

Shielding Federated Learning: Robust Aggregation with Adaptive Client Selection

Wei Wan^{1,3,4,5}, Shengshan Hu^{1,3,4,5}, Jianrong Lu^{1,3,4,5}, Leo Yu Zhang⁷,
Hai Jin^{2,3,4,6} and Yuanyuan He^{1*}

¹School of Cyber Science and Engineering, Huazhong University of Science and Technology

²School of Computer Science and Technology, Huazhong University of Science and Technology

³National Engineering Research Center for Big Data Technology and System

⁴Services Computing Technology and System Lab

⁵Hubei Engineering Research Center on Big Data Security

⁶Cluster and Grid Computing Lab

⁷School of Information Technology, Deakin University

{wanwei_0303, hushengshan, lujianrong, hjin, yuanyuan_cse}@hust.edu.cn,
leo.zhang@deakin.edu.au

Abstract

Federated learning (FL) enables multiple clients to collaboratively train an accurate global model while protecting clients' data privacy. However, FL is susceptible to Byzantine attacks from malicious participants. Although the problem has gained significant attention, existing defenses have several flaws: the server irrationally chooses malicious clients for aggregation even after they have been detected in previous rounds; the defenses perform ineffectively against sybil attacks or in the heterogeneous data setting.

To overcome these issues, we propose MAB-RFL, a new method for robust aggregation in FL. By modelling the client selection as an extended *multi-armed bandit* (MAB) problem, we propose an adaptive client selection strategy to choose honest clients that are more likely to contribute high-quality updates. We then propose two approaches to identify malicious updates from sybil and non-sybil attacks, based on which rewards for each client selection decision can be accurately evaluated to discourage malicious behaviors. MAB-RFL achieves a satisfying balance between exploration and exploitation on the potential benign clients. Extensive experimental results show that MAB-RFL outperforms existing defenses in three attack scenarios under different percentages of attackers.

1 Introduction

Federated learning (FL) [McMahan *et al.*, 2017] is a newly emerging machine learning paradigm that aims to train a shared global model on decentralized data with user's privacy guarantee. A FL system consists of a central server and multiple clients. The clients do not share their raw data with the

central server, instead, each client trains a local model via fine-tuning the global model over its local data set and then sends the local model update (*i.e.*, the difference between local model and the global model) to the server. The server aggregates all the local updates to obtain a new global model and shares it with the clients for next iteration. FL could not only construct a model with high accuracy but also address the privacy concerns, thus drawing extensive research attention recently.

However, FL is susceptible to Byzantine attacks from malicious clients due to the distributed nature. For example, attackers can upload poisoned local updates via contaminating the training data (*i.e.*, data poisoning attacks [Bagdasaryan *et al.*, 2020; Tolpegin *et al.*, 2020]) or tampering with the local model updates directly (*i.e.*, model poisoning attacks [Baruch *et al.*, 2019; Shejwalkar and Houmansadr, 2021; Fang *et al.*, 2020]), in order to damage the global model or even leave a backdoor.

To mitigate Byzantine attacks, plenty of defensive schemes have been proposed. Their key ideas include eliminating updates that are far from the overall distribution through similarity detection [Blanchard *et al.*, 2017; Xia *et al.*, 2019], circumventing the malicious parameters through statistical characteristics [Yin *et al.*, 2018; Xie *et al.*, 2018; 2019a], detecting outliers based on the extracted features of updates [Tolpegin *et al.*, 2020; Shejwalkar and Houmansadr, 2021], relying a clean dataset to detect abnormality [Xie *et al.*, 2019b; Cao *et al.*, 2021], and directly clipping the local updates with large magnitudes [Karimireddy *et al.*, 2021].

Albeit tremendous in types, all the existing defenses suffer from one or multiple limitations as follows. Firstly, they paid little attention to the client selection step. In each iteration, the server is set to randomly select a subset of clients [Cao *et al.*, 2021] or simply choose all of them to participate in the aggregation [Blanchard *et al.*, 2017; Xia *et al.*, 2019; Karimireddy *et al.*, 2021]. Second, they cannot effectively defend against sybil attacks, in which all the malicious updates are similar or even identical [Baruch *et al.*, 2019; Shejwalkar and Houmansadr, 2021; Fang *et al.*, 2020]. Third,

*Corresponding Author

they perform badly in non-IID (*independently identically distribution*) scenarios, where all the updates are dramatically different from each other [Cao *et al.*, 2021; Prakash and Avestimehr, 2020]. Lastly, they assume that the server holds a validation dataset that has the same distribution with the training set of clients [Cao *et al.*, 2021; Xie *et al.*, 2019b; Prakash and Avestimehr, 2020], which obviously violates the privacy requirements of FL.

To address these issues, we propose a new multi-armed bandit based robust FL scheme, named MAB-RFL. Our scheme models the client selection process in federated learning as an extended MAB problem enabling the server to adaptively select updates that are more likely to be benign. Before applying existing MAB solutions to FL, there are two key challenges needed to be addressed. First, in standard MAB problem, only one arm is chosen at each time step, however, in FL scenarios, a mass of clients are selected to guarantee the high accuracy of the global model. Inspired by the Thompson sampling [Agrawal and Goyal, 2012], we propose an adaptive client selection strategy to decide the current round participants. Second, in the standard MAB scenario, the reward for each decision comes from the slot machine directly once an arm is played. Nevertheless, in FL the server has no idea how to evaluate each decision of client selection. Therefore, we propose two approaches to identify malicious updates from different attacks (*i.e.*, sybil and non-sybil), based on which rewards can be easily computed. Note that MAB methods have been applied in FL in other works, but they focus on reducing communication overheads [Cho *et al.*, 2020] or constructing high-quality models [Huang *et al.*, 2020], it is still unclear how to apply MAB algorithms to detect Byzantine attacks.

In summary, we make the following contributions:

- We model the client selection in Byzantine settings as a MAB problem and propose a novel adaptive client selection strategy, enabling the server to rationally select participants that are more likely to be honest and useful.
- We propose two new solutions to identify malicious clients from sybil and non-sybil attacks, such that the rewards for client selection decisions can be accurately evaluated to discourage misbehaviour.
- We conduct extensive experiments to evaluate MAB-RFL. The results show that the proposed scheme is effective for three attack scenarios under different percentages of attackers.

2 Background

2.1 Federated Learning

We consider a general FL system, consisting of a central server and K clients. Each client k ($k = 1, 2, \dots, K$) has a dataset D_k , the size of which is denoted as $|D_k| = n_k$. It is worth noting that each local dataset may be subject to a different distribution, that is, the clients' data may be distributed in a non-IID way. The clients aim to collaboratively train a shared global model w . Apparently, the problem can be solved via minimizing the empirical loss, *i.e.*, $\arg \min_w f(D, w)$, where $D = \bigcup_{k=1}^K D_k$ and $f(D, w)$ is a

loss function (*e.g.*, mean absolute error, cross-entropy). However, the optimization requires all the clients to share their raw data to a central server, which would result in a serious threat to client's privacy. Instead, FL obtains w by optimizing $\arg \min_w \sum_{k=1}^K f(D_k, w)$. Specifically, the FL system iteratively performs the following three steps until the global model converges:

- **Step I:** In the t -th iteration, the central server broadcasts a global model w_t to the clients;
- **Step II:** After receiving w_t , each client k trains a new local model w_t^k over D_k by solving the optimization problem $\arg \min_{w_t^k} f(D_k, w_t^k)$ and then uploads the local model update $g_t^k := w_t^k - w_t$ to the server;
- **Step III:** The server aggregates all the local updates according to client's proportional dataset size as follow:

$$w_{t+1} \leftarrow w_t + \sum_{k=1}^K \frac{n_k}{n} g_t^k, \text{ where } n = \sum_{k=1}^K n_k. \quad (1)$$

2.2 Multi-Armed Bandit Problem and Thompson Sampling

The *multi-armed bandit* (MAB) problem is a classical framework which studies the exploration/exploitation trade-off in sequential decision problems. Specifically, assuming that we are given a slot machine with K arms, at each time step $t = 1, 2, 3, \dots, T$, we play an arm $k(t) \in \{1, 2, \dots, K\}$ and observe the corresponding reward $R_{k(t)}$. In this work we consider the Bernoulli bandit, *i.e.*, $R_{k(t)}$ is either 0 or 1. The MAB aims to minimize the expected total regret over T rounds, where the regret is the amount we lose because of not playing optimal arm in each round. Formally, the T period expected cumulative regret is given by:

$$\mathbb{E}[\text{Regret}(T)] = TR_{k^*} - \mathbb{E}\left[\sum_{t=1}^T R_{k(t)}\right] = \sum_{t=1}^T \mathbf{1}\{k_t \neq k^*\}, \quad (2)$$

where k^* denotes the best arm (*i.e.*, $R_{k^*} = \max_k R_k$). The Thompson sampling algorithm [Agrawal and Goyal, 2012] is a popular solution to the MAB problem which models the mean reward of each arm using a Beta distribution and then chooses the arm with the largest mean reward.

Alg. 1 provides the details of Thompson sampling for the Bernoulli bandit. Note that S_k and F_k denote the number of successes and failures outcome of the arm k respectively, and both of them are initialized to 1, because Beta(1, 1) is the uniform distribution on (0, 1), which means that each arm has the same probability of being chosen at the beginning.

3 Problem Setup

3.1 Attack Model

We employ a similar attack model in previous works [Cao *et al.*, 2021; Prakash and Avestimehr, 2020]. To be specific, an adversary controls several clients (less than half of all participants), and it can poison the local data and updates of the controlled clients arbitrarily. However, the central server and benign clients are not under the control of the adversary. The

Algorithm 1 Thompson sampling for the Bernoulli bandit

```

1: for  $k = 1, 2, \dots, K$  do
2:    $S_k \leftarrow 1, F_k \leftarrow 1.$ 
3: end for
4: for  $t = 1, 2, \dots, T$  do
5:   for  $k = 1, 2, \dots, K$  do
6:     Draw  $\hat{R}_k$  according to  $\text{Beta}(S_k, F_k).$ 
7:   end for
8:   Draw arm  $k(t) \leftarrow \arg \max_k \hat{R}_k$  and observe reward
      $R_{k(t)}$ .
9:   if  $R_{k(t)} = 1$  then
10:     $S_{k(t)} \leftarrow S_{k(t)} + 1.$ 
11:   else
12:     $F_{k(t)} \leftarrow F_{k(t)} + 1.$ 
13:   end if
14: end for
    
```

only way the adversary can degrade the performance of the global model is by uploading poisoned local updates through compromised clients. We consider three kinds of attacks according to the adversary’s knowledge about the FL system:

Agnostic attack: the adversary knows neither benign updates nor the aggregation algorithm.

Partial knowledge attack: the adversary only has the knowledge of benign updates.


Full knowledge attack: the adversary knows both benign updates and aggregation algorithm.

Although the agnostic attack is the weakest type, it is the most realistic attack. Partial knowledge attack and full knowledge attack are too powerful to be applied in practice, but they have been widely investigated in previous studies [Cao *et al.*, 2021; Shejwalkar and Houmansadr, 2021; Karimireddy *et al.*, 2021; Mhamdi *et al.*, 2021] to fully understand the threats of Byzantine attacks.

3.2 Defense Model

Our defense is deployed on the central server, and does not require any privacy information about participants. Specifically, the central server does not access the raw local training data, and is unknown about the number of compromised clients and even the training dataset size of each client. Our defense is expected to achieve the following three goals simultaneously:

- **Rationality.** The method should avoid choosing clients who continually uploaded low-quality updates in previous iterations. Note that the low-quality updates may be caused by adversary’s meticulous design [Shejwalkar and Houmansadr, 2021; Fang *et al.*, 2020] or benign participants’ poor training data [Wan *et al.*, 2021].
- **Robustness.** The method should prevent malicious updates from degrading the accuracy or the convergence rate of the global model in all kinds of attack scenarios, *i.e.*, the agnostic attack, the partial knowledge attack, and the full knowledge attack.
- **Fidelity.** The method should not incur a drop in non-adversarial scenarios. In other words, the defense should be as accurate as FedAvg when there is no attack.



(a) Evaluation of sampling

(b) Evaluation of similarity

Figure 1: (a) The comparison between random sampling and rational sampling; (b) Maximum cosine similarity between benign updates.

To the best of our knowledge, none of existing Byzantine-robust schemes satisfy the rationality goal. They sample either a random subset of clients or all participants in each iteration. Besides, some schemes such as Krum, Median, FABA, and *Centered Clip* (CC) violate the robustness goal, because they perform poorly in full-knowledge settings. Moreover, Krum does not satisfy the fidelity goal.

4 MAB-RFL: Multi-Armed Bandit Based Robust Federated Learning

4.1 Motivation Behind Rational Client Selection

After reviewing the existing defenses against Byzantine attacks in FL, we find that they choose either all clients [Xia *et al.*, 2019; Blanchard *et al.*, 2017; Karimireddy *et al.*, 2021] or a random subset of them [Cao *et al.*, 2021]. This irrational client selection strategy may continually select malicious clients to participate in the aggregation, even after they have been identified in previous iterations, leading to a slow convergence and a waste of communication and computation resources. Fig. 1(a) gives a toy example. We consider a FL system containing 50 clients, among which 15 are malicious, and 10 clients will be selected at each iteration. Random sampling *A* randomly selects 10 clients and aggregates all the updates directly. Random sampling *B* also randomly selects 10 clients but discards all the malicious updates before aggregation. Rational sampling samples 10 benign clients and aggregates all the updates directly. The results show that random sampling *B* converges slower and the accuracy of the final global model is about 5% lower compared with the rational sampling. Therefore, a rational server should always choose honest clients who are more likely to submit high-quality updates.

4.2 A Complete Description of MAB-RFL

Alg. 2 shows our complete MAB-RFL scheme. In Step I (Lines 2-4), the server adaptively selects clients and broadcasts the current global model to the clients. In Step II (Lines 5-11), the server identifies sybil updates (Lines 5-8) as well as non-sybil but malicious updates (Lines 9-11). In Step III (Lines 12-14), the server assigns reward to each selected client. In Step IV (Lines 15-19), the server updates the global model based on the normalized momentum vectors of the remaining clients. Note that we set the learning rate as the average of the magnitudes of all the left updates (Line 16) to make

Algorithm 2 A Complete Description of MAB-RFL

Input: K clients; estimated maximum cosine similarity c_{max} and minimum cosine similarity c_{min} ; importance of historical information λ ; acceptable difference between clusters α .

Output: Global model w_{t+1} .

- 1: **Initialization:** $B_k \leftarrow 1, M_k \leftarrow 1, \forall k \in [K]$.
 - 2: //Step I: Select reliable clients.
 - 3: $L_{selected} \leftarrow ACS(\{B_k, k \in [K]\}, \{M_k, k \in [K]\})$
 - 4: Broadcast the current global model w_t to the clients in $L_{selected}$ and gather the corresponding local updates $g = \{g_k, k \in L_{selected}\}$.
 - 5: //Step II: Identify sybil attacks.
 - 6: $L_{sybil} \leftarrow ISA(L_{selected}, g)$.
 - 7: $L_{selected} \leftarrow L_{selected} - L_{sybil}$.
 - 8: $g \leftarrow \{g_k, k \in L_{selected}\}$.
 - 9: //Step II: Identify non-sybil attacks.
 - 10: $L_{nsybil} \leftarrow INSA(L_{selected}, g)$.
 - 11: $L_{selected} \leftarrow L_{selected} - L_{nsybil}$.
 - 12: //Step III: Distribute rewards.
 - 13: $M_k \leftarrow M_k + 1, \forall k \in L_{sybil} \cup L_{nsybil}$.
 - 14: $B_k \leftarrow B_k + 1, \forall k \in L_{selected}$.
 - 15: //Step IV: Aggregate updates.
 - 16: Calculate the learning rate:
 $\eta_t \leftarrow Mean(\{|g_k|, k \in L_{selected}\})$.
 - 17: Calculate the normalized momentum vectors:
 $mm \leftarrow \{m_t^k, k \in L_{selected}\}$ through Eq. (5) and Eq. (6).
 - 18: $w_{t+1} \leftarrow w_t + \eta_t \cdot Mean(mm)$.
 - 19: **return** w_{t+1}
-

up for the slow or fast learning caused by the normalization. Next, we present details for the first three steps.

4.3 Adaptive Client Selection

In the standard Thompson sampling algorithm, only the arm with the maximal mean reward (estimated by a Beta distribution) will be played at each time step. An intuitive idea for extending it to FL is to select clients with top- m (m is the number of benign clients) mean rewards. However, it relies on the assumption that the number of benign clients is fixed and known to the server. This assumption is unrealistic since FL is a dynamic distributed network where honest and compromised clients can drop out arbitrarily.

To get rid of the dependence on m , we propose selecting each client with a changing probability (*i.e.*, the probability varies with each client's past performance, rather than using a fixed probability) as shown in Alg. 3. Specifically, we set the probability of each client k being selected to be p_k , which is drawn from Beta(B_k, M_k) distribution (Lines 1-4). Here B_k and M_k denote the frequency that the local updates from client k are identified as benign and malicious (through our methods described in Section 4.4), respectively. Initially, we set $B_k = M_k = 1$. If none of the clients are selected in a round, we will randomly choose a subset of clients (Lines 5-8).

This method can achieve a well balance between exploration and exploitation. A client who consistently uploads benign updates (*i.e.*, a large B_k and a small M_k) will obtain a large p_k from Beta(B_k, M_k), and it is easier to be selected

Algorithm 3 ACS (Adaptive Client Selection)

Input: The frequencies that the local updates from each client $k \in [K]$ are regarded as benign B_k and malicious M_k .

Output: The selected clients list $L_{selected}$.

- 1: **for** $k = 1, 2, \dots, K$ **do**
 - 2: $p_k \leftarrow \text{Beta}(B_k, M_k)$.
 - 3: Add k to the selected clients list $L_{selected}$ at the probability p_k .
 - 4: **end for**
 - 5: **if** $L_{selected} = \emptyset$ **then**
 - 6: Assign a random nonempty subset of $[K]$ to $L_{selected}$.
 - 7: **end if**
 - 8: **return** $L_{selected}$
-

Algorithm 4 ISA (Identifying Sybil Attacks)

Input: The selected clients list $L_{selected}$ and the corresponding local updates $\{g_k, k \in L_{selected}\}$.

Output: Rejected clients list L_{sybil} .

- 1: Construct a graph $G = (V, E)$, where the vertex set $V = L_{selected}$, and the edge set E is obtained through Eq. (3).
 - 2: Assign the maximal connected subgraph of G to $G' = (V', E')$.
 - 3: $L_{sybil} \leftarrow V'$.
 - 4: **return** L_{sybil}
-

(exploitation), while a rarely selected client (*i.e.*, both B_k and M_k are small) may also get a large p_k (exploration), because the Beta distribution has a large variance when the parameters are small.

4.4 Abnormal Clients Identification

Identifying Sybil Attacks. Sybil attack is a recently proposed special Byzantine attack method [Fung *et al.*, 2020], where all the compromised clients upload totally identical local updates to launch covert attacks. This kind of attack poses a significant threat to existing defenses because existing defenses assume that benign updates are much similar and discard updates that are far from the overall distribution.

Before presenting our defensive method against sybil attacks, we introduce a basic fact that has been long misunderstood in the literature. In FL, the similarity between benign updates is as high as we expected when we use Euclidean distance as the evaluation metric. If we use cosine similarity as the evaluation metric, however, the benign updates have low similarity even in IID setting. This phenomenon is also observed by our experiments shown in Fig. 1(b), where the training data among clients are IID and each client has the same dataset size. We can see that after several iterations, the maximum cosine similarity reduces sharply to 0 (indicating a larger dissimilarity), while the Euclidean distances (computed after normalization) get smaller (indicating a larger similarity). We owe this phenomenon to the fact that when the model converges, benign gradients will gradually become zero, thus making their cosine similarities close to zero.

Based on the findings above, we propose mitigating sybil attacks by using undirected graph. Specifically, we regard




Figure 2: Illustration of sybil attack detection

Algorithm 5 INSA (Identifying Non-Sybil Attacks)

Input: The selected clients list $L_{selected}$ and the corresponding local updates $\{g_k, k \in L_{selected}\}$.

Output: Rejected clients list L_{nsybil} .

- 1: **for all** $k \in L_{selected}$ **do**
 - 2: Calculate the momentum vector m_t^k according to Eq. (5) and normalize m_t^k to be \overline{m}_t^k through Eq. (6).
 - 3: **end for**
 - 4: Extract the key features of each \overline{m}_t^k by PCA.
 - 5: Divide $L_{selected}$ into a larger cluster C_l and a smaller cluster C_s by applying agglomerative clustering algorithm to the key features.
 - 6: Calculate the mean momentum vector of each cluster:
 $m_l \leftarrow Mean(\{m_t^k, k \in C_l\})$;
 $m_s \leftarrow Mean(\{m_t^k, k \in C_s\})$.
 - 7: **if** $\cos(m_l, m_s) > \alpha$ **then**
 - 8: $L_{nsybil} \leftarrow \emptyset$.
 - 9: **else**
 - 10: $L_{nsybil} \leftarrow C_s$.
 - 11: **end if**
 - 12: **return** L_{nsybil}
-

each uploaded local update as a vertex and construct a graph $G = (V, E)$ through the following rule: if two updates are too similar in direction (in terms of cosine similarity), an undirected edge will be added. Formally, an undirected edge is added between vertex j and k if and only if the updates g_t^j and g_t^k satisfy the following property:

$$\cos(g_t^j, g_t^k) \geq \max\{c_{max}e^{\frac{1-t}{20}}, c_{min}\}, \quad (3)$$

where \cos denotes the cosine similarity, c_{max} and c_{min} ($-1 < c_{min} < c_{max} < 1$) represent the estimated maximum and minimum cosine similarities, and t is the current iteration. Note that the tolerable cosine similarity (right of the above inequality) decreases as iterations go on, which indicates that it is harder for the adversary to launch covert sybil attack.

After constructing the graph, the updates (vertexes) in the maximal connected subgraph $G' = (V', E')$ of G will be considered as malicious ones. Fig. 2 gives a simple illustration for this process and Alg. 4 shows the details.

Identifying Non-sybil Attacks. In non-sybil settings where malicious updates obviously differ from each other, it becomes difficult to identify them. A pivotal reason for

the difficulty lies in the extremely high dimension of local models, which contains deciding features as well as redundant features that will mask the differences between benign and malicious updates. Moreover, the non-IID property of FL naturally incurs disparity among benign updates, making it more difficult to find out abnormal ones. In light of this, we propose squeezing the feature space before filtering out updates. In order to counter the non-IID issue, we further make use of Momentum [Mhamdi *et al.*, 2021] to reduce the variance between updates such that an IID-like setting can be built. Specifically, our approach contains three steps as shown in Alg. 5: 1) calculating the momentum vector of each client (Lines 1-3); 2) using *principal component analysis* (PCA) to extract the key features of each momentum vector (Line 4); 3) applying agglomerative clustering algorithm to divide the momentum vectors and choosing the appropriate vectors for aggregation (Lines 5-12).

Traditionally, the momentum vector of client k at iteration t is given by:

$$m_t^k = g_t^k + \lambda m_{t-1}^k, \quad (4)$$

where $\lambda \in (0, 1)$ indicates the importance of historical information, g_t^k is the local update, and the initial momentum m_0^k is generally set to be 0. Nevertheless, it is not suitable for our case because clients are not selected continuously in FL. We thus slightly revise the momentum vector calculation as:

$$m_t^k = g_t^k + \lambda^{t-t_k} m_{t_k}^k, \quad (5)$$

where t_k is the latest iteration when client k is selected. The intuition behind Eq. (5) is that the historical information $m_{t_k}^k$ of a chronically unselected client is obsolete, thus multiplied by a smaller discount factor. Note that Eq. (5) and Eq. (4) are equivalent without using client selection strategy (*i.e.*, $t_k = t - 1$).

Besides, attackers may upload updates with large magnitudes to dominate the global update. We thus normalize the momentum vector as:

$$\overline{m}_t^k = \frac{m_t^k}{\|m_t^k\|}. \quad (6)$$

This step ensures that a single malicious update has limited impact on the aggregated global update even though it could circumvent the detection.

In order to prevent MAB-RFL from discarding updates excessively, we set a threshold α to determine if all the momentum vectors will be used for aggregation (Lines 7-9).




Figure 3: Model accuracy under label flipping attack




Figure 4: Model accuracy under LIE attack

4.5 Reward Distribution Mechanism

A correct reward distribution mechanism is important for MAB-RFL to reduce the possibility of selecting abnormal clients in subsequent iterations. As shown in Alg. 2, after identifying abnormal clients who conduct sybil and non-sybil attacks, they will get a reward 0 (*i.e.*, $M_k \leftarrow M_k + 1$, Line 13) and be discarded in the current round (Line 7 and Line 11). The remaining selected clients get a reward 1 (*i.e.*, $B_k \leftarrow B_k + 1$, Line 14) and their updates are used for final aggregation (Lines 15-19).

5 Experiments

5.1 Experimental Setup

Datasets and Models. We evaluate MAB-RFL on MNIST and CIFAR-10. For MNIST, we use a *convolutional neural network* (CNN) with 2 convolutional layers followed by 2 fully connected layers as the global model. For CIFAR-10, we employ a simplified version of VGG16 as the global model.

Data Distribution. Following previous works [Cao *et al.*, 2021; Shejwalkar and Houmansadr, 2021], we assign each client a random dominated label l . Then each training sample with label l is distributed to the client with probability $q \in [0, 1]$ and any other label with probability $\frac{1-q}{L-1}$, where L is the total kinds of labels and q is the non-IID degree. A larger q indicates a higher degree of non-IID. Note that $q = \frac{1}{L}$ corresponds to the IID scenario. In our experiments, we set $q = 0.5$ by default, which is the highest non-IID degree existing works considered. We assign different training set sizes for clients. Specifically, the training set size of each client is randomly chosen from 10 to 500 for MNIST and from 1200 to 2000 for CIFAR-10.




Figure 5: Model accuracy under AGRT attack

Evaluated Attacks. We consider *label flipping* (LF) attack [Tolpegin *et al.*, 2020], *little is enough* (LIE) attack [Baruch *et al.*, 2019], and *AGR-tailored* (AGRT) attack [Shejwalkar and Houmansadr, 2021], which belong to agnostic attack, partial knowledge attack, and full knowledge attack, respectively.

LF attack. For each controlled client, the adversary flips the local sample labels to generate misleading update. Specifically, the adversary flips the label l of each training sample to $L - l - 1$, where L is the number of classes in the classification problem and $l = 0, 1, \dots, L - 1$.

LIE attack. The adversary first estimates coordinate-wise mean μ_j and deviation σ_j through benign updates. Then the malicious update \tilde{g} can be crafted as follow:

$$(\tilde{g})_j = \mu_j - z^{max} \cdot \sigma_j, j \in [d], \quad (7)$$

where d is the total dimension of model update, and the positive attack factor z^{max} determines the amount of noise added. A large z^{max} makes the constructed update harmful to the global model but easy to be detected by the central server, while a small one is covert but less baleful. The original paper [Baruch *et al.*, 2019] advises that z^{max} can be obtained by the Cumulative Standard Normal Function $\phi(z)$:

$$z^{max} = \max_z \left(\phi(z) < \frac{K - \lfloor \frac{K}{2} + 1 \rfloor}{m} \right). \quad (8)$$

Note that all the controlled clients upload identical malicious update \tilde{g} to the central server, making the attack inconspicuous.

AGRT attack. AGRT attack aims to maximize the damage to the global model in each FL iteration in an optimized manner. Formally, the adversary performs the following optimization problem to construct malicious updates:

$$\begin{aligned} \operatorname{argmax}_{\gamma} \quad & \|\hat{g} - \mathcal{A}(g_{\{i \in [m]\}} \cup \tilde{g}_{\{i \in [m+1, K]\}})\| \\ \tilde{g}_{\{i \in [m+1, K]\}} = & \hat{g} + \gamma \cdot \tilde{\Delta}; \hat{g} = \operatorname{FedAvg}(g_{\{i \in [m]\}}), \end{aligned} \quad (9)$$

where \mathcal{A} is the known defense method (*i.e.*, Krum) applied by the central server, $g_{\{i \in [m]\}}$ are the benign updates that the adversary knows (without loss of generality, we assume the first m clients are benign and the rest are controlled by the adversary), \hat{g} is a reference benign aggregation obtained by FedAvg [McMahan *et al.*, 2017] that averages all the benign updates that the adversary knows, $\tilde{\Delta}$ is a perturbation vector, γ is a scaling coefficient, $\tilde{g}_{\{i \in [m+1, K]\}}$ are the final malicious updates.

It is worth noting that the above optimization incurs totally identical malicious updates for all controlled clients, thus MAB-RFL can easily discard all the malicious updates through Alg. 4. Nevertheless, considering that the adversary has the knowledge of the defense scheme in advance, we set the updates for half of the controlled clients to \tilde{g} and the other half to be the same as any benign update. This ensures that half of the malicious updates can circumvent the sybil detection and that one benign update is removed.

Evaluated Defenses. We compare MAB-RFL with three classical defenses, *i.e.*, Krum [Blanchard *et al.*, 2017], FABA [Xia *et al.*, 2019], and Median [Yin *et al.*, 2018], as well as two newly proposed defenses, *i.e.*, DnC [Shejwalkar and Houmansadr, 2021] and *Centered Clip* (CC) [Karimireddy *et al.*, 2021].

Krum. Krum calculates the Euclidean distance between any two local gradients and selects the one that is closest to the $m - 2$ neighboring local gradients.

FABA. FABA removes the local update that is farthest from the average of the local updates repeatedly until the number of eliminated updates reaches a predefined threshold.

Median. Median directly takes the coordinate-wise median value in each dimension of all local updates as the new global update.

DnC. DnC leverages *singular value decomposition* (SVD) to extract the common features between benign and poisoned updates, and randomly samples a subset of parameters of each local update as its substitution, which will be projected along their top right singular eigenvector. Then an outlier score is obtained by computing the inner product of substitution and the projection, and the $K - m$ local updates with the highest scores will be removed.

CC. CC clips the local updates with large magnitudes, with the intuition that attackers may upload such updates to dominate the global model.

Evaluation Metrics and Parameter Setting. We use *accuracy*, which is the proportion of correctly predicted testing samples to total testing samples, to evaluate the final global model. We set the number of clients $K = 50$ for both datasets. To reduce the total communication rounds between clients and the server, we set the local epoch of each client to be 3. The total iteration $T = 100$. The importance of historical information $\lambda = 0.1$. For MNIST, we set the estimated maximum cosine similarity $c_{max} = 0.7$, minimum cosine similarity $c_{min} = 0.3$, and the acceptable difference between clusters $\alpha = -0.1$. For CIFAR-10, we set $c_{max} = 0.3$, $c_{min} = 0.1$, $\alpha = 0$. For AGRT attack, we choose $\tilde{\Delta}$ as $-std(g_{\{i \in [m]\}})$, *i.e.*, the inverse standard deviation.

Comparison with the Existing Defenses. MAB-RFL achieves the three defense goals as shown in Table 1.

Rationality: MAB-RFL performs sampling with a more sophisticated method (*i.e.*, rational sampling) than the naive methods used in previous works, which improves the probability of selecting benign clients. The toy example in section 4.1 well justifies this.

Robustness: MAB-RFL achieves the robustness objective, *i.e.*, shielding FL from adversaries. As shown in Table 1, MAB-RFL outperforms the existing defenses on three strong

Dataset (Model)	Percentage of Attackers	Attack Type	Accuracy (%)							
			CC	DnC	FABA	Krum	Median	MAB-RFL	FedAvg	
CIFAR-10 (VGG16)	0%	—	76.57	75.25	75.17	63.40	75.07	76.33	77.02	
		AGRT	76.39	70.68	62.36	10.00	65.33	76.25	—	
		LIE	72.45	74.26	71.87	61.26	71.56	75.18	—	
	10%	LF	69.86	75.07	73.47	50.27	74.86	76.09	—	
		AGRT	54.06	70.38	60.25	10.00	62.43	73.70	—	
		LIE	67.34	72.17	68.62	56.45	62.93	75.72	—	
	20%	LF	64.62	66.57	70.33	46.25	71.19	77.27	—	
		AGRT	10.00	70.28	10.00	10.00	48.49	70.35	—	
		LIE	61.67	74.14	59.83	38.81	52.91	74.84	—	
	30%	LF	61.15	67.16	68.33	47.61	68.03	75.61	—	
		AGRT	10.00	65.22	10.00	10.00	10.00	67.35	—	
		LIE	53.01	70.28	10.00	10.00	10.00	72.90	—	
	40%	LF	53.71	56.54	56.00	38.61	60.75	73.28	—	
		AGRT	10.00	58.08	10.00	10.00	10.00	67.53	—	
		LIE	34.67	70.92	10.00	10.00	10.00	74.14	—	
	44%	LF	46.89	60.22	45.66	30.52	52.93	74.75	—	
		AGRT	10.00	60.35	10.00	10.00	10.00	67.90	—	
		LIE	10.00	40.59	10.00	10.00	10.00	73.68	—	
	48%	LF	40.42	67.45	34.29	26.82	46.41	73.78	—	
		—	97.45	97.53	97.68	95.56	97.76	97.82	98.01	
		AGRT	93.81	97.44	96.96	88.44	97.89	97.77	—	
	MNIST (CNN)	0%	LIE	97.36	97.58	95.56	95.50	97.70	97.83	—
			LF	95.51	97.65	95.61	67.41	96.03	97.11	—
			AGRT	73.97	96.96	92.83	63.23	95.16	97.29	—
10%		LIE	97.05	97.43	95.00	93.64	96.21	97.54	—	
		LF	91.09	96.94	95.06	65.87	96.36	96.68	—	
		AGRT	9.74	97.32	10.10	9.58	90.30	97.24	—	
20%		LIE	97.00	97.23	93.63	79.73	93.68	97.05	—	
		LF	83.77	96.64	94.00	59.45	95.09	97.10	—	
		AGRT	10.09	95.46	10.32	10.12	86.09	96.70	—	
30%		LIE	94.71	96.63	86.82	60.97	70.07	97.01	—	
		LF	69.57	92.83	92.27	45.49	92.05	97.03	—	
		AGRT	9.61	95.39	9.82	10.07	65.21	96.61	—	
40%		LIE	93.11	96.45	63.73	50.49	62.28	96.97	—	
		LF	58.12	89.39	70.37	50.37	80.32	96.48	—	
		AGRT	10.32	93.16	10.09	9.93	30.44	96.09	—	
48%		LIE	88.85	91.70	10.51	9.96	33.85	96.86	—	
		LF	48.10	77.91	63.94	3.68	64.47	95.26	—	

Table 1: A comparison with existing defense methods on three attacks over two non-IID datasets

attacks over two non-IID dataset. For instance, in CIFAR-10, all the attacks pose no impact on MAB-RFL.

Fidelity: The accuracy of MAB-RFL under all attacks is always comparable to that of baseline (*i.e.*, FedAvg) on two datasets, while the accuracy of Krum and Median are reduced even without attackers, which reflects that MAB-RFL is applicable for non-adversarial scenarios.

Impact of the Percentage of Attackers. Fig. 3, Fig. 4, and Fig. 5 show that MAB-RFL always performs better or similar compared with state-of-the-art defenses in all settings. CC, FABA, and Median perform slightly worse than MAB-RFL when there are no more than 30% attackers, however, their accuracy drops dramatically with 40% attackers. DnC performs much better, but its accuracy drops obviously when there are nearly 50% attackers.

6 Conclusion

We propose a new robust federated learning scheme to make the server perform rational client selection decisions. Two approaches are designed to identify malicious updates from sybil and non-sybil attacks, based on which rewards for each client selection decision can be accurately evaluated. Extensive experimental results show that MAB-RFL outperforms existing defenses.

Acknowledgments

Shengshan’s work is supported in part by the National Natural Science Foundation of China (Grant No. U20A20177). Leo’s work is supported in part by the National Natural Science Foundation of China (Grant No. 61702221). Yuanyuan’s work is supported in part by the National Natural Science Foundation of China (Grant No. 62002127).

References

- [Agrawal and Goyal, 2012] Shipra Agrawal and Navin Goyal. Analysis of thompson sampling for the multi-armed bandit problem. In *Proceedings of the 25th Annual Conference on Learning Theory (COLT'12)*, pages 39.1–39.26, 2012.
- [Bagdasaryan *et al.*, 2020] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS'20)*, volume 108, pages 2938–2948, 2020.
- [Baruch *et al.*, 2019] Gilad Baruch, Moran Baruch, and Yoav Goldberg. A little is enough: Circumventing defenses for distributed learning. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS'19)*, pages 8632–8642, 2019.
- [Blanchard *et al.*, 2017] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems (NeurIPS'17)*, pages 119–129, 2017.
- [Cao *et al.*, 2021] Xiaoyu Cao, Minghong Fang, Jia Liu, and Neil Zhenqiang Gong. Fltrust: Byzantine-robust federated learning via trust bootstrapping. In *Proceedings of the 28th Annual Network and Distributed System Security Symposium (NDSS'21)*, 2021.
- [Cho *et al.*, 2020] Yae Jee Cho, Samarth Gupta, Gauri Joshi, and Osman Yagan. Bandit-based communication-efficient client selection strategies for federated learning. In *Proceedings of the 54th Asilomar Conference on Signals, Systems, and Computers (ACSCC'20)*, pages 1066–1069, 2020.
- [Fang *et al.*, 2020] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. Local model poisoning attacks to byzantine-robust federated learning. In *Proceedings of the 29th USENIX Security Symposium (USENIX Security'20)*, pages 1605–1622, 2020.
- [Fung *et al.*, 2020] Clement Fung, Chris J. M. Yoon, and Ivan Beschastnikh. The limitations of federated learning in sybil settings. In *Proceedings of the 23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID'20)*, pages 301–316, 2020.
- [Huang *et al.*, 2020] Tiansheng Huang, Weiwei Lin, Keqin Li, and Albert Y. Zomaya. Stochastic client selection for federated learning with volatile clients. *CoRR*, abs/2011.08756, 2020.
- [Karimireddy *et al.*, 2021] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. Learning from history for byzantine robust optimization. In *Proceedings of the 38th International Conference on Machine Learning (ICML'21)*, volume 139, pages 5311–5319, 2021.
- [McMahan *et al.*, 2017] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS'17)*, volume 54, pages 1273–1282, 2017.
- [Mhamdi *et al.*, 2021] El Mahdi El Mhamdi, Rachid Guerraoui, and Sébastien Rouault. Distributed momentum for byzantine-resilient stochastic gradient descent. In *Proceedings of the 9th International Conference on Learning Representations (ICLR'21)*, 2021.
- [Prakash and Avestimehr, 2020] Saurav Prakash and Amir Salman Avestimehr. Mitigating byzantine attacks in federated learning. *CoRR*, abs/2010.07541, 2020.
- [Shejwalkar and Houmansadr, 2021] Virat Shejwalkar and Amir Houmansadr. Manipulating the byzantine: Optimizing model poisoning attacks and defenses for federated learning. In *Proceedings of the 28th Annual Network and Distributed System Security Symposium (NDSS'21)*, 2021.
- [Tolpegin *et al.*, 2020] Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data poisoning attacks against federated learning systems. In *Proceedings of the 25th European Symposium on Research in Computer Security (ESORICS'20)*, volume 12308, pages 480–501, 2020.
- [Wan *et al.*, 2021] Wei Wan, Jianrong Lu, Shengshan Hu, Leo Yu Zhang, and Xiaobing Pei. Shielding federated learning: A new attack approach and its defense. In *Proceedings of IEEE Wireless Communications and Networking Conference (WCNC'21)*, pages 1–7, 2021.
- [Xia *et al.*, 2019] Qi Xia, Zeyi Tao, Zijiang Hao, and Qun Li. FABFA: an algorithm for fast aggregation against byzantine attacks in distributed neural networks. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI'19)*, pages 4824–4830, 2019.
- [Xie *et al.*, 2018] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Generalized byzantine-tolerant sgd. *Journal of Environmental Sciences (China) English Ed*, 2018.
- [Xie *et al.*, 2019a] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. SLSGD: secure and efficient distributed on-device machine learning. In *Proceedings of Machine Learning and Knowledge Discovery in Databases - European Conference (ECML PKDD'19)*, volume 11907, pages 213–228, 2019.
- [Xie *et al.*, 2019b] Cong Xie, Sanmi Koyejo, and Indranil Gupta. Zeno: Distributed stochastic gradient descent with suspicion-based fault-tolerance. In *Proceedings of the 36th International Conference on Machine Learning (ICML'19)*, volume 97, pages 6893–6901, 2019.
- [Yin *et al.*, 2018] Dong Yin, Yudong Chen, Kannan Ramchandran, and Peter L. Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *Proceedings of the 35th International Conference on Machine Learning (ICML'18)*, volume 80, pages 5636–5645, 2018.