

Multi-metrics adaptively identifies backdoors in Federated learning

Siquan Huang¹

cssiquanhuang@mail.scut.edu.cn

Yijiang Li²

yli556@jhu.edu

Chong Chen¹

cschenchong@mail.scut.edu.cn

Leyu Shi¹

csshileyu@mail.scut.edu.cn

Ying Gao^{1*}

gaoying@scut.edu.cn

¹School of Computer Science and Engineering, South China University of Technology

²Johns Hopkins University

Abstract

The decentralized and privacy-preserving nature of federated learning (FL) makes it vulnerable to backdoor attacks aiming to manipulate the behavior of the resulting model on specific adversary-chosen inputs. However, most existing defenses based on statistical differences take effect only against specific attacks, especially when the malicious gradients are similar to benign ones or the data are highly non-independent and identically distributed (non-IID). In this paper, we revisit the distance-based defense methods and discover that i) Euclidean distance becomes meaningless in high dimensions and ii) malicious gradients with diverse characteristics cannot be identified by a single metric. To this end, we present a simple yet effective defense strategy with multi-metrics and dynamic weighting to identify backdoors adaptively. Furthermore, our novel defense has no reliance on predefined assumptions over attack settings or data distributions and little impact on benign performance. To evaluate the effectiveness of our approach, we conduct comprehensive experiments on different datasets under various attack settings, where our method achieves the best defensive performance. For instance, we achieve the lowest backdoor accuracy of 3.06% under the most difficult Edge-case PGD, showing significant superiority over previous defenses. The experiments also demonstrate that our method can be well-adapted to a wide range of non-IID degrees without sacrificing the benign performance.

1. Introduction

Federated learning (FL) [26, 35] is a distributed machine learning paradigm that enables multiple participants to train a quality model collaboratively without exchanging their local data. During each round of training, the central server

distributes the global model to a subset of clients. Each client updates the model with the local data and submits the parameters to the server for aggregation. By training an efficient and quality model in a decentralized manner without sacrificing the privacy of participants, FL alleviates the possible conflict between technological developments and regulations for data privacy protection (e.g., General Data Protection Regulation). FL has already been applied and achieved success in multiple fields, such as image processing [32], word prediction [36], medical imaging [30], and edge computing [56, 51].

However, FL is vulnerable to backdoors manipulating the model towards the targeted behavior on specific adversary-chosen inputs [24, 4, 6, 47, 50, 11, 33, 39, 19, 31]. The backdoor attack is more difficult to be detected compared with the untargeted poisoning attacks [6, 7, 33, 11, 19] since it does not affect the regular function of the model and its gradient is more similar to benign ones [47, 38]. Multiple defenses have been proposed to improve the robustness of FL, such as the scoring-based methods [17, 8, 53, 41, 10, 55, 16, 29, 2, 20], which leverage a particular metric to distinguish malicious gradients from the benign ones. Despite its effectiveness against some backdoors, researchers discovered that well-designed attacks (termed *stealthy backdoors*) whose gradients are indistinguishable from benign ones (through scaling [4, 47] or trigger split [50]) could easily bypass these defenses. Differential privacy (DP)-based method [46, 40, 38, 49] builds upon the observation that the DP method [14], traditionally used against DP attacks, is also effective against backdoors. By adding Gaussian noise to the global model, these methods can dilute the impact of potentially poisoned model updates. Surprisingly, DP-based methods show great ability in resisting *stealthy backdoors* (e.g., Edge-case PGD [47]). Despite its ability to resist the *stealthy backdoors*, the noises added by the DP significantly decrease the overall performance and the convergence speed. In comparison, the distance-based method less

*Corresponding Author

impacts the global model by aggregating only the benign gradients. Consequently, a natural question arises: *can we defend the stealthy backdoors without sacrificing the performance of the FL model?* To accomplish this, we turn to the distance-based methods that don't sacrifice the benign performance and promote the research question: *how can we successfully leverage distance metrics to discriminate hostile updates from benign ones?*

To this end, we revisit the distance-based defense and discover two limitations: 1. Euclidean distance (*i.e.*, L_2 distance) suffers from the curse of dimensionality. Parameters of Neural Networks (NNs) can be viewed as high-dimensional vectors, and Euclidean distance generally fails to discriminate between malicious and benign gradients in high-dimensional space. 2. Single metric takes effect only against particular attacks with detailed assumptions regarding the malicious gradients. For instance, cosine distance detects malicious gradients with sizeable angular deviations while euclidean distance detects malicious gradients with a large L_2 norm scaled by the attacker to impact the global models. Moreover, backdoor attacks are conducted with different data and scenarios, resulting in malicious gradients with various characteristics that a single metric cannot handle (*i.e.*, both gradients with a large norm and gradients with sizeable angular deviations exist in one round of aggregation). What's worse is that the defender has no knowledge regarding the attacker and the underlying data distributions due to privacy requirements by FL, which makes detecting with a single metric even more difficult.

To address the above two problems, we first introduce the Manhattan distance, which we theoretically prove more meaningful in high dimensional space than Euclidean distance. Empirically, with the Manhattan distance, our proposed distance-based defense shows remarkable performance against the *stealthy backdoors*. To cope with gradients with various properties, we leverage multiple metrics cooperatively to identify the malicious gradients. We demonstrate in Section 5.4 and Figure 7 that malicious gradients of some characteristics are better identified by specific metrics, which justify our motivation. To handle the different attacks and environments, we further propose to apply a whitening transformation and generate dynamic weights to handle the non-IID distribution of participants and different scales brought by the different distances. Finally, we compute the score for each submitted gradient and aggregate only the benign ones based on the scoring. Extensive experiments illustrate that our defense maintains a high model performance and robustness simultaneously, which has never been achieved before. At a high level, we summarize the main contributions of this paper as follows:

- We present a novel defense with multi-metrics to adaptively identify backdoors, which is applicable in a generic adversary model without predefined assumptions

over the attack strategy or data distribution.

- We show that by introducing the Manhattan distance, our defense alleviates the “meaningfulness” problem of Euclidean distance in high dimensions. By utilizing multiple metrics with dynamic weighting, our defense can resist backdoor attacks under a wide range of attack settings and data distributions.
- We empirically evaluate our method on various datasets and different attack settings. By showing that our defense maintains both high robustness and benign performance under the *stealthy backdoors* that previously have not been successfully defended, we demonstrate the effectiveness of our approach.

2. Related work

Backdoor Attacks on Federated Learning. FL enables multiple clients collaboratively train a model without local data exchange. In an FL system, clients update the local model with private data and submit the parameters to the server for aggregation. Given a total of N participants, each with $n^{(i)}$ samples. In each round t , the server randomly selects K clients to participate. FedAvg [35] is a widely-adopted FL algorithm, which can be formulated as:

$$\mathbf{w}_t = \mathbf{w}_{t-1} + \eta \cdot \frac{\sum_{i=1}^K n^{(i)} \Delta \mathbf{w}_t^{(i)}}{\sum_{i=1}^K n^{(i)}} \quad (1)$$

where \mathbf{w}_t is the global model parameters at round t , $\Delta \mathbf{w}_t^{(i)}$ is the local model updates from the client i in round t , and η is the global model learning rate.

Empirical studies from Bagdasaryan et al. [4] demonstrate that FedAvg [35] is vulnerable to backdoor attacks since FL has no authority over local data and the training process due to privacy concerns. The model replacement attack [4] successfully injects a backdoor into the global model by the single-shot attack. Some well-designed attacks strategically target the weakness of the defenses, such as the PGD attack [47] scaling and projecting the gradients, or the DBA attack [50] splitting the triggers before uploading them. Moreover, the Edge-case PGD attack [47] modifies the poisoned data and model. Undoubtedly, these attacks present a tremendous challenge to the security of the FL system.

Defense Against the Backdoor Attack. Various methods to defend against backdoor attacks improve the robustness and security of FL. In general, we can divide these methods into two categories. The first type of defense strategies tries to distinguish malicious gradients from benign ones by classifying or scoring with the assumption that benign and malicious gradients show distinct differences in the vector space [17, 8, 53, 41, 10, 55, 16, 29, 2, 20]. Krum

and Multi-Krum [8] select the gradient with the smallest Euclidean norm from the others to aggregate and update the global model at each round. Foolsgold [17] assumes that the gradients of attackers are always consistent and distinguishes them by the Cosine similarity between historical updates. Trimmed Mean and Median [53] remove a fraction of the gradients with maximum and minimum values and then take the median value on each dimension of the remaining gradients as the parameters of the global model. RFA [41] leverages the geometric median of the gradients to keep the aggregated model close to the benign model and away from the backdoor one. The strength lies in that, if successful, the gradients of attackers are excluded from the aggregation process, which solves the core and fundamental problem of defending the attack.

However, the above methods take effect only under particular attacks since they make detailed assumptions about the attack or the data distributions. Foolsgold [17] assumes the gradients of attackers are consistent while benign gradients from Non-IID data differ. Krum [8] assumes that the data of benign clients is similar under IID settings and excludes abnormal malicious ones. These predefined assumptions limit the applicability of these defenses under various attacks and data distributions and can be easily bypassed by elaborate attacks. For example, PGD [47] breaks any Euclidean-based defense (*e.g.* Krum and Multi-Krum) by scaling the norm to a small value. Other methods, such as FLTrust [10], use a root dataset and aggregate only those gradients with large cosine similarities with the root dataset. FedInv [55] inverts the gradients into dummy datasets and then identifies the malicious gradient by comparing these dummy datasets. These methods are limited as well. FLTrust cannot defend against backdoors crafted on edge cases (*e.g.* Edge-case PGD) since they appear less often and are probably not in the root dataset. Inverting the gradients into dummy datasets violates the privacy requirement of FL. Essentially, this method can be characterized as another type of attack [3, 21].

The other category of defenses [46, 40, 38, 49] builds upon the observation that the differential privacy (DP) method [14], traditionally used against DP attacks is also effective against backdoor attacks. Earlier studies [4] have demonstrated that clipping and adding noise (typically used in DP) can resist backdoor attacks without any specific assumptions over the data distribution. Following these findings, Weak-DP [46] clips the norm and adds the Gaussian noise to the global model for mitigating the backdoors. RLR [40] and Flame [38] apply the clustering methods before the DP process. However, the noise added significantly decreases the performance and the convergence speed of the FL system. In a nutshell, using DP sacrifices regular functionality for robustness. Despite its deficiency, the DP-based method has been the only method that has taken effect

against *stealthy backdoors* [38].

3. Methodology

We first provide the threat model in Section 3.1 and then detail our design of the proposed defense method. Our goal is to design an efficient defense with the following properties: (1) Effectiveness: To prevent the adversary from achieving its hostile purpose, we aim to design a defense that identifies as many malicious updates as possible to eliminate its impact on the global model. (2) Performance: Benign performance of the global model must be preserved to maintain its utility. (3) Independence: The defense must be applicable to generic adversary models under different data distributions and attack strategies. Our defense consists of three major components: defining and calculating the gradient feature, computing the dynamic weight and score for each update, and aggregating the benign gradients with the score computed, as shown in Figure 1 and Algorithm 1.

3.1. Threat Model

Adversarial goals. Adversary aims to alter the behavior of the model on some specific examples with no negative impact on the benign performance. We consider the adversary effective if this is met. In addition, the adversary will make their attack as covert as possible, *e.g.*, by scaling down. Formally, adversary \mathcal{A} manipulates the global model G to the poisoned model G' . For any inputs x , we have the outputs below:

$$G'(x) = \begin{cases} l' \neq G(x) & \forall x \in D_{\mathcal{A}} \\ G(x) & \forall x \notin D_{\mathcal{A}} \end{cases} \quad (2)$$

where $D_{\mathcal{A}}$ denotes the poisoned dataset belonging to adversary \mathcal{A} , and l' denotes the specific label modified by adversary \mathcal{A} of the input x .

Adversarial capabilities. This paper assumes the white-box setting with full knowledge of aggregation operations, including potential defenses. Also, adversaries can freely alter the parameters in the local training, ensure participation in each round and compromise less than 50% of total FL clients. However, the adversary cannot control any process or parameter executed in the aggregator.

3.2. Feature of gradient

The essential logic of distance-based defense involves defining some indicative metric that can well discriminate the malicious gradients from the benign ones and removes the hostile updates from the aggregation. Consequently, the core problem becomes how to define a metric that identifies the characteristics of hostile gradients. For instance, NDC [46] utilizes the length of the gradient (*i.e.* Euclidean distance) while FoolGold [17] leverages Cosine similarity to measure the difference between updates.

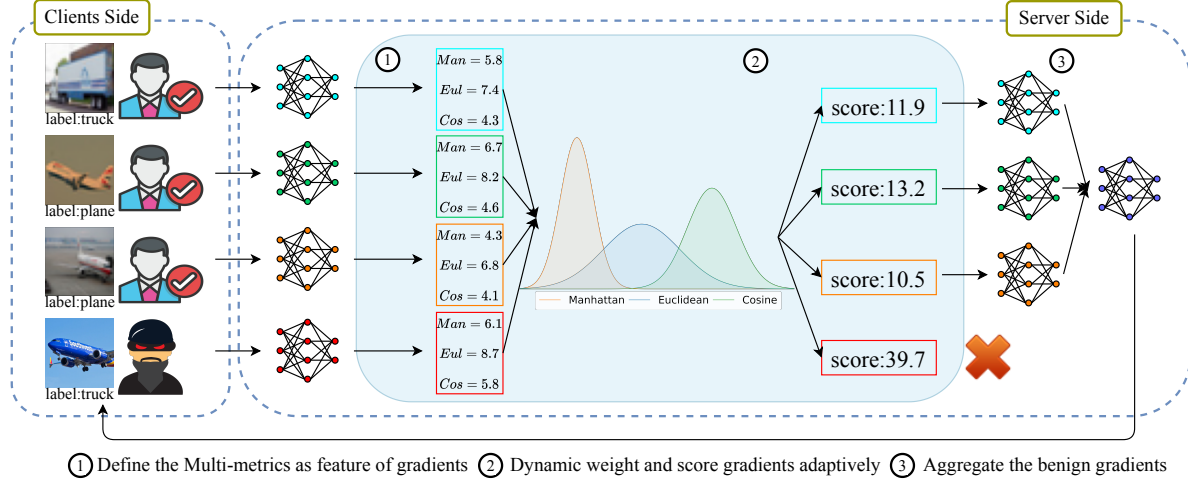


Figure 1: The overview of our defense process. Step 1 and step 2 are our core contributions.

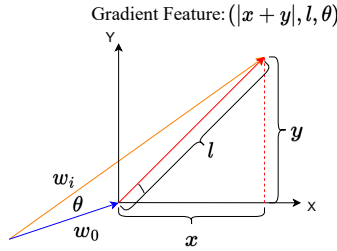


Figure 2: Demonstration of *feature of gradient* on two-dimensional space. w_0 is initial global model w_0 and w_i is the local model of the i -th client. $w_i - w_0$ is the gradient and we define the *feature of gradient* as $(|x + y|, l, \theta)$.

Manhattan Distance. However, the Euclidean distance suffers from the so-called “curse of dimensionality” [48], which renders distance metrics less sensitive in high dimensional space, especially in the case of Euclidean distance. Theorem 1 uses $\frac{Dmax_d^k - Dmin_d^k}{Dmin_d^k}$ to indicate the “meaningfulness” of a distance referred as the *Relative Contrast*. It demonstrates that the distance metric is meaningless in high-dimensional space because as dimensionality increases, the distance between the nearest and the furthest neighbor approximates to be the same.

Theorem 1 (Beyer et. al.[5] The curse of dimensionality)

$$\text{If } \lim_{d \rightarrow \infty} \text{var} \left(\frac{\|X_d\|_k}{E[\|X_d\|_k]} \right) = 0, \text{ then}$$

$$\frac{Dmax_d^k - Dmin_d^k}{Dmin_d^k} \rightarrow_p 0,$$

where d denotes the dimensionality of the data space, $E[X]$ and $\text{var}[X]$ denotes the expected value and variance of a

random variable X , $Dmax_d^k$ and $Dmin_d^k$ means the farthest/nearest distance of the N points to the origin using the distance metric L_k .

To find out which distance metric is more meaningful in high-dimensional space, we size up the relation between the dimension d and the distance L_k . We have Lemma 1, which shows that $Dmax_d^k - Dmin_d^k$ increases at the ratio of $d^{(1/k)-(1/2)}$.

Lemma 1 (Hinneburg et.al.[22]) Let \mathcal{F} be an arbitrary distribution of two points and the distance function $\|\cdot\|$ be an L_k metric. Then,

$$\lim_{d \rightarrow \infty} E \left[\frac{Dmax_d^k - Dmin_d^k}{d^{(1/k)-(1/2)}} \right] = C_k,$$

where C_k is some constant dependent on k .

Subsequently, we use the value of $Dmax_d^k - Dmin_d^k$ as a criterion to evaluate the ability of a particular metric to discriminate different vectors in high dimensional space.

Proposition 1 Let $M_d = Dmax_d^1 - Dmin_d^1$ reflect the discriminating ability of Manhattan distance and $U_d = Dmax_d^2 - Dmin_d^2$ reflect the discriminating ability of Euclidean distance. Then,

$$\lim_{d \rightarrow \infty} E \left[\frac{M_d}{U_d \cdot d^{\frac{1}{2}}} \right] = C'.$$

where C' is a constant. Proof see in Appendix C.

Proposition 1 shows that M_d and $U_d \cdot d^{\frac{1}{2}}$ are infinite of the same order. Since $d^{\frac{1}{2}}$ is also an infinite value, M_d towers above U_d , suggesting that Manhattan can discriminate more than the Euclidean distance in high-dimension space. Given

this, we conclude that the Manhattan distance is better than the Euclidean metric, withstanding the curse of dimensionality, which chimes in with the previous experiment [1, 37]. The parameters of neural networks can be viewed as a typical high-dimensional space, especially when the number of parameters increases as the trend in deep learning applications [9, 44, 42, 15, 34, 54, 43]. In light of this and Theorem 1, we argue that Euclidean distance is insufficient to discriminate between malicious and benign gradients. With Proposition 1, we introduce the Manhattan distance metric to describe the characteristics of gradients which captures the difference in high dimensional space better.

Multiple Metrics as Feature of Gradient. Another problem that we identify as discussed in Section 1 is that current methods conduct defense on a single metric basis. With only one metric, a sophisticated attacker could effortlessly bypass them with a well-designed gradient. For instance, PGD attack [47] manage to break the defense based on Euclidean distances by clipping the gradient norm to a small value. Moreover, attackers conduct attacks under different environments and data distributions, leading to malicious gradients with diverse characteristics that a single metric cannot handle. To this end, we propose multiple metrics to cooperatively identify the malicious gradient by defining the *feature of gradient* as the angle by the Cosine similarity, length by the Euclidean distances, and its Manhattan norm by the Manhattan distance, namely, $\mathbf{x} = (x_{Man}, x_{Eul}, x_{Cosine})$. Here, x_{Man} denotes the Manhattan distance, x_{Eul} denotes the Euclidean distance and x_{Cosine} denotes the Cosine similarity. Figure 2 provides a schematic diagram of the *gradient feature*. We compute the gradient feature x_{Man} , x_{Eul} , and x_{Cosine} of the i -th client of each round as follows: $x_{Man}^{(i)} = \|\mathbf{w}_i - \mathbf{w}_0\|_1$, $x_{Eul}^{(i)} = \|\mathbf{w}_i - \mathbf{w}_0\|_2$, $x_{Cosine}^{(i)} = \frac{\langle \mathbf{w}_0, \mathbf{w}_i \rangle}{\|\mathbf{w}_0\| \cdot \|\mathbf{w}_i\|}$, where \mathbf{w}_0 denotes the global model before federated training and \mathbf{w}_i denotes the i -th client’s local model after training, $\|\cdot\|_1$ means the L_1 norm of a vector, $\|\cdot\|_2$ means the L_2 norm of a vector, $\langle \cdot, \cdot \rangle$ and represents inner product. After defining the *feature of gradient*, we utilize it for malicious identification. The goal is to identify the outlier among the gradients. We use the sum of the distance between each gradient as the indicator. Formally, we define it as Equation 3.

$$\mathbf{x}'^{(i)} = \left(\sum_{j,j \neq i}^K |x_{Man}^{(i)} - x_{Man}^{(j)}|, \sum_{j,j \neq i}^K |x_{Eul}^{(i)} - x_{Eul}^{(j)}|, \sum_{j,j \neq i}^K |x_{Cosine}^{(i)} - x_{Cosine}^{(j)}| \right), \quad (3)$$

3.3. Dynamic Weighting Through Whitening

After defining the *feature of gradient* and indicator, we now provide an approach to score each gradient. Two fac-

tors are taken into account when considering an efficient way to turn the indicator into a unified score that works under various data distributions and attacks.

- Different scales of the three distance metrics are the first obstacle before utilizing these metrics collaboratively. Since each metric is correlated, a novel regularization is required instead of the usual normalization by the maximum value.
- Different data distributions (e.g. different degrees of non-IID) render the gradients of both malicious and benign clients different. Thus, dynamic weighting is required to cope with various environments and attacks to achieve a universal defense.

Considering the above factors, we propose first to project the value of the metrics in the *feature of gradient* to its corresponding principal axis by applying a whitening process:

$$\delta^{(i)} = \sqrt{\mathbf{x}'^{(i)\top} \Sigma^{-1} \mathbf{x}'^{(i)}}. \quad (4)$$

where Σ is the the covariance matrix of $\mathbf{X} = [\mathbf{x}'^{(1)} \mathbf{x}'^{(2)} \dots \mathbf{x}'^{(K)}]^\top$. Since we need to calculate the inverse of the covariance matrix Σ , the number of samples must be larger than the number of features, i.e., $K > 3$. Please notice that the inverse of the covariance matrix is calculated based on the selected gradients, which will change the weights of the features dynamically depending on the feature distribution, hence ”dynamic weighting”. With such a dynamic weight, our method can better adapt to different environments and resist various attacks.

3.4. Aggregation of Benign Gradients

After obtaining the score of each gradient, we aggregate the gradients with high scoring. A higher score indicates that the gradient is less divergent among all. After the selection of ”benign” gradients, we can aggregate them with any existing aggregation methods (e.g. Averaging of gradients). We follow [35] and perform standard FedAvg on the ”benign” gradients. A detailed analysis of aggregation of benign gradients is in Appendix A.1.

4. Experiment Setup

We will first introduce the settings and then the objective of our experiments.

Datasets. We conduct our experiments on multiple datasets, including two visual datasets, CIFAR-10 [27] and EMNIST [12]. We also experiment on other datasets to prove the generalization ability, including CINIC10 [13], LOAN [25] and Sentiment140 [18] as follows:

- **CIFAR-10 dataset [27]** is a vision classification dataset containing 50,000 training samples and

Algorithm 1 Our aggregation algorithm

Input: Total number of clients in each round K , models of clients $w^{(1)}, w^{(2)}, \dots, w^{(K)}$, last global model w_0 , fraction selected to aggregation p , global learning rate η , size of the i -th client dataset $n^{(i)}$

Output: Global model w^*

```
1: for  $i \in \{1, 2, \dots, K\}$  do           ▷ compute the gradients
   features
2:    $x_{Man}^{(i)} \leftarrow \|w_0 - w_i\|_1$ 
3:    $x_{Eul}^{(i)} \leftarrow \|w_0 - w_i\|_2$ 
4:    $x_{Cos}^{(i)} \leftarrow \frac{\langle w_0, w_i \rangle}{\|w_0\| \cdot \|w_i\|}$ 
5:    $x^{(i)} \leftarrow (x_{Man}^{(i)}, x_{Eul}^{(i)}, x_{Cos}^{(i)})$ 
6: end for
7: for  $i, j \in \{1, 2, \dots, K\}$  do       ▷ compute the sum of the
   distance between each gradient
8:    $x'_{Man}{}^{(i)} \leftarrow \sum_{j, j \neq i}^K |x_{Man}^{(i)} - x_{Man}^{(j)}|$ 
9:    $x'_{Eul}{}^{(i)} \leftarrow \sum_{j, j \neq i}^K |x_{Eul}^{(i)} - x_{Eul}^{(j)}|$ 
10:   $x'_{Cos}{}^{(i)} \leftarrow \sum_{j, j \neq i}^K |x_{Cos}^{(i)} - x_{Cos}^{(j)}|$ 
11:   $x'^{(i)} \leftarrow (x'_{Man}{}^{(i)}, x'_{Eul}{}^{(i)}, x'_{Cos}{}^{(i)})$ 
12: end for
13:  $X \leftarrow [x'^{(1)}, x'^{(2)}, \dots, x'^{(K)}]^\top$ 
14:  $\Sigma \leftarrow$  the Covariance of  $X$        ▷ compute the adaptive
   weight matrix in this round
15: for  $i \in \{1, 2, \dots, K\}$  do
16:    $\delta^{(i)} \leftarrow \sqrt{x'^{(i)\top} \cdot \Sigma^{-1} \cdot x'^{(i)}}$    ▷ dynamically
   compute the divergency
17: end for
18: Remove the  $K \cdot (1-p)$  models with the high divergency
    $\delta^{(i)}$ , remaining models as  $\mathbb{B}$ 
19:  $w^* \leftarrow w_0 + \eta \cdot \frac{\sum_{i \in \mathbb{B}} n^{(i)} \cdot (w^{(i)} - w_0)}{\sum_{i \in \mathbb{B}} n^{(i)}}$ 
```

10,000 testing samples with ten classes. To simulate the FL environment, we set the a of the Dirichlet distribution to be 0.5, representing a non-IID distribution. Smaller a means a larger degree of non-IID.

- **EMNIST dataset [12]** is a digit classification dataset consisting of 280,000 real-world handwritten images of digits from 0 to 9. We also set the degree of non-IID to be $a = 0.5$.
- **CINIC10 [13]** has a total of 270,000 images, 4.5 times that of CIFAR-10, which is constructed from two different sources: ImageNet and CIFAR-10. We set the non-IID degree as $a = 0.5$.
- **Lending Club Loan dataset (LOAN) [25]** contains financial information such as loan information and credit scores for loan status prediction. There are 2,260,668 samples with a total of 9 classes. To simulate a more realistic situation, we divide the data by US states, where each state represents one client. We

set the non-IID parameter $a = 0.9$.

- **Sentiment140 [18]** is a publicly available English Twitter sentiment classification dataset containing 1.6 million tweets, of which 800,000 are for training, and the other 800,000 are for testing. Each tweet in the dataset is labeled with a sentiment polarity (positive, neutral, or negative).

Models. Different models are used on different datasets to demonstrate the generality of our approach. A 9-layer VGG style network (VGG-9) [45] is trained on the CIFAR-10. On the EMNIST, we train LeNet-5 model [28]. We train the LSTM model [23] for Sentiment140 and a network with three fully-connected layers for LOAN.

Backdoor Attacks. We implement both pixel-pattern and semantic backdoors. To make the attack more challenging, we use the multiple-shot attack instead of the single-shot. The following are the attack methods we use:

- **Model Replacement Attack [4]:** Malicious clients generate backdoor gradients during their local training. Attackers scale up their gradients to ensure that the malicious model will replace the global model. Usually, the scaling factor is N/K ;
- **DBA Attack [50]:** DBA decomposes a global trigger pattern into separate local patterns and embeds them into the training sets of different attackers, resulting in a more negligible difference between benign and backdoor gradients;
- **PGD Attack [47]:** To bypass the Euclidean-based defense, PGD Attack projects the model parameter within the ball centered around the global model;
- **Edge-case PGD Attack [47]:** Attackers have a set of edge-case samples and some benign samples as their poisoned datasets. Edge-case samples are unlikely to exist in the benign training set. Thus, the benign gradients cannot dilute the backdoor. We launch this attack in the form of PGD. The Edge-case PGD attack is so stealthy that most defenses cannot defend against it.

Evaluation. We leverage the following two metrics for evaluating the effectiveness of our defense:

- **Backdoor Accuracy (BA)** indicates whether the attacker succeeds. It shows the accuracy of the model on the backdoor task. The attackers aim to maximize this metric, and the lower the BA, the more efficient the defense.
- **Main Task Accuracy (MA)** is the accuracy on the main task. The overall objective of all participants in an FL system is to maximize MA. The benign clients maximize the model performance, while the attackers seek to be stealthy and maintain MA. Thus, the defense should not cause a significant drop in MA, which we report to demonstrate that our defense does not affect

the functionality of the FL system. We report all BA and MA values in percentages.

Hyperparameters are detailed in Appendix B.

Objectives of Experiments. Firstly, we aim to illustrate the effectiveness of our proposed defense under different attacks by comparing it with the current and previous state-of-the-art defenses. Secondly, we demonstrate that the performance of our defense is consistent under different scenarios (*e.g.*, different degrees of non-IID a and different attack frequencies). Thirdly, we figure out the reason for our defense works through the ablation study in Section 5.4.

5. Experiment Results

5.1. Robustness against Different Attacks and Comparison with SOTA

We first demonstrate the robustness of our approach under different attacks and then compare our methods with the state-of-the-art defense methods such as Krum [8], Multi-Krum [8], Foolsgold [17], RFA [41], Weak-DP [46], and Flame [38] in Table 1. All the defenses except RFA and Weak-DP successfully defended against the model replacement attack since RFA and Weak-DP still incorporate some malicious gradient into the global model. Our defense achieves the lowest BA on both datasets (0.56% and 0.00%) with little impact on the MA. Under the DBA attack, Multi-Krum almost has no effect. Flame and our method also deteriorate because 40% of the participants in each round are malicious, which significantly impacts the method of selecting multiple clients for aggregation. Although Krum and Foolsgold successfully defend against the DBA attack, they also negatively affect the overall performance of the global model. Our defense can reduce the BA to less than 10% without affecting the MA.

PGD attack is much stealthier than the above two which Krum, RFA, and Multi-Krum fail to detect. Moreover, once they select the wrong gradient for aggregation, the global model is replaced, leading to a worse BA than FedAvg. Ours and Flame achieve the best effect of 0.56% and 0.00% on the two datasets, respectively. Edge-case PGD is the hardest of all since it is both stealthy and effective since it attacks with edge-case samples and projects the gradient back to the L_2 norm ball. All the previous defenses have limited effect on this attack. Flame reduces the BA to 5.12%. However, Flame suffers from a 6% drop on MA, which is a result of adding noise. In comparison, ours achieves a much lower BA of 3.06% with almost no impact on MA.

On EMNIST, we also achieve better defense compared to up-to-date methods. Remarkably, we achieve 0.0% BA under both PGD and Edge-case PGD attacks meaning that we defend entirely against these two adversaries. We also achieve similar performance without sacrificing the benign

performance against model replacement and DBA. To provide an overall comparison with previous methods, we follow [52] and design an improved version of the ranking score for each defense with respect to baseline FedAvg. We use the percentage of improvement over the baseline to score and compute it on MA/BA respectively, formally as $\frac{K_{MA/BA} - B_{MA/BA}}{B_{MA/BA}}$ where K denotes some methods and B denotes baseline. Adding up $(score_{MA} - score_{BA})$ across all attacks produces the ranking score for each defense. As shown in the last column of Table 1, we show that our method obtains the highest ranking score with almost 400% better than the baseline and outperforms the second-best Flame by around 0.5. We detail the ranking score in Appendix D.

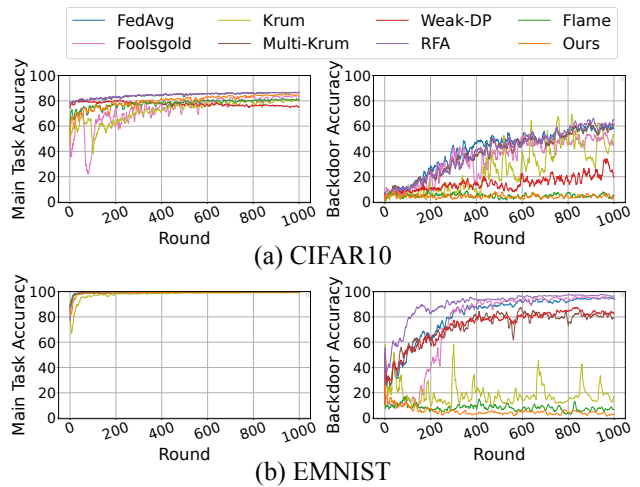


Figure 3: MA(%) and BA(%) of various defenses under Edge-case PGD attack.

We show the MA and BA along the training process of the seven defenses under the Edge-case PGD attack in Figure 3. Only our method and Flame successfully resist the attack during the entire training process, and Flame also dampens the MA. The disadvantage of our approach is that the convergence speed is slightly slowed down compared to FedAvg, but it is still much better than the compared defenses, *e.g.*, Krum and Foolsgold.

5.2. Impact of Training Environment

In this section, we demonstrate that our defense is invariant to the different training environments, *i.e.* degrees of non-IID, and attack frequency.

Impact of Different Degrees of Non-IID. As mentioned abundantly in Section 3.3, the different data distributions contribute to one of the hardest problems in distance-based defense. Thus, to illustrate the effectiveness of our method, we demonstrate that our defense is invariant to the different degrees of non-IID. The experiment is conducted with the

Table 1: Robustness of our approach compared to the SOTA defenses for various challenging attacks.

Dataset	Defense	Model Replacement [4]		DBA [50]		PGD [47]		Edge-case PGD [47]		Ranking Score \uparrow
		MA \uparrow	BA \downarrow	MA \uparrow	BA \downarrow	MA \uparrow	BA \downarrow	MA \uparrow	BA \downarrow	
CIFAR10	FedAvg [35]	86.95	64.80	79.23	90.44	87.04	14.44	87.14	55.10	0
	RFA [41]	86.69(+0.00)	25.56(-0.61)	79.6(+0.00)	57.69(-0.36)	87.1(+0.00)	52.56(+2.64)	86.47(-0.01)	65.31(+0.19)	-1.86
	Foolsgold [17]	85.71(-0.01)	6.67(-0.90)	77.56(-0.02)	3.43(-0.96)	84.92(-0.02)	14.44(+0.00)	85.72(-0.02)	45.41(-0.18)	+1.96
	Krum [8]	82.17(-0.05)	6.11(-0.91)	78.18(-0.01)	6.01(-0.93)	82.32(-0.05)	66.67(+3.62)	81.23(-0.07)	59.18(+0.07)	-2.04
	Multi-Krum [8]	86.55(+0.00)	1.67(-0.97)	79.33(0.00)	91.39(+0.01)	86.52(-0.01)	17.78(+0.23)	87.4(+0.00)	60.2(+0.09)	+0.63
	Weak-DP [46]	74.41(-0.14)	46.11(-0.29)	10.00(-0.87)	0.00(-1.00)	73.61(-0.15)	12.78(-0.11)	73.84(-0.15)	53.06(-0.04)	+0.12
	Flame [38]	80.58(-0.07)	0.56(-0.99)	76.78(-0.03)	37.24(-0.59)	81.24(-0.07)	0.56(-0.96)	81.41(-0.07)	5.12(-0.91)	+3.21
	Ours	86.34(-0.01)	0.56(-0.99)	79.61(+0.00)	9.98(-0.89)	86.44(-0.01)	0.56(-0.96)	86.86(+0.00)	3.06(-0.94)	+3.77
EMNIST	FedAvg [35]	99.54	96.00	97.68	94.13	99.55	10.00	99.37	96.00	0
	RFA [41]	99.57(+0.00)	6.00(-0.94)	97.87(+0.00)	1.39(-0.99)	99.32(+0.00)	4.00(-0.60)	99.29(+0.00)	97.00(+0.01)	+2.51
	Foolsgold [17]	96.42(-0.03)	98.00(+0.02)	97.24(+0.00)	0.64(-0.99)	99.07(+0.00)	94.00(+8.40)	99.13(+0.00)	98.00(+0.02)	-7.49
	Krum [8]	99.22(+0.00)	0.00(-1.00)	97.7(+0.00)	0.56(-0.99)	99.12(+0.00)	1.00(-0.90)	99.14(+0.00)	12.00(-0.88)	+3.76
	Multi-Krum [8]	99.58(+0.00)	0.00(-1.00)	97.85(+0.00)	47.43(-0.50)	99.54(+0.00)	0.00(-1.00)	99.57(+0.00)	84.00(-0.13)	+2.63
	Weak-DP [46]	99.37(+0.00)	86.00(-0.10)	10.00(-0.90)	0.00(-1.00)	99.41(+0.00)	14.00(+0.40)	99.39(+0.00)	89.00(-0.07)	-0.12
	Flame [38]	99.39(+0.00)	0.00(-1.00)	97.12(-0.01)	17.38(-0.82)	99.39(+0.00)	0.00(-1.00)	99.44(+0.00)	13.00(-0.86)	+3.67
	Ours	99.53(+0.00)	0.00(-1.00)	97.39(+0.00)	4.23(-0.96)	99.54(+0.00)	0.00(-1.00)	99.58(+0.00)	0.00(-1.00)	+3.95

Edge-case PGD attack. Figure 4 shows that the BA remains low in all cases. Our method only slightly impacts the MA with $a = 0.2$.

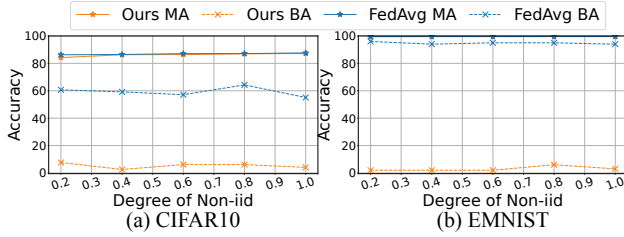


Figure 4: Accuracy(%) of our defense under Edge-case PGD attack on different degrees of non-IID, which is indicated by the a parameter of the Dirichlet distribution.

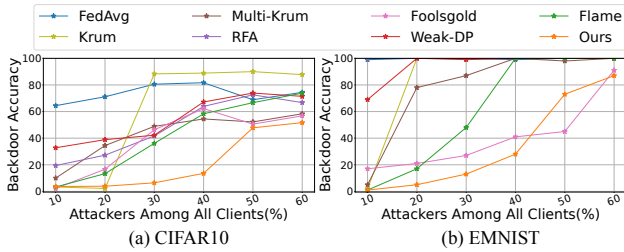


Figure 5: BA(%) of various defenses under the backdoor attack versus percentage of attackers among all clients.

Impact of Attacker Percentage. We show the BA performance with different ratios of attackers under the backdoor attack in Figure 5, where we compare ours with other methods. As a result of our rigorous experiment settings, it usually selects more than $K/2$ adversarial models in each round, so all defenses, especially the clustering-based methods, fail to defend when 50% of clients are compromised.

For EMNIST, BA rises steeply as the simplicity of the task, achieving high accuracy with few parameters injected. Nonetheless, our method performs significantly superior to other methods, and exhibits invariance to the number of attackers encountered.

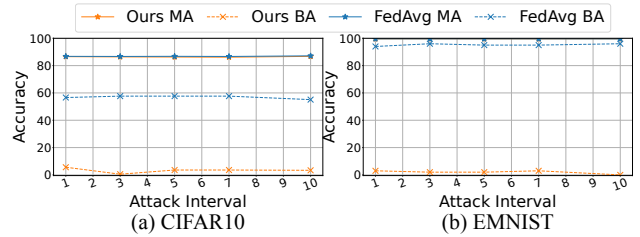


Figure 6: Accuracy(%) of our defense under Edge-case PGD attack versus attack interval.

Impact of Attack Frequency. The frequency of the attack affects the strength and persistence of the attack. We demonstrate that our defense is invariant to the frequency of attack in Figure 6 where we show that our defense achieves consistently low BA under all attack intervals. Since the other defenses cannot resist the edge-case PGD with the minimum frequency(attack interval of 10), we do not plot their results.

5.3. Generalization to Different Datasets

We experiment on the LOAN and Sentiment140, as shown in Table 2. Our method achieves high MA and low BA on both datasets, proving the adaptivity to different models and tasks.

5.4. Ablation Study

The above experiments present the remarkable defense ability of our approach. In this section, we conduct an abla-

Table 2: Effectiveness of our method against backdoor attack on other types of datasets.

Defense	CINIC10		LOAN		Sentiment140	
	MA \uparrow	BA \downarrow	MA \uparrow	BA \downarrow	MA \uparrow	BA \downarrow
FedAvg	80.02	36.22	89.05	61.36	82.59	89.17
Ours	76.24	4.59	88.52	0	81.67	5.83

Table 3: Effectiveness of the metrics in our approach to defending against the various attacks on CIFAR10.

Defenses	Model Replacement	PGD	Edge-case PGD
	MA/BA	MA/BA	MA/BA
	Man	83.86/ 0.56	83.74/25.56
Eul	86.24/ 0.56	85.52/17.78	87.12 /54.08
Cos	84.22/2.22	83.84/30.00	85.38/66.84
Man+Eul	84.09/1.11	84.17/28.63	84.3/67.35
Man+Cosine	85.74/1.67	85.16/23.68	85.86/6.63
Cosine+Eul	86.31/ 0.56	85.44/16.11	85.14/63.78
Man+Cosine+Eul	86.34 / 0.56	86.44 / 0.56	86.86/ 3.06

tion study to understand the role of each component.

Multi-metrics Identifies the Stealthiest Attacks. We first conduct an ablation study on the multi-metrics, as shown in Table 3. We observe the following: 1. All three metrics perform poorly individually. 2. Manhattan distance improves the defense ability when combined with the other metrics. 3. Cosine similarity combined with Manhattan distance provides the best defense against the Edge-case PGD. Given the above observations, we conclude that multi-metrics adaptively defends against universal backdoors. The introduced Manhattan metric contributes to the defense against the stealthiest attack (*e.g.*, Edge-case PGD). To further demonstrate the claim that a single metric cannot identify malicious gradients of diverse characteristics, we show the dominance of each metric under different attacks. We take the metric as the dominant one with the most considerable contribution to this round’s scoring and plot their frequency in Figure 7.

Firstly, we show that Manhattan distance contributes to the defense of PGD and Edge-case PGD, complementing the Cosine similarity justifying the above observations 2 and 3. We also notice that Euclidean is never dominant under PGD and Edge-case PGD, consistent with the analysis that these two attacks bypass L_2 norm through projection back to L_2 norm ball. We also notice that Manhattan distance plays no role under the model replacement attack, consistent with the above experiment that Manhattan does not affect the final BA. It is because model replacement attack scales the gradient of the malicious client, making it susceptible to Cosine and Euclidean distances. Finally, we

conclude that better defense can be achieved if multiple instead of single metrics are utilized since different metrics dominantly detect different malicious gradients.

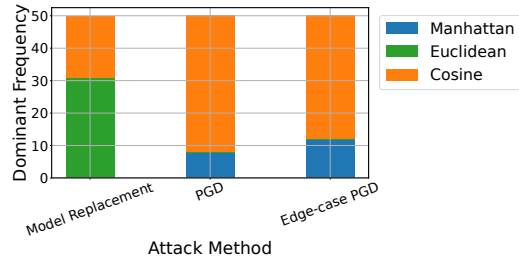


Figure 7: Frequency of dominance of each metric under different attacks on CIFAR10.

Dynamic Weighting through Whitening Process Adaptively Adjust the Scoring. We illustrate the effectiveness of our whitening process discussed in Section 3.3 in Table 4 by comparing it with the standard max normalization. BA increases by a large margin if only max normalization is applied, which illustrates that the whitening process plays an essential role in identifying the malicious under non-IID distributions. Moreover, we notice the MA increases as we apply the whitening process, which shows that it can flexibly identify the benign but different gradients (due to non-IID) and the true malicious ones.

Table 4: Effectiveness of different weighting methods on defense against various attacks on CIFAR10.

Weighting	Model Replacement		PGD		Edge-case PGD	
	MA \uparrow	BA \downarrow	MA \uparrow	BA \downarrow	MA \uparrow	BA \downarrow
Max Norm	83.86	0.56	83.74	25.56	84.08	62.24
Whitening	86.34	0.56	86.44	0.56	86.86	3.06

6. Conclusion

Existing defense methods fail to handle backdoor attacks with malicious updates similar to benign ones, especially when benign clients possess data of non-independent and identical distribution. In this paper, we propose a novel adaptive multi-metrics defense method by leveraging multiple metrics with dynamic weighting to defend against backdoor attacks in FL, withstanding a wide range of stealthy and elaborate attacks. We conduct exhaustive experiments to evaluate our approach proving its effectiveness over a wide range of attack settings.

Acknowledgements

This work is supported by the Key-Area Research and Development Program of Guangdong Province under Grant 2019B010137001.

References

- [1] Charu C Aggarwal, Alexander Hinneburg, and Daniel A Keim. On the surprising behavior of distance metrics in high dimensional space. In *International conference on database theory*, pages 420–434. Springer, 2001.
- [2] Sebastien Andreina, Giorgia Azzurra Marson, Helen Möllering, and Ghassan Karame. Baffle: Backdoor detection via feedback-based federated learning. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*, pages 852–863. IEEE, 2021.
- [3] Sean Augenstein, H Brendan McMahan, Daniel Ramage, Swaroop Ramaswamy, Peter Kairouz, Mingqing Chen, Rajiv Mathews, et al. Generative models for effective ml on private, decentralized datasets. *arXiv preprint arXiv:1911.06679*, 2019.
- [4] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 2938–2948. PMLR, 2020.
- [5] Kevin Beyer, Jonathan Goldstein, Raghu Ramakrishnan, and Uri Shaft. When is “nearest neighbor” meaningful? In *International conference on database theory*, pages 217–235. Springer, 1999.
- [6] Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *International Conference on Machine Learning*, pages 634–643. PMLR, 2019.
- [7] Battista Biggio, Blaine Nelson, and Pavel Laskov. Poisoning attacks against support vector machines. *arXiv preprint arXiv:1206.6389*, 2012.
- [8] Peva Blanchard, El Mahdi El Mhamdi, Rachid Guerraoui, and Julien Stainer. Machine learning with adversaries: Byzantine tolerant gradient descent. *Advances in Neural Information Processing Systems*, 30, 2017.
- [9] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [10] Xiaoyu Cao, Minghong Fang, Jia Liu, and Neil Zhenqiang Gong. Fltrust: Byzantine-robust federated learning via trust bootstrapping. *arXiv preprint arXiv:2012.13995*, 2020.
- [11] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- [12] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, pages 2921–2926. IEEE, 2017.
- [13] Luke N Darlow, Elliot J Crowley, Antreas Antoniou, and Amos J Storkey. Cinc-10 is not imagenet or cifar-10. *arXiv preprint arXiv:1810.03505*, 2018.
- [14] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- [15] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity, 2021.
- [16] Shuhao Fu, Chulin Xie, Bo Li, and Qifeng Chen. Attack-resistant federated learning with residual-based reweighting. *arXiv preprint arXiv:1912.11464*, 2019.
- [17] Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. The limitations of federated learning in sybil settings. In *23rd International Symposium on Research in Attacks, Intrusions and Defenses (RAID 2020)*, pages 301–316, 2020.
- [18] Alec Go, Richa Bhayani, and Lei Huang. Twitter sentiment classification using distant supervision. *CS224N project report, Stanford*, 1(12):2009, 2009.
- [19] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733*, 2017.
- [20] Rachid Guerraoui, Sébastien Rouault, et al. The hidden vulnerability of distributed learning in byzantium. In *International Conference on Machine Learning*, pages 3521–3530. PMLR, 2018.
- [21] Meng Hao, Hongwei Li, Guowen Xu, Sen Liu, and Haomiao Yang. Towards efficient and privacy-preserving federated deep learning. In *ICC 2019-2019 IEEE international conference on communications (ICC)*, pages 1–6. IEEE, 2019.
- [22] Alexander Hinneburg, Charu C Aggarwal, and Daniel A Keim. What is the nearest neighbor in high dimensional spaces? In *26th Internat. Conference on Very Large Databases*, pages 506–515, 2000.
- [23] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [24] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021.
- [25] Wendy Kan. Lending club loan data. <https://www.kaggle.com/wendykan/lending-club-loan-data>. . 2019.
- [26] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- [27] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [28] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [29] Suyi Li, Yong Cheng, Wei Wang, Yang Liu, and Tianjian Chen. Learning to detect malicious clients for robust federated learning. *arXiv preprint arXiv:2002.00211*, 2020.
- [30] Yijiang Li, Wentian Cai, Ying Gao, Chengming Li, and Xiping Hu. More than encoder: Introducing transformer decoder to upsample. In *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 1597–1602. IEEE, 2022.

- [31] Yuezun Li, Yiming Li, Baoyuan Wu, Longkang Li, Ran He, and Siwei Lyu. Invisible backdoor attack with sample-specific triggers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16463–16472, 2021.
- [32] Yijiang Li, Xinjiang Wang, Lihe Yang, Litong Feng, Wayne Zhang, and Ying Gao. Diverse cotraining makes strong semi-supervised segmentor. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023.
- [33] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. 2017.
- [34] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12009–12019, 2022.
- [35] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguerre y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017.
- [36] Brendan McMahan and Daniel Ramage. Federated learning: Collaborative machine learning without centralized training data. <https://ai.googleblog.com/2017/04/federated-learning-collaborative.html>, 2017.
- [37] Evgeny M Mirkes, Jeza Allohbi, and Alexander Gorban. Fractional norms and quasinnorms do not help to overcome the curse of dimensionality. *Entropy*, 22(10):1105, 2020.
- [38] Thien Duc Nguyen, Phillip Rieger, Huili Chen, Hossein Yalame, Helen Möllering, Hossein Fereidooni, Samuel Marchal, Markus Miettinen, Azalia Mirhoseini, Shaza Zeitouni, et al. Flame: Taming backdoors in federated learning. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 1415–1432, 2022.
- [39] Thien Duc Nguyen, Phillip Rieger, Markus Miettinen, and Ahmad-Reza Sadeghi. Poisoning attacks on federated learning-based iot intrusion detection system. In *Proc. Workshop Decentralized IoT Syst. Secur.(DISS)*, pages 1–7, 2020.
- [40] Mustafa Safa Ozdayi, Murat Kantarcioglu, and Yulia R Gel. Defending against backdoors in federated learning with robust learning rate. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 9268–9276, 2021.
- [41] Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning. *IEEE Transactions on Signal Processing*, 70:1142–1154, 2022.
- [42] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [43] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems*, 34:8583–8595, 2021.
- [44] Jing Shao, Siyu Chen, Yangguang Li, Kun Wang, Zhenfei Yin, Yinan He, Jianing Teng, Qinghong Sun, Mengya Gao, Jihao Liu, et al. Intern: A new learning paradigm towards general vision. *arXiv preprint arXiv:2111.08687*, 2021.
- [45] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [46] Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? *arXiv preprint arXiv:1911.07963*, 2019.
- [47] Hongyi Wang, Kartik Sreenivasan, Shashank Rajput, Harit Vishwakarma, Saurabh Agarwal, Jy-yong Sohn, Kangwook Lee, and Dimitris Papailiopoulos. Attack of the tails: Yes, you really can backdoor federated learning. *Advances in Neural Information Processing Systems*, 33:16070–16084, 2020.
- [48] EM Wright. Adaptive control processes: a guided tour. by richard bellman. 1961. 42s. pp. xvi+ 255.(princeton university press). *The Mathematical Gazette*, 46(356):160–161, 1962.
- [49] Chulin Xie, Minghao Chen, Pin-Yu Chen, and Bo Li. Crfl: Certifiably robust federated learning against backdoor attacks. In *International Conference on Machine Learning*, pages 11372–11382. PMLR, 2021.
- [50] Chulin Xie, Keli Huang, Pin-Yu Chen, and Bo Li. Dba: Distributed backdoor attacks against federated learning. In *International Conference on Learning Representations*, 2019.
- [51] Xinyan Xie, Yijiang Li, Ying Gao, Chaojie Wu, Ping Gao, Binjie Song, Wei Wang, and Yiqin Lu. Weakly supervised object localization with soft guidance and channel erasing for auto labelling in autonomous driving systems. *ISA transactions*, 132:39–51, 2023.
- [52] Nanyang Ye, Kaican Li, Haoyue Bai, Runpeng Yu, Lanqing Hong, Fengwei Zhou, Zhenguo Li, and Jun Zhu. Ood-bench: Quantifying and understanding two dimensions of out-of-distribution generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7947–7958, 2022.
- [53] Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. Byzantine-robust distributed learning: Towards optimal statistical rates. In *International Conference on Machine Learning*, pages 5650–5659. PMLR, 2018.
- [54] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- [55] Bo Zhao, Peng Sun, Tao Wang, and Keyu Jiang. Fedinv: Byzantine-robust federated learning by inverting local model updates. In *Thirty-Sixth AAAI Conference on Artificial Intelligence*, pages 9171–9179. AAAI Press, 2022.
- [56] Weiming Zhuang, Xin Gan, Yonggang Wen, Shuai Zhang, and Shuai Yi. Collaborative unsupervised visual representation learning from decentralized data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4912–4921, 2021.