gradients obtained by training on real data. Instead, they upload their own designed malicious gradients. As two-trapdoor AHE algorithm are operated on integers, we use the precision parameter deg to convert all parameters of gradients from floating-point to integer format such that: $g' = \lfloor g \cdot deg \rfloor$. Then, each user encrypts the gradient and submits it to S_1 for aggregation.

C. Privacy-Preserving Defense Strategy

To avoid the risk of a single untrusted server leaking secret keys or decrypting users' gradients, PROFL utilizes a combination of two-trapdoor AHE mechanism and blinding techniques to execute defensive strategies on two servers without revealing plaintext gradients. Furthermore, in order to prevent malicious users from colluding with either of the two servers to decrypt gradients submitted by other users, we set up two sets of public-secret key pairs. This enables users to encrypt their own submitted gradients, while preventing them from decrypting gradients submitted by other users.

To defend against poisoning attacks, PROFL proposes a secure defense algorithm based on the Multi-Krum algorithm. During the defense process, the secure distance function SecDis is used to calculate the Euclidean distance between two encrypted gradients. Then, the secure selection function SecSel selects $\lfloor \frac{n}{2} \rfloor$ gradients that are closest to other local gradients in geometric space as representatives. Based on the proportion of malicious users, these representative gradients are closer to the center of the group, and it is believed that they can represent the characteristics of the vast majority of benign user gradients, thereby enabling the global gradient to converge towards the optimal aggregated gradient.

To address the issue of cunning malicious users whose gradients cannot be distinguished using defense methods based on similarity principles, we propose the Secure Represent function SecRep to mitigate this vulnerability. SecRep applies the Pauta criterion to remove outliers in each dimension and then selects the median as the representative for that dimension. Through these operations, we obtain a representative global gradient. The procedure is shown in Algorithm 1, and the three security protocols are as follows:

SecDis: SecDis employs additive homomorphic algorithm and blinding technique to calculate the Euclidean distance

Algorithm 1: Privacy-preserving Defense Strategy

Input: The encrypted local gradients $\{[[\mathbf{g}_1]], [[\mathbf{g}_2]], \dots, [[\mathbf{g}_n]]\}$.
Output: Representative gradient $[[\mathbf{g}_{agg}]]$.

- 1 /*Calculate the sum of the distances of each gradient to the other gradients*/
- 2 **for** $i \in [1, n]$ **do**
- 3 $sum_i = 0$;
- 4 **end**
- 5 **for** $i \in [1, n]$ **do**
- 6 **for** $j \in [i, n]$ **do**
- 7 $dis(i, j) = \text{SecDis}([[g_i]], [[g_j]])$;
- 8 /*Accumulate $dis(i, j)$ using homomorphic addition*/
- 9 $[[sum_i]] = [[sum_i]] \oplus dis(i, j)$,
- 10 $[[sum_j]] = [[sum_j]] \oplus dis(i, j)$;
- 11 **end**
- 12 /*Screening Representatives*/
- 13 Using SecSel , selects $\lfloor \frac{n}{2} \rfloor$ gradients with the largest corresponding value in (sum_i) ($i \in [1, n]$) among n gradients, as the representative of benign gradients;
- 14 /*Remove Outliers*/
- 15 Arrange all gradients into sets $\{d_i\}_{i \in [1, m]}$ by dimension, where $d_i = \{[[g_1[i]]], [[g_2[i]]], \dots, [[g_n[i]]]\}$ (here we use $g[i]$ to denote the i -th element of \mathbf{g});
- 16 **for** $i \in [1 : m]$ **do**
- 17 Use SecRep algorithm to remove outliers in set d_i and calculate the median med_i of set d_i as a representative of the i -th dimension;
- 18 **end**
- 19 **return:** aggregated gradient $[[\mathbf{g}_{agg}]] = [[(med_1, med_2, \dots, med_m)]]$.

between two encrypted gradients. Given two encrypted gradients $[[\mathbf{g}_x]]$ and $[[\mathbf{g}_y]]$, the servers S_1 and S_2 jointly calculate the Euclidean distance $dis(i, j) \leftarrow \text{SecDis}([[g_x]], [[g_y]])$. The specified procedure is demonstrated in Fig.2.

SecSel: To identify gradients that can represent benign users, S_1 and S_2 jointly call decryption algorithm PartDec1 and PartDec2 to obtain $(sum_{i \in [1, n]})$ and find the $\lfloor \frac{n}{2} \rfloor$ gradients with the largest corresponding value in (sum_i) among the n gradients. To minimize the time spent in the selection process, we use the Randomized-Selection (RS) algorithm, which can solve the this top-k problem in $\mathcal{O}(n)$ time complexity.

SecRep: To supplement the defense vulnerability, PROFL removes outliers in each dimension using the Pauta criterion, also known as the 3σ criterion. The Pauta criterion is a very common statistical method for dealing with outliers. It can accurately rejects outliers and preserves benign value. Submitted gradients has a symmetrical distribution with high in the middle and low on both sides, which can be approximated as a normal distribution, thus it meets the preconditions of the Pauta criterion. Finally, for each dimension, servers re-statistics and select the median as the representative value of the current dimension, as the final result of the aggregation. To protect the privacy of gradients in the above process, we need to use blinding techniques to perform statistics on the decrypted gradient value without exposing the plaintext. The specified procedure is shown in Fig.3.

Implementation of SecDis

Input: S_1 holds two gradients $[[\mathbf{g}_x]], [[\mathbf{g}_y]]$, the public key pk_g and the secret key sk_1 ; S_2 holds the secret key sk_2 .

Output: Encrypted Euclidean distance $[[d]]$ between $[[\mathbf{g}_x]]$ and $[[\mathbf{g}_y]]$.

Procedure:

- @ S_1 : To protect privacy of gradients, for $i \in \{x, y\}$, S_1 randomly selects two nonzero noise vectors $\mathbf{r}_i \leftarrow Z_N^*$ to blind the gradients using Equation (2) such that $[[\mathbf{g}'_i]] = [[\mathbf{g}_i]] \cdot [[\mathbf{r}_i]] = [[\mathbf{g}_i + \mathbf{r}_i]]$. Then, S_1 partially decrypts the gradient by calling $[[\mathbf{g}'_i]]_{sk_1} \leftarrow \text{PartDec1}(sk_1, [[\mathbf{g}'_i]])$, and send $[[\mathbf{g}'_i]]_{sk_1}$ and $[[\mathbf{g}'_i]]$ to S_2 .
- @ S_2 : To fully decrypt gradients, S_2 calls algorithm $\mathbf{g}'_i \leftarrow \text{PartDec2}(sk_2, [[\mathbf{g}'_i]], [[\mathbf{g}'_i]]_{sk_1})$ ($i \in \{x, y\}$). Then S_2 calculates the Euclidean distance \mathbf{dis}' between \mathbf{g}'_x and \mathbf{g}'_y such that $\mathbf{dis}' = (\mathbf{g}'_x - \mathbf{g}'_y)^2 = ((\mathbf{g}_x + \mathbf{r}_x) - (\mathbf{g}_y + \mathbf{r}_y))^2$. Then S_2 sends \mathbf{dis}' to S_1 .
- @ S_1 : To remove the random noises and get $[[\mathbf{dis}]]$ using Equation (1) such that:

$$[[\mathbf{dis}]] = [[\mathbf{dis}' + \mathbf{t}]] = [[\mathbf{dis}']] \cdot [[\mathbf{t}]] \quad (3)$$

which needs intermediate results $[[\mathbf{dis}']]$ and $[[\mathbf{t}]]$. $[[\mathbf{t}]]$ can be obtained using Equation (1) and (2) such that:

$$\begin{aligned} [[\mathbf{t}]] &= [[\mathbf{dis} - \mathbf{dis}']] = [[(\mathbf{g}_x - \mathbf{g}_y)^2 - (\mathbf{g}'_x - \mathbf{g}'_y)^2]] \\ &= [[\mathbf{g}_x]]^{-2r_x} \cdot [[-\mathbf{r}_x^2]] \cdot [[\mathbf{g}_x]]^{2r_x} \cdot [[\mathbf{g}_y]]^{2r_x} \cdot [[2r_x r_y]] \\ &\quad \cdot [[\mathbf{g}_y]]^{-2r_y} \cdot [[-\mathbf{r}_y^2]] \end{aligned} \quad (4)$$

where S_1 owns $[[\mathbf{g}_x]]$, $[[\mathbf{g}_y]]$, \mathbf{r}_x and \mathbf{r}_y . Then, removes the random noises \mathbf{r}_x and \mathbf{r}_y and gets $[[\mathbf{dis}]]$ using Equation (3), (4) and $[[\mathbf{dis}']] \leftarrow \text{Enc}(pk_g, \mathbf{dis}')$. Finally, accumulate each dimension in \mathbf{dis} to obtain a numerical value $[[d]]$ for the distance between gradients using Equation (1) such that $\prod_{i=1}^m [[\mathbf{dis}(i)]]$.

Fig. 2. Detailed procedure of SecDis.

D. Integrate and Distribute

After executing the SecRep in each dimension, S_1 concatenates the representatives of each dimension to get the global gradient $\mathbf{g}_{agg} = (med_1, \dots, med_m)$. Then, convert the global gradient from integer type to floating-point type and use it to update the current global model according to $\mathbf{w}' := \mathbf{w} - lr \cdot \mathbf{g}$. Finally, the latest global model \mathbf{w}' is encrypted with the public key pk_m using $[[\mathbf{w}']] \leftarrow \text{Enc}(pk_m, \mathbf{w}')$, and distributed to each user.

IV. THEORETICAL ANALYSIS

A. Security Analysis

We provide a standard hybrid argument so as to prove that the joint view of servers S_1 and S_2 does not disclose any sensitive information about local gradients during protocol execution. In other words, in the honest-but-curious setting, PROFL achieves IND-CPA security [12]. The security theorem is shown in Theorem 1, and due to space constraints, the proof is omitted here.

Implementation of SecRep

Input: S_1 holds the encrypted value of all gradients in the t -th dimension $d_t = \{[[\mathbf{g}_1(t)]], [[\mathbf{g}_2(t)]], \dots, [[\mathbf{g}_{n'}(t)]]\}$ ($n' = \lceil \frac{n}{2} \rceil$) and the secret key sk_1 ; S_2 holds the secret key sk_2 .

Output: The representative of all gradients in the t -th dimension med_t .

Procedure:

- @ S_1 : To protect privacy, S_1 randomly selects a nonzero noise vector $\mathbf{r} \leftarrow Z_N^*$ to blind the gradients $\mathbf{g}_i(t)$ ($i \in [1, n']$) using Equation(2) such that $[[\mathbf{g}_i(t)']] = [[\mathbf{g}_i(t)]] \cdot [[\mathbf{r}]] = [[\mathbf{g}_i(t) + \mathbf{r}]]$. Then, S_1 partially decrypts the gradient by calling $[[\mathbf{g}_i(t)']]_{sk_1} \leftarrow \text{PartDec1}(sk_1, [[\mathbf{g}_i(t)']])$, and send $[[\mathbf{g}_i(t)']]_{sk_1}$ and $[[\mathbf{g}_i(t)']]$ to S_2 .
- @ S_2 : To fully decrypt the gradients $[[\mathbf{g}_i(t)']]$ ($i \in [1, n']$), S_2 calls $\mathbf{g}_i(t)' \leftarrow \text{PartDec2}(sk_2, [[\mathbf{g}_i(t)']], [[\mathbf{g}_i(t)']]_{sk_1})$. Then, S_2 calculates the mean μ and variance σ of $\mathbf{g}_i(t)'$ and eliminates the gradient value outside the range of $\mu \pm 3\sigma$ according to the Pauta criterion. Finally, S_2 re-statistics the gradient value to find the median med'_t and sends it to S_1 .
- @ S_1 : To obtain the median med_t , S_1 removes the blinding r of the median med'_t using $med_t = med'_t - r$.

Fig. 3. Detailed procedure of SecRep.

Theorem 1. Given a security parameter κ , two non-colluding servers S_1 and S_2 , users U , define a random variable $REAL_{\Pi}^{U, \kappa}$ to represent the joint view of servers S_1 and S_2 during the execution of the protocol Π . There exists a probabilistic polynomial-time (PPT) simulator SIM , whose view $SIM_{\Pi}^{U, \kappa}$ is computationally indistinguishable from $REAL_{\Pi}^{U, \kappa}$.

B. Complexity Analysis

To evaluate the efficiency of PROFL, we discuss the communication and computational complexity for each training iteration, as shown in Table I.

TABLE I
COMPUTATION AND COMMUNICATION COMPLEXITY IN PROFL

Phase	PROFL	
	Compu.	Comm.
SecDis*	$\mathcal{O}(n^2 m (T_{Enc} + T_{Add}) + nm (T_{Dec} + T_{Mul}))$	$\mathcal{O}(2nm X)$
SecSel	$\mathcal{O}(nT_{Dec})$	$\mathcal{O}(n)$
SecRep	$\mathcal{O}(nm(T_{Add} + T_{Dec}))$	$\mathcal{O}(2m X)$

Note: SecDis* indicates that SecDis is executed multiple times.

Let n and m be the number of users participating in the aggregation and the number of dimensions of gradients. $|X|$ represents the communication complexity of an encrypted number. T_{Add} , T_{Mul} , T_{Enc} , and T_{Dec} are the time complexities of homomorphic operations, where homomorphic decryption consists of partial decryption step one and two. The time complexities $T_{Enc} > T_{Dec} \gg T_{Mul} > T_{Add}$.

When calculating the centrality of n gradients, each gradient must calculate the Euclidean distance with the rest of the gradients, so SecDis needs to be called $\sum_{i=1}^{n-1} i$ times. However, in this process, many intermediate variables are repeatedly calculated, which can be avoided by storing them. The same approach can be taken for the communication.

Therefore, the computational and communication complexity can be greatly reduced. It's worth noting that we observe a positive correlation between the computational complexity of the SecDis phase and n^2 . However, considering that in practical applications, such as in the medical and banking fields, each user represents an institution, the quantity of users n is often small. Therefore, the computational overhead of SecDis won't become unacceptable due to the growth of n .

In the SecSel protocol, the complexity mainly depends on the number of gradients n . In the SepRep process, the complexity is limited to the number of benign gradients $\lceil \frac{n}{2} \rceil$ and the gradient dimension m .

V. PERFORMANCE EVALUATION

In this section, we conducted experiments on two datasets using multiple different settings to evaluate the performance of PROFL: (1) evaluating the performance of PROFL under various attack settings (attack method, attack ratio); (2) comparing the performance of PROFL with four other typical defense schemes, including Krum [4], Trimmed Mean [13], PEFL [7], and ShieldFL [8]. To highlight the advantages of PROFL, we used FedAvg [14] without any defense as a baseline.

A. Testbed and Methodology

All experiments were run on a Macbook pro computer equipped with a M2 pro processor. We implemented the two-trapdoor AHE protocol at the 80-bit security level [15] using python 3.8.11. The PROFL algorithm was tested using the pytorch library and lightweight threads were used to simulate the local learning processes of multiple users.

Data Setting: To evaluate the performance of PROFL, we conducted experiments on three benchmark datasets, MINIST, FashionMINIST and CIFAR-10. To simulate real situations, the local learning model of each user should be underfitting, which can reflect the value of FL. Therefore, we believe that the total number of training data should be set at about 10%. PROFL is designed for independently and identically distributed data, so we choose to randomly shuffle the dataset and distribute it equally to each user.

Models and Hyperparameters: In this task, we use a classical logistic regression as the model structure. We set the total number of users in the experiment to 20. For each experiment, we set the batch size to 256, momentum to 0.5, and perform 1000 iterations on datasets. The data point for each experiment is the average of 10 experimental results.

Poisoning Attack Settings: To evaluate the ability of PROFL to resist poisoning attacks, we designed two types of attacks: target attacks [16] and non-target attacks. In addition, after the gradient generation of attacks, we added the insidious gradient attack described in the threat model. We simulate the process of a Byzantine adversary evade the robust defense method and poison the model by carefully designing the value of certain dimensions of the malicious gradient. To evaluate the performance of PROFL, we conducted experiments under different proportions of malicious users. We used attack ratio $Att_{ratio} = |U^*|/|U|$ to describe the percentage of malicious

users, where $|U|$ is the number of all users and $|U^*|$ is the number of malicious users.

B. Accuracy Evaluation

Table II measures the performance of PROFL on target and non-target attacks. In targeted attacks, we marked the difference in the testing accuracy of the global model for the specific source class Acc_{source} between PROFL and baseline. In non-targeted attacks, we marked the difference in the testing accuracy Acc between PROFL and baseline. $AI = Acc - Acc^*$ is the improvement in accuracy of the global model, where Acc^* is the testing accuracy of the baseline model. $AI_{source} = Acc_{source} - Acc_{source}^*$ is the improvement in accuracy of the global model for a specific source class.

TABLE II
ACCURACY COMPARISON ($Att_{ratio} = 50\%$)

Datasets	Attack	Baseline		PROFL			
		Acc^*	Acc^*	Acc_s	Acc	AI_s	AI
Minist	Targeted	22.3%	*	97.4%	92.8%	75.1%	*
	Untargeted	*	46.5%	*	92.1%	*	45.6%
Fashion	Targeted	19.9%	*	73.9%	82.3%	54.0%	*
	Untargeted	*	43.5%	*	82.5%	*	39.0%

Acc_s and AI_s denotes Acc_{source} and AI_{source} , respectively.

Fig.4(a) show the accuracy trajectories with $Att_{ratio}=50\%$ under the targeted attack, where the source class is '0' and the target class is '6'. As the figures indicate, PROFL has significantly improved $Acc_{sources}$ compared to the baseline.

To evaluate the impact of the ratio of attacks on PROFL, Fig.4(b) depicts the change of Acc and AI with $Att_{ratio} \in [0\%, 30\%, 50\%]$. The image indicates that PROFL maintains stable accuracy and remarkable accuracy improvement at different Att_{ratio} . Even under the worst case scenario of $Att_{ratio} = 50\%$, PROFL can still achieve high accuracy rates of up to 92% and 82% in the datasets MINIST and FashionMINIST respectively, with accuracy improvements of 48% and 43%. This demonstrates the powerful robustness of PROFL.

As shown in Fig.4(c)-4(d), We compared PROFL with existing excellent solutions at different Att_{ratio} . It can be observed that, as the proportion of malicious users increases, the accuracy of all defense measures except PROFL decreases significantly. We believe that this is because with the increase of Att_{ratio} , not only valuable data during the training process is reduced, but also the ability of defense methods to identify high-proportion malicious gradients weakens, leading to a deterioration of the robust aggregation effect. However, PROFL maintains excellent robustness throughout the process. We attribute this to PROFL's ability to accurately exclude poisoning gradients, making the contribution of benign gradients to the global gradient is purer.

C. Communication and Computational Overhead

We evaluated the communication overhead of PROFL. Additionally, we introduced the total communication volume provided by ShieldFL [8], which employs similar encryption approach as ours. This allows for a clear comparison with PROFL under the same security level. The communication volume comparison is illustrated in the Table.III. We can observe

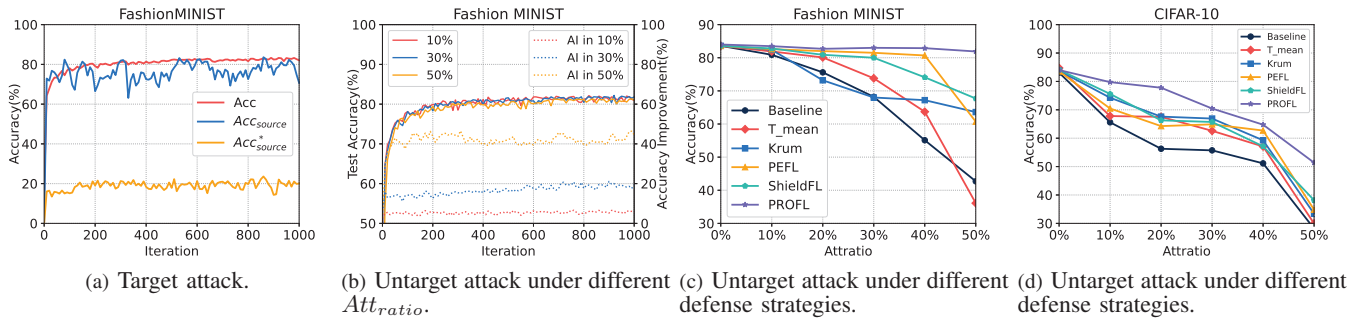


Fig. 4. EXPERIMENTAL RESULT.

TABLE III

TOTAL COMPUTATION COST AND RUNNING TIME OF THE DEFENSE STRATEGY IN ONE ITERATION

Methods	Total transmitted data (MB)					
	3 Users	5 Users	10 Users	15 Users	20 Users	25 Users
PROFL	18	31	63	96	129	162
ShieldFL	87	146	291	436	582	728
Methods	Running time					
	3 Users	5 Users	10 Users	15 Users	20 Users	25 Users
PROFL	4.79 s	9.68 s	34.08 s	75.95 s	135.23 s	211.95 s

that PROFL demonstrates better control over communication overhead compared to ShieldFL. This is attributed to the fact that in PROFL, only during the SecDis process do servers transmit ciphertexts of gradients, while in other phases, data transmitted between servers consists of intermediate variables in plaintext or a small amount of encrypted data.

During the security defense process, we optimize task execution using 12-core multi-threading technique. We tested the average runtime of each algorithmic components in PROFL under different numbers of participants, as shown in the Table III. In the case of 20 participants, after optimizing scheduling, the runtime for one iteration of the privacy-preserving defense strategy is 34.08 seconds. In the same setup, ShieldFL provides data of around 1.2 seconds. However, it should be noted that this data is obtained from running on devices different from ours. Compared to ShieldFL, although PROFL does not exhibit significant advantages in terms of communication and computation overheads, it provide greater robustness and security to federated learning in scenarios with a higher proportion of poisoning attackers or instances of covert poisoning attacks.

VI. CONCLUSION

In this paper, we propose PROFL to defend against various poisoning attacks while protecting privacy. We use a composite algorithm to identify and eliminate the interference from poisoning attacks. We use the two-trapdoor AHE algorithm to protect the privacy of the entire FL process. Security analysis and experimental evaluation show that PROFL is superior to similar works in terms of security and robustness. In future work, we will further explore methods to optimize efficiency and reduce the computational and communication costs.

ACKNOWLEDGMENT

This research was supported by the National Natural Science Foundation of China under Grant No. 62372446 and the National Key Research and Development Program of China under Grant No.2018YFA0704703.

REFERENCES

- [1] I. Dayan, H. R. Roth, *et al.*, "Federated learning for predicting clinical outcomes in patients with covid-19," *Nature medicine*, vol. 27, no. 10, pp. 1735–1743, 2021.
- [2] W. Zheng, L. Yan, *et al.*, "Federated meta-learning for fraudulent credit card detection," in *IJCAI*, pp. 4654–4660, 2020.
- [3] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," in *NeurIPS*, pp. 14747–14756, 2019.
- [4] P. Blanchard, E. M. E. Mhamdi, *et al.*, "Machine learning with adversaries: Byzantine tolerant gradient descent," in *NIPS*, pp. 119–129, 2017.
- [5] C. Xie, K. Huang, P. Chen, *et al.*, "DBA: distributed backdoor attacks against federated learning," in *ICLR*, 2020.
- [6] G. Baruch, M. Baruch, and Y. Goldberg, "A little is enough: Circumventing defenses for distributed learning," in *NeurIPS*, pp. 8632–8642, 2019.
- [7] X. Liu, H. Li, *et al.*, "Privacy-enhanced federated learning against poisoning adversaries," *IEEE Trans. Inf. Forensics Secur.*, vol. 16, pp. 4574–4588, 2021.
- [8] Z. Ma, J. Ma, *et al.*, "Shieldfl: Mitigating model poisoning attacks in privacy-preserving federated learning," *IEEE Trans. Inf. Forensics Secur.*, vol. 17, pp. 1639–1654, 2022.
- [9] X. Liu, R. H. Deng, *et al.*, "An efficient privacy-preserving outsourced calculation toolkit with multiple keys," *IEEE Trans. Inf. Forensics Secur.*, vol. 11, no. 11, pp. 2401–2414, 2016.
- [10] P. Mohassel and Y. Zhang, "Secureml: A system for scalable privacy-preserving machine learning," in *2017 IEEE symposium on security and privacy (SP)*, pp. 19–38, IEEE, 2017.
- [11] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *EUROCRYPT*, pp. 223–238, Springer, 1999.
- [12] P. Paillier and D. Pointcheval, "Efficient public-key cryptosystems provably secure against active adversaries," in *ASIACRYPT*, pp. 165–179, 1999.
- [13] D. Yin, Y. Chen, R. Kannan, and P. Bartlett, "Byzantine-robust distributed learning: Towards optimal statistical rates," in *International Conference on Machine Learning*, pp. 5650–5659, PMLR, 2018.
- [14] B. McMahan, E. Moore, *et al.*, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*, pp. 1273–1282, PMLR, 2017.
- [15] D. Catalano, R. Gennaro, and N. Howgrave-Graham, "The bit security of paillier's encryption scheme and its applications," in *EUROCRYPT*, pp. 229–243, Springer, 2001.
- [16] E. Bagdasaryan, A. Veit, *et al.*, "How to backdoor federated learning," in *International Conference on Artificial Intelligence and Statistics*, pp. 2938–2948, PMLR, 2020.