ROBUST FEDERATED LEARNING FOR NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

In federated learning, data is distributed among local clients which collaboratively train a prediction model using secure aggregation. To preserve the privacy of the clients, the federated learning paradigm requires each client to maintain a private local training data set, and only uploads its summarized model updates to the server. In this work, we show that this paradigm could lead to a degenerate model, which collapses in performance when the corrupted updates (under adversarial manipulations) are uploaded and aggregated at the server. To improve model robustness, we first decompose the aggregation error of the central server into bias and variance, and then, propose a robust federated learning framework, named `Fed_BVA`, that performs on-device adversarial training using the bias-variance oriented adversarial examples supplied by the server via asymmetrical communications. The experiments are conducted on multiple benchmark data sets using several prevalent neural network models, and the empirical results show that our framework is robust against white-box and black-box adversarial corruptions under both IID and non-IID settings.

1 INTRODUCTION

The explosive amount of decentralized user data collected from the ever-growing usage of smart devices, e.g., smartphones, wearable devices, home sensors, etc., has led to a surge of interest in the field of decentralized learning. To protect the privacy-sensitive data of the clients, federated learning (McMahan et al., 2017; Yang et al., 2019) has been proposed. Federated learning only allows a group of clients to train local models using their own data, and then collectively merges the model updates on a central server using secure aggregation (Acar et al., 2018). Due to its high privacy-preserving property, federated learning has attracted much attention in recent years along with the prevalence of efficient light-weight deep models (Howard et al., 2017) and low-cost network communications (Wen et al., 2017; Konečnỳ et al., 2016).

In federated learning, the central server only inspects the secure aggregation of the local models as a whole. Consequently, it is susceptible to clients' corrupted updates (e.g., system failures, adversarial manipulations, etc.). Recently, multiple robust federated learning models (Fang et al., 2019; Pillutla et al., 2019; Portnoy & Hendler, 2020; Mostafa, 2019) have been proposed. However, these works only focus on performing client-level robust training or designing server-level aggregation variants with hyper-parameter tuning, and largely ignore the underlying true cause of federated learning's vulnerability coming from the server's generalization error, as we will show in Section 4.1.

Our work bridges this gap by investigating the error incurred during secure aggregation of federated learning from the perspective of bias-variance decomposition (Pedro, 2000; Valentini & Dietterich, 2004). Specifically, we show that the generalization error of the aggregated model on the central server can be decomposed as the combination of **bias** (triggered by the main prediction of these clients) and **variance** (triggered by the variations among clients' predictions). Next, we propose to perform the local robust training on clients by supplying them with a tiny amount of the bias-variance perturbed examples generated from the central server via asymmetrical communications. The experiments are conducted on neural networks with cross-entropy loss, however, other loss functions are also applicable as long as their gradients w.r.t. bias and variance are tractable to estimate. In this way, any gradient-based adversarial training strategies (Goodfellow et al., 2015; Madry et al., 2018) could be used. Compared with previous work, our major contributions include:

- We provide the exact solution of bias-variance analysis w.r.t. the generalization error which is perfectly suitable for neural network based federated learning. As a comparison, performing

adversarial attacks or training with conventional federated learning methods will only focus on the bias of the central model but ignore the variance.
- We demonstrate that the conventional federated learning framework is vulnerable to the strong attacking methods with increasing communication rounds even if the adversarial training using the locally generated adversarial examples is performed on each client.
- Without violating the clients' privacy, we show that providing a tiny amount of bias-variance perturbed data from the central server to the clients through asymmetrical communication could dramatically improve the robustness of the training model under various settings.

2 PRELIMINARIES

2.1 SETTINGS

In federated learning, there is a central server and K different clients, each with access to a private training set $\mathcal{D}_k = \{(x_i^k, t_i^k)\}_{i=1}^{n_k}$, where x_i^k, t_i^k, and n_k are the features, label, and number of training examples in the kth client $(k = 1, \cdots, K)$. Each data $\mathcal{D}_k$ is exclusively owned by client k and will not be shared with the central server or other clients. In addition, there is a small public training set $\mathcal{D}_s = \{(x_j^s, t_j^s)\}_{j=1}^{n_s}$ with n_s training examples from the server that is shared with clients, where $n_s \ll \sum_{k=1}^{K} n_k$. Note that this will not break the privacy constraints, for example, hospitals (local devices) that contribute to a federated learned medical image diagnosis system could take a few publicly accessible images as additional inputs. The goal of federated learning is to train a global classifier $f(\cdot)$ using knowledge from all the clients such that it generalizes well over test data $\mathcal{D}_{test}$. The notation used in this paper is summarized in the Appendix (see Table 4).

2.2 PROBLEM DEFINITION

Privacy-preserving federated learning trains the machine learning models without directly accessing to the clients' raw data. A group of clients could collaborate in achieving the learning objective under the coordination of a central server that aggregates the local clients' model updates asynchronously (Kairouz et al., 2019). In this paper, we study the adversarial robustness of neural networks[1] in federated learning setting, and we define robust decentralized learning as follows.

Definition 2.1. (Robust Federated Learning)
Input: (1) A set of private training data $\{\mathcal{D}_k\}_{k=1}^{K}$ on K different clients; (2) Tiny amount of training data $\mathcal{D}_s$ on the central server; (3) Learning algorithm $f(\cdot)$ and loss function $L(\cdot, \cdot)$.
Output: A trained model on the central server that is robust against adversarial perturbation.

We would like to point out that our problem definition has the following properties: **Asymmetrical communication:** The asymmetrical communication between each client and server cloud is allowed: the server provides both global model parameters and limited shared data to the clients; while each client only uploads its local model parameters back to the server. **Data distribution:** All training examples on the clients and the server are assumed to follow the same data distribution. However, the experiments show that our proposed algorithm also achieves outstanding performance under the non-IID setting, which could be common among personalized clients in real scenarios. **Shared learning algorithm:** All the clients are assumed to use the identical model $f(\cdot)$, including architectures as well as hyper-parameters (e.g., learning rate, local epochs, local batch size).

2.3 BIAS-VARIANCE TRADE-OFF

Following (Pedro, 2000; Valentini & Dietterich, 2004), we define the optimal prediction, main prediction as well as the bias, variance, and noise for any real-valued loss function $L(\cdot, \cdot)$ as follows:

Definition 2.2. (**Optimal Prediction and Main Prediction**) Given a loss function $L(\cdot, \cdot)$ and a learning algorithm $f(\cdot)$, the optimal prediction y_* and main prediction y_m for an example are defined as:

$$y_*(x) = \arg\min_{y} \mathbb{E}_t[L(y, t)] \quad \text{and} \quad y_m(x) = \arg\min_{y'} \mathbb{E}_{\mathcal{D}}[L(f_{\mathcal{D}}(x), y')] \tag{1}$$

where t and $\mathcal{D}$ are viewed as the random variables to denote the class label and training set, and $f_{\mathcal{D}}$ denotes the model trained on $\mathcal{D}$. In short, the main prediction is the prediction whose average loss

[1] Our theoretical contribution mainly focuses on classification using neural networks with cross-entropy loss and mean squared loss. However, the proposed framework is generic to allow the use of other classification loss functions as well.

relative to all the predictions over data distributions is minimum, e.g., the main prediction for zero-one loss is the mode of predictions. In this work, we show that the main prediction is the average prediction of client models for mean squared (MSE) loss and cross-entropy (CE) loss in Section 4.1.

Definition 2.3. (**Bias, Variance and Noise**) Given a loss function $L(\cdot,\cdot)$ and a learning algorithm $f(\cdot)$, the expected loss $\mathbb{E}_{\mathcal{D},t}[L(f_{\mathcal{D}}(x),t)]$ for an example x can be decomposed[2] into bias, variance and noise as follows:

$$B(x) = L(y_m, y_*) \quad \text{and} \quad V(x) = \mathbb{E}_{\mathcal{D}}[L(f_{\mathcal{D}}(x), y_m)] \quad \text{and} \quad N(x) = \mathbb{E}_t[L(y_*, t)] \tag{2}$$

In short, bias is the loss incurred by the main prediction w.r.t. the optimal prediction, and variance is the average loss incurred by predictions w.r.t. the main prediction. Noise is conventionally assumed to be irreducible and independent to $f(\cdot)$.

Remark. *Our definitions on optimal prediction, main prediction, bias, variance and noise slightly differ from previous ones (Pedro, 2000; Valentini & Dietterich, 2004). For example, conventional optimal prediction was defined as $y_*(x) = \arg\min_y \mathbb{E}_t[L(t,y)]$, and it is equivalent to our definition when loss function is symmetric over its arguments, i.e., $L(y_1,y_2) = L(y_2,y_1)$.*

3 The Proposed Framework

A typical framework (Kairouz et al., 2019) of privacy-preserving federated learning can be summarized as follows: (1) *Client Update:* Each client updates local model parameters w_k by minimizing the empirical loss over its own training set; (2) *Forward Communication:* Each client uploads its model parameter update to the central server; (3) *Server Update:* It synchronously aggregates the received parameters; (4) *Backward Communication:* The global parameters are sent back to the clients. Our framework follows the same paradigm but with substantial modifications as below.

Server Update. The server has two components: The first one uses FedAvg (McMahan et al., 2017) algorithm to aggregate the local models' parameters, i.e., $w_G = \text{Aggregate}(w_1, \cdots, w_K) = \sum_{k=1}^{K} \frac{n_k}{n} w_k$ where $n = \sum_{k=1}^{K} n_k$ and w_k is the model parameters in the kth client. Meanwhile, another component is designed to produce adversarially perturbed examples which could be induced by a poisoning attack algorithm for the usage of robust adversarial training.

It has been well studied (Belkin et al., 2019; Pedro, 2000; Valentini & Dietterich, 2004) that in the classification setting, the generalization error of a learning algorithm on an example is determined by the bias, variance, and irreducible noise as defined in Eq. (2). Similar to the previous work, we also assume a noise-free learning scenario where the class label t is a deterministic function of x (i.e., if x is sampled repeatedly, the same values of its class t will be observed). This motivates us to generate the adversarial examples by attacking the bias and variance induced by clients' models as:

$$\max_{\hat{x} \in \Omega(x)} B(\hat{x}; w_1, \cdots, w_K) + \lambda V(\hat{x}; w_1, \cdots, w_K) \quad \forall (x,t) \in \mathcal{D}_s \tag{3}$$

where $B(\hat{x}; w_1, \cdots, w_K)$ and $V(\hat{x}; w_1, \cdots, w_K)$ could be empirically estimated from a finite number of clients' parameters trained on local training sets $\{\mathcal{D}_1, \mathcal{D}_2, \cdots, \mathcal{D}_K\}$. Here λ is a hyper-parameter to measure the trade-off of bias and variance, and $\Omega(x)$ is the perturbation constraint.

Note that $\mathcal{D}_s$ (on the server) is the candidate subset of all available training examples that would lead to their perturbed counterparts. This is a more feasible setting as compared to generating adversarial examples on clients' devices because the server usually has much powerful computational capacity in real scenarios that allows the usage of flexible poisoning attack algorithms. In this case, both poisoned examples and server model parameters would be sent back to each client (*Backward Communication*), while only clients' local parameters would be uploaded to the server (*Forward Communication*), i.e., the *asymmetrical communication* as discussed in Section 2.2.

Client Update. The robust training of one client's prediction model (i.e., w_k) can be formulated as the following minimization problem.

$$\min_{w_k} \left(\sum_{i=1}^{n_k} L(f_{\mathcal{D}_k}(x_i^k; w_k), t_i^k) + \sum_{j=1}^{n_s} L(f_{\mathcal{D}_k}(\hat{x}_j^s; w_k), t_j^s) \right) \tag{4}$$

where $\hat{x}_j^s \in \Omega(x_j^s)$ is the perturbed examples that is asymmetrically transmitted from the server.

[2] This decomposition is based on the weighted sum of bias, variance, and noise. In general, t is a non-deterministic function (Pedro, 2000) of x when the irreducible noise is considered. Namely, if x is sampled repeatedly, different values of t will be observed.

Remark. *Intuitively, the bias measures the systematic loss of a learning algorithm, and the variance measures the prediction consistency of the learner over different training sets. Therefore, our robust federated learning framework has the following advantages: (i) it encourages the clients to consistently produce the optimal prediction for perturbed examples, thereby leading to a better generalization performance; (ii) local adversarial training on perturbed examples allows to learn a robust local model, and thus a robust global model could be aggregated from clients.*

Theoretically, we could still have another alternative robust federated training strategy:

$$\min_{w_k} \sum_{i=1}^{n_k} \max_{\hat{x}_i^k \in \Omega(x_i^k)} L(f(\hat{x}_i^k; w_k), t_i^k) \quad \forall k \in \{1, 2, \cdots, K\} \tag{5}$$

where the perturbed training examples of each client k is generated on local devices from $\mathcal{D}_k$ instead of transmitted from the server. This min-max formula is similar to (Madry et al., 2018; Tramèr et al., 2018) where the inner maximization problem synthesizes the adversarial counterparts of clean examples, while the outer minimization problem finds the optimal model parameters over perturbed training examples. Thus, each local robust model is trained individually, nevertheless, poisoning attacks on device will largely increase the computational cost and memory usage. Meanwhile, it only considers the client-specific loss and is still vulnerable against adversarial examples with increasing communication rounds. Both phenomena are observed in our experiments (see Fig. 4 and Fig. 5).

4 Algorithm

4.1 Bias-Variance Attack

We first consider the maximization problem in Eq. (3) using bias-variance based adversarial attacks. It aims to find the adversarial example $\hat{x}$ (from the original example x) that would produce large bias and variance values w.r.t. clients' local models. Specifically, perturbation constraint $\hat{x} \in \Omega(x)$ forces the adversarial example $\hat{x}$ to be visually indistinguishable w.r.t. x. Here we consider the well-studied l_∞-bounded adversaries[3] (Goodfellow et al., 2015; Madry et al., 2018; Tramèr et al., 2018) such that $\Omega(x) := \{\hat{x} \big| ||\hat{x} - x||_\infty \leq \epsilon\}$ for a perturbation magnitude ϵ. Furthermore, we propose to consider the following two gradient-based algorithms to generate adversarial examples.

Bias-variance based Fast Gradient Sign Method (BV-FGSM): Following FGSM (Goodfellow et al., 2015), it linearizes the maximization problem in Eq. (3) with one-step attack as follows.

$$\hat{x}_{\text{BV-FGSM}} := x + \epsilon \cdot \text{sign}\left(\nabla_x \left(B(x; w_1, \cdots, w_K) + \lambda V(x; w_1, \cdots, w_K)\right)\right) \tag{6}$$

Bias-variance based Projected Gradient Descent (BV-PGD): PGD can be considered as a multi-step variant of FGSM (Kurakin et al., 2017) and might generate powerful adversarial examples. This motivated us to derive a BV-based PGD attack:

$$\hat{x}_{\text{BV-PGD}}^{l+1} := \text{Proj}_{\Omega(x)}\left(\hat{x}^l + \epsilon \cdot \text{sign}\left(\nabla_{\hat{x}^l}\left(B(\hat{x}^l; w_1, \cdots, w_K) + \lambda V(\hat{x}^l; w_1, \cdots, w_K)\right)\right)\right) \tag{7}$$

where $\hat{x}^l$ is the adversarial example at the lth step with the initialization $\hat{x}^0 = x$ and $\text{Proj}_{\Omega(x)}(\cdot)$ projects each step onto $\Omega(x)$.

Remark. *The proposed framework could be naturally generalized to any gradient-based adversarial attack algorithms where the gradients of bias $B(\cdot)$ and variance $V(\cdot)$ w.r.t. x are tractable when estimated from finite training sets. Compared with the existing attack methods (Carlini & Wagner 2017; Goodfellow et al., 2015; Kurakin et al., 2017; Moosavi-Dezfooli et al., 2016), our loss function the adversary aims to optimize is a linear combination of bias and variance, whereas existing work mainly focused on attacking the overall classification error that considers bias only.*

The following theorem states that bias $B(\cdot)$ and variance $V(\cdot)$ as well as their gradients over input x could be estimated using the clients' models.

Theorem 4.1. Assume that $L(\cdot, \cdot)$ is the cross-entropy loss function, then, the empirical estimated main prediction y_m for an input example (x, t) has the following closed-form expression:

$$y_m(x; w_1, \cdots, w_K) = \frac{1}{K} \sum_{k=1}^{K} f_{\mathcal{D}_k}(x; w_k)$$

[3] l_∞ robustness is surely not the only option for robustness learning. However, we use this standard approach to show the limitations of prior federated learning, and evaluate the improvements of our proposed framework.

Furthermore, the empirical bias and variance, as well as their gradients over an input x are estimated as follows:

$$B(x; w_1, \cdots, w_K) = \frac{1}{K}\sum_{k=1}^{K} L(f_{\mathcal{D}_k}(x; w_k), t); \qquad V(x; w_1, \cdots, w_K) = L(y_m, y_m) = H(y_m)$$

Here, $H(y_m) = -\sum_{j=1}^{C} y_m^{(j)} \log y_m^{(j)}$ is the entropy of the main prediction y_m and C is the number of classes. Easily, we can have their gradients in terms of the bias and variance as $\nabla_x B(x; w_1, \cdots, w_K) = \frac{1}{K}\sum_{k=1}^{K} \nabla_x L(f_{\mathcal{D}_k}(x; w_k), t)$ and $\nabla_x V(x; w_1, \cdots, w_K) = -\frac{1}{K}\sum_{k=1}^{K}\sum_{j=1}^{C}(\log y_m^{(j)} + 1)\nabla_x f_{\mathcal{D}_k}^{(j)}(x; w_k)$. Details of the proof is elaborated in A.2.

In addition, we also consider the case where $L(\cdot, \cdot)$ is the MSE loss function. Although its main prediction y_m for an input example (x, t) has a closed-form expression which is exactly the same as the cross-entropy loss, its empirical bias and unbiased variance can only be estimated in the following formulas: $B(x; w_1, \cdots, w_K) = ||\frac{1}{K}\sum_{k=1}^{K} f_{\mathcal{D}_k}(x; w_k) - t||_2^2$ and $V(x; w_1, \cdots, w_K) = \frac{1}{K-1}\sum_{k=1}^{K} ||f_{\mathcal{D}_k}(x; w_k) - \frac{1}{K}\sum_{k=1}^{K} f_{\mathcal{D}_k}(x; w_k)||_2^2$. Note that the gradients of MSE's bias and variance are much more computational demanding comparing with the concise formulas that cross-entropy ends up with. More comparisons are illustrated in Appendix A.5.1.

Algorithm 1 Fed_BVA

1: **Input:** K (number of clients, with local data sets $\{\mathcal{D}_k\}_{k=1}^{K}$); f (learning model), E (number of local epochs); F (fraction of clients selected on each round); B (batch size of local client); η (learning rate); $\mathcal{D}_s$ (shared data set on server); ϵ (perturbation magnitude).
2: **Initialization:** Initialize w_G^0 and $\hat{\mathcal{D}}_s = \emptyset$
3: **for** each round $r = 1, 2, \cdots$ **do**
4: $\quad m = \max(F \cdot K, 1)$
5: $\quad S_r \leftarrow$ randomly sampled m clients
6: $\quad$ **for** each client $k \in S_r$ in parallel **do**
7: $\quad\quad w_k^r \leftarrow$ **ClientUpdate**$(w_G^{r-1}, \hat{\mathcal{D}}_s, k)$
8: $\quad$ **end for**
9: $\quad \hat{\mathcal{D}}_s \leftarrow$ **BVAttack**$(\mathcal{D}_s, \{w_k^r\}|k \in S_r)$
10: $\quad w_G^r \leftarrow$ **Aggregate**$(w_k^r|k \in S_r)$
11: **end for**
12: **return** w_G

Algorithm 2 ClientUpdate$(w, \hat{\mathcal{D}}_s, k)$

Initialize k^{th} client's model with w
$\mathcal{B} \leftarrow$ split $\mathcal{D}_k \cup \hat{\mathcal{D}}_s$ into batches of size B
for each local epoch $i = 1, 2, \cdots, E$ **do**
$\quad$ **for** local batch $(x, t) \in \mathcal{B}$ **do**
$\quad\quad w \leftarrow w - \eta\nabla L(f_{\mathcal{D}_k}(x; w), t)$
$\quad$ **end for**
end for
return w

Algorithm 3 BVAttack$(\mathcal{D}_s, \{w_k^r\}|k \in S_r)$

Initialize $\hat{\mathcal{D}}_s = \emptyset$
for $(x, t) \in \mathcal{D}_s$ **do**
$\quad$ Estimate the gradients $\nabla_x B(x)$ and $\nabla_x V(x)$ using Theorem 4.1
$\quad$ Calculate $\hat{x}$ using Eq. (6) or (7) and add to $\hat{\mathcal{D}}_s$
end for
return $\hat{\mathcal{D}}_s$

4.2 FED_BVA

We present a novel robust federated learning algorithm with our proposed bias-variance attacks, named `Fed_BVA`. Following the framework defined in Eq. (3) and Eq. (4), key components of our algorithm are (1) bias-variance attacks for generating adversarial examples on the server, and (2) adversarial training using poisoned server examples together with clean local examples on each client. Therefore, we optimize these two objectives by producing the adversarial examples $\hat{\mathcal{D}}_s$ and updating the local model parameters w iteratively.

The proposed algorithm is summarized in Alg. 1. Given the server's $\mathcal{D}_s$ and clients' training data $\{\mathcal{D}_k\}_{k=1}^{K}$ as input, the output is a robust global model on the server. First, it initializes the server's model parameter w_G and perturbed data $\hat{\mathcal{D}}_s$, and then assigns to the randomly selected clients (Steps 4-5). Next, each client optimizes its own local model (Steps 6-8) with the received global parameters w_G as well as its own clean data $\mathcal{D}_k$, and uploads the updated parameters back to the server. At last, the server generates the perturbed data $\hat{\mathcal{D}}_s$ (Step 9) using the proposed bias-variance attack algorithm (see Alg. 3) and aggregates the global model parameters (Step 10) using an aggregation method, e.g., FedAvg (McMahan et al., 2017).

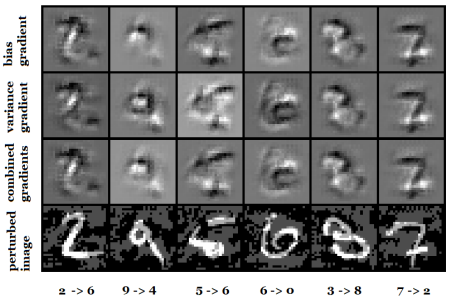

Figure 1: Visualizations of bias, variance, bias+variance, and perturbed images for MNIST.

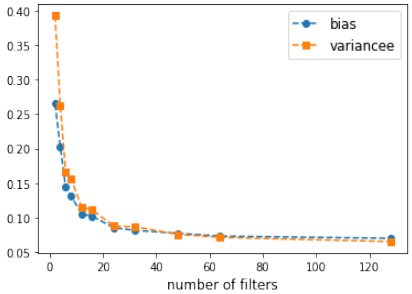

Figure 2: Bias-variance curve w.r.t. the CNN model complexity on MNIST.

5 Experiments

5.1 Settings

In this section, we evaluate the adversarial robustness of our proposed algorithm on four benchmark data sets: MNIST[4], Fashion-MNIST[5], CIFAR-10[6] and CIFAR-100[6]. The baseline models we used include: (1). `FedAvg`: the classical federated averaging model (McMahan et al., 2017). (2). `FedAvg_AT`: The simplified version of our proposed method where the local clients perform adversarial training with the asymmetrical transmitted perturbed data generated on top of `FedAvg`'s aggregation. (3) - (5). `Fed_Bias, Fed_Variance, Fed_BVA`: Our proposed methods where the asymmetrical transmitted perturbed data is generated using the gradients of bias-only attack, variance-only attack, and bias-variance attack, respectively. (6). `EAT`: Ensemble adversarial training (Tramèr et al., 2018), where each client performs local adversarial training using Eq. (5), and their model updates are aggregated on server using `FedAvg`. For fair comparisons, all baselines are modified to the asymmetrical communications setting (`FedAvg` and `EAT` have clean $\mathcal{D}_s$ received), and all their initializations are set to be the same. (7). `EAT+Fed_BVA`: A combination of baselines (5) and (6). Note that baselines (5) and (6) have high computational requirements on client devices, and are usually not preferred in real scenarios.

For the defense model, we use a 4-layer CNN model for MNIST and Fashion-MNIST, and VGG9 architecture for CIFAR-10 and CIFAR-100. Regarding blackbox attacks, we apply ResNet18 (He et al., 2016), VGG11 (Simonyan & Zisserman, 2015), Xception (Chollet, 2017), and MobileNetV2 (Sandler et al., 2018) for CIFAR data, and provide a variety of models for MNIST and Fashion-MNIST by following the design of (Tramèr et al., 2018). The training is performed using the SGD optimizer with fixed learning rate of 0.01 and momentum of value 0.9. All hyper-parameters of federated learning are presented in Table 5 in the Appendix. We empirically demonstrate that these hyper-parameter settings are preferable in terms of both training accuracy and robustness (see the details of Fig. 6 - Fig. 14 in the Appendix). To evaluate the robustness of our federated learning algorithm against adversarial attacks, except for the clean model training, we perform FGSM (Goodfellow et al., 2015), PGD (Kurakin et al., 2017) with 10 and 20 steps towards the aggregated server model on the $\mathcal{D}_{test}$. Following (Tramèr et al., 2018; Wang et al., 2019), the maximum perturbations allowed are $\epsilon = 0.3$ on MNIST and Fashion-MNIST, and $\epsilon = \frac{16}{255}$ on CIFAR-10 and CIFAR-100 for both threat and defense models. For IID sampling, the data is shuffled and uniformly partitioned into each client; For non-IID setting, data is divided into $2F \cdot K$ shards based on sorted labels, then assigns each client with 2 shards. Thereby, each client will have data with at most two classes.

5.2 Result Analysis

To analyze the properties of our proposed `Fed_BVA` framework, in Fig. 1, we visualize the extracted gradients using adversarial attack from bias, variance, and bias-variance. Notice that the gradients of bias and variance are similar but with subtle differences in local pixel areas. However, according to Theorem 4.1, the gradient calculation of these two are quite different: bias requires the target label as input, but variance only needs the model output and main prediction. From another perspective, we also investigate the bias-variance magnitude relationship with varying model complexity. As shown in Fig. 2, with increasing model complexity (more convolutional filters in CNN), both bias and variance decrease. This result is different from the double-descent curve or bell-shape

[4] http://yann.lecun.com/exdb/mnist
[5] https://github.com/zalandoresearch/fashion-mnist
[6] https://www.cs.toronto.edu/~kriz/cifar.html

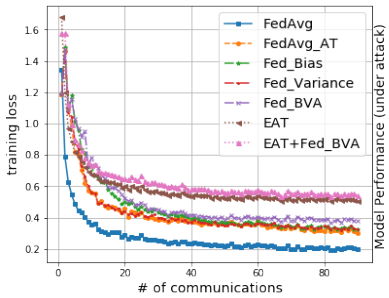

Figure 3: Convergence on Fashion-MNIST

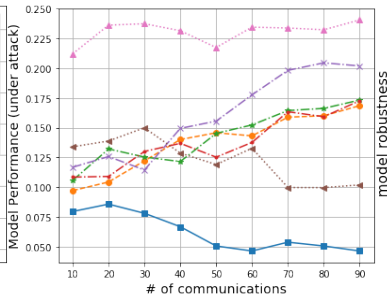

Figure 4: Performance on Fashion-MNIST

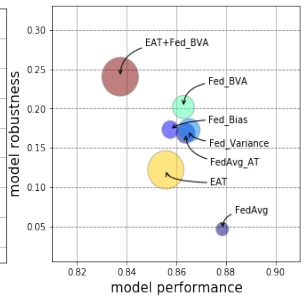

Figure 5: Efficiency on Fashion-MNIST

	IID				non-IID			
	Clean	FGSM	PGD-10	PGD-20	Clean	FGSM	PGD-10	PGD-20
FedAvg	**0.9885**	0.6610	0.5717	0.2513	**0.9776**	0.4910	0.4678	0.1409
FedAvg_AT	0.9876	0.8008	0.7383	0.5009	0.9756	0.7043	0.6423	0.3929
Fed_Bias	0.9855	0.8033	0.7814	0.5794	0.9711	0.6564	0.5750	0.3271
Fed_Variance	0.9843	0.8090	0.7906	0.5826	0.9741	0.6798	0.6067	0.3919
Fed_BVA	0.9859	0.8204	0.8105	0.6253	0.9701	0.7130	0.6747	0.4842
EAT	0.9808	**0.9015**	0.9084	0.8130	0.9737	**0.8068**	0.7281	0.3875
EAT+Fed_BVA	0.9805	0.8991	**0.9116**	**0.8258**	0.9606	0.7947	**0.8194**	**0.6779**

Table 1: Accuracy of MNIST under white-box attacks in IID and non-IID settings

variance curve claimed in (Belkin et al., 2019; Yang et al., 2020). The reasons are twofold: First, their bias-variance definitions are from the MSE regression decomposition perspective, whereas our decomposition utilizes the concept of main prediction, and the generalization error is decomposed from the classification perspective; Second, their implementations only evaluate the bias and variance using training batches on one central model and thus is different from the definition which requires the variance to be estimated from multiple sub-models (in our scenario, client models).

The convergence plot of all baselines is presented in Fig. 3. We observe that `FedAvg` has the best convergence, and all robust training will have a slightly higher loss upon convergence. This matches the observations in (Madry et al., 2018) which state that training performance may be sacrificed in order to provide robustness for small capacity networks. For the model performance shown in Fig. 4, we observe that the aggregation of federated learning is vulnerable to adversarial attacks since both `FedAvg` and `EAT` have decreased performance with an increasing number of server-client communications. Other baselines that utilized the asymmetrical communications have increasing robustness with more communication rounds although only a small number of perturbed examples ($n_s = 64$) are transmitted. We also observe that when communication rounds reach 40, `Fed_BVA` starts to outperform `EAT` while the latter is even more resource-demanding than `Fed_BVA` (shown in Fig. 5, where the pie plot size represents the running time). Overall, bias-variance based adversarial training via asymmetric communication is both effective and efficient for robust federated learning.

For the comprehensive experiments shown in Table 1 and Table 2, it is easy to verify that our proposed model outperforms all other baselines regardless of the source of the perturbed examples (i.e., locally generated like `EAT+Fed_BVA` or asymmetrically transmitted from the server like `Fed_BVA`). Comparing with standard robust federated learning `FedAvg_AT`, the performance of `Fed_BVA` against adversarial attacks still increases $4\% - 13\%$ and $2\% - 9\%$ on IID and non-IID settings respectively, although `Fed_BVA` is theoretically suitable for the cases that clients have IID samples. In Table 3, we observe a similar trend where `Fed_BVA` outperforms `FedAvg_AT` on CIFAR-10 and CIFAR-100 (with $0.2\% - 10\%$ increases) when defending different types of adversarial examples. Comparing with strong local adversarial training baseline `EAT`, we also observe a maximum 13% accuracy increase when applying its bias-variance oriented baseline `EAT+Fed_BVA`.

We also conducted various additional experiments in Appendix A.5 which includes: (1) Comparison of efficiency and effectiveness of `Fed_BVA` using cross-entropy loss and MSE loss; (2) Comparison of single-step `Fed_BVA` and multi-step `Fed_BVA` in terms of the generation of $\hat{\mathcal{D}}_s$; (3) Three training scenarios of `Fed_BVA` that use client-specific adversarial examples or universal adversarial examples; (4) Ablation study in terms of the number of shared perturb examples n_s, optimizer's momentum, and the number of local epochs E; (5) Blackbox attacking transferability between various models on all four data sets under multiple settings.

	IID				non-IID			
	Clean	FGSM	PGD-10	PGD-20	Clean	FGSM	PGD-10	PGD-20
FedAvg	**0.8786**	0.3103	0.0851	0.0463	**0.8157**	0.2093	0.0587	0.0166
FedAvg_AT	0.8638	0.4995	0.198	0.1685	0.7355	0.4955	0.1811	0.1283
Fed_Bias	0.8576	0.5115	0.2161	0.1732	0.719	0.4506	**0.2256**	**0.1745**
Fed_Variance	0.8652	0.4878	0.2095	0.1726	0.6804	0.4365	0.2159	0.1642
Fed_BVA	0.8629	0.5300	0.2431	0.2019	0.6624	0.4925	0.1504	0.0876
EAT	0.8558	**0.8055**	0.1879	0.1019	0.7841	0.5605	0.0774	0.0273
EAT+Fed_BVA	0.8324	0.7149	**0.3668**	**0.2406**	0.7180	**0.6194**	0.1528	0.1022

Table 2: Accuracy of Fashion-MNIST under white-box attacks in IID and non-IID settings

	CIFAR-10				CIFAR-100			
	Clean	FGSM	PGD-10	PGD-20	Clean	FGSM	PGD-10	PGD-20
FedAvg	0.8953	0.2867	0.2274	0.0805	**0.7298**	0.1497	0.1224	0.0422
FedAvg_AT	0.8937	0.3118	0.2914	0.1124	0.7039	0.1663	0.0721	0.0518
Fed_Bias	0.8884	0.3154	0.2892	0.1136	0.7036	0.1601	0.1654	0.0628
Fed_Variance	0.8931	0.3098	0.2772	0.1025	0.7121	0.1672	0.1683	0.0643
Fed_BVA	**0.8964**	0.3254	0.3026	0.1165	0.7074	0.1677	0.1717	0.0648
EAT	0.8331	0.5903	0.6630	0.5530	0.6613	0.2648	0.2033	0.1865
EAT+Fed_BVA	0.8337	**0.5928**	**0.6649**	**0.5617**	0.6588	**0.2708**	**0.3316**	**0.2092**

Table 3: Accuracy of CIFAR-10 and CIFAR-100 under white-box attacks

6 Related Work

Adversarial Machine Learning: While machine learning models have achieved remarkable performance over clean inputs, recent work (Goodfellow et al., 2015; Szegedy et al., 2014) showed that those trained models are vulnerable to adversarially chosen examples by adding the imperceptive noise to the clean inputs. In general, the adversarial robustness of centralized machine learning models have been explored from the following aspects: adversarial attacks (Biggio et al., 2012; Carlini & Wagner, 2017; Guo et al., 2019; Athalye et al., 2018; Zhu et al., 2019), defense (or robust model training) (Madry et al., 2018; Shafahi et al., 2019; Tramèr & Boneh, 2019; Carlini et al., 2019; Tramèr et al., 2018) and interpretable adversarial robustness (Schmidt et al., 2018; Cullina et al., 2018; Tsipras et al., 2018).

Federated Learning: Federated learning with preserved privacy (Konečnỳ et al., 2016; McMahan et al., 2017; Hard et al., 2018) has become prevalent in recent years. Meanwhile, the vulnerability of federated learning to backdoor attacks has also been explored by (Bagdasaryan et al., 2018; Bhagoji et al., 2019; Xie et al., 2019). Following their work, multiple robust federated learning models (Fang et al., 2019; Pillutla et al., 2019; Portnoy & Hendler, 2020; Mostafa, 2019) are also proposed and studied. In this paper, we studied the underlying true cause of federated learning's vulnerability from the perspective of bias-variance analysis. This is in sharp contrast to the existing work that focused on performing client-level model poisoning attacks or designing server-level aggregation variants with hyper-parameter tuning.

Bias-Variance Decomposition: Bias-variance decomposition (Geman et al., 1992) was originally introduced to analyze the generalization error of a learning algorithm. Then, a generalized bias-variance decomposition (Pedro, 2000; Valentini & Dietterich, 2004) was studied in the classification setting which enabled flexible loss functions (e.g., squared loss, zero-one loss). More recently, bias-variance trade-off was experimentally evaluated on modern neural network models (Neal et al., 2018; Belkin et al., 2019; Yang et al., 2020).

7 Conclusion

In this paper, we proposed a novel robust federated learning framework, in which the aggregation incurred loss during the server's aggregation is dissected into a bias part and a variance part. Our approach improves the model robustness through adversarial training by supplying a few bias-variance perturbed samples to the clients via asymmetrical communications. Extensive experiments have been conducted where we evaluated its performance from various aspects on several benchmark data sets. We believe the further exploration of this direction will lead to more findings on the robustness of federated learning.

REFERENCES

Abbas Acar, Hidayet Aksu, A. Selcuk Uluagac, and Mauro Conti. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Comput. Surv.*, 51(4):79:1–79:35, 2018.

Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok. Synthesizing robust adversarial examples. In *International Conference on Machine Learning*, pp. 284–293, 2018.

Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. *arXiv preprint arXiv:1807.00459*, 2018.

Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine-learning practice and the classical bias–variance trade-off. *Proceedings of the National Academy of Sciences*, 116(32):15849–15854, 2019.

Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *International Conference on Machine Learning*, pp. 634–643, 2019.

Battista Biggio, B Nelson, and P Laskov. Poisoning attacks against support vector machines. In *International Conference on Machine Learning*, pp. 1807–1814, 2012.

Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *2017 ieee symposium on security and privacy*, pp. 39–57. IEEE, 2017.

Nicholas Carlini, Anish Athalye, Nicolas Papernot, Wieland Brendel, Jonas Rauber, Dimitris Tsipras, Ian J. Goodfellow, Aleksander Madry, and Alexey Kurakin. On evaluating adversarial robustness. *CoRR*, abs/1902.06705, 2019.

François Chollet. Xception: Deep learning with depthwise separable convolutions. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pp. 1800–1807, 2017.

Daniel Cullina, Arjun Nitin Bhagoji, and Prateek Mittal. PAC-learning in the presence of adversaries. In *Advances in Neural Information Processing Systems*, pp. 230–241, 2018.

Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Zhenqiang Gong. Local model poisoning attacks to byzantine-robust federated learning. *CoRR*, abs/1911.11815, 2019.

Stuart Geman, Elie Bienenstock, and René Doursat. Neural networks and the bias/variance dilemma. *Neural computation*, 4(1):1–58, 1992.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *International Conference on Learning Representations*, 2015.

Chuan Guo, Jacob Gardner, Yurong You, Andrew Gordon Wilson, and Kilian Weinberger. Simple black-box adversarial attacks. In *International Conference on Machine Learning*, pp. 2484–2493, 2019.

Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. *arXiv preprint arXiv:1811.03604*, 2018.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pp. 770–778. IEEE Computer Society, 2016.

Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.

Jakub Konečnỳ, 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.

Alexey Kurakin, Ian Goodfellow, and Samy Bengio. Adversarial machine learning at scale. In *International Conference on Learning Representations*, 2017.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pp. 1273–1282, 2017.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. DeepFool: A simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2574–2582, 2016.

Hesham Mostafa. Robust federated learning through representation matching and adaptive hyper-parameters. *CoRR*, abs/1912.13075, 2019.

Brady Neal, Sarthak Mittal, Aristide Baratin, Vinayak Tantia, Matthew Scicluna, Simon Lacoste-Julien, and Ioannis Mitliagkas. A modern take on the bias-variance tradeoff in neural networks. *arXiv preprint arXiv:1810.08591*, 2018.

Domingos Pedro. A unified bias-variance decomposition and its applications. In *International Conference on Machine Learning*, pp. 231–238, 2000.

Venkata Krishna Pillutla, Sham M. Kakade, and Zaïd Harchaoui. Robust aggregation for federated learning. *CoRR*, abs/1912.13445, 2019.

Amit Portnoy and Danny Hendler. Towards realistic byzantine-robust federated learning. *CoRR*, abs/2004.04986, 2020.

Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, pp. 4510–4520, 2018.

Ludwig Schmidt, Shibani Santurkar, Dimitris Tsipras, Kunal Talwar, and Aleksander Madry. Adversarially robust generalization requires more data. In *Advances in Neural Information Processing Systems*, pp. 5014–5026, 2018.

Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *Advances in Neural Information Processing Systems*, pp. 3353–3364, 2019.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR*, 2015.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna Estrach, Dumitru Erhan, Ian Goodfellow, and Robert Fergus. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014.

Florian Tramèr and Dan Boneh. Adversarial training and robustness for multiple perturbations. In *Advances in Neural Information Processing Systems*, pp. 5858–5868, 2019.

Florian Tramèr, Alexey Kurakin, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick Drew McDaniel. Ensemble adversarial training: Attacks and defenses. In *International Conference on Learning Representations*, 2018.

Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2018.

Giorgio Valentini and Thomas G Dietterich. Bias-variance analysis of support vector machines for the development of svm-based ensemble methods. *Journal of Machine Learning Research*, pp. 725–775, 2004.

Yisen Wang, Xingjun Ma, James Bailey, Jinfeng Yi, Bowen Zhou, and Quanquan Gu. On the convergence and robustness of adversarial training. In *Proceedings of the 36th International Conference on Machine Learning,*, volume 97, pp. 6586–6595, 2019.

Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in Neural Information Processing Systems*, pp. 1509–1519, 2017.

Chulin Xie, Keli Huang, Pin-Yu Chen, and Bo Li. DBA: Distributed backdoor attacks against federated learning. In *International Conference on Learning Representations*, 2019.

Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Trans. Intell. Syst. Technol.*, 10(2):12:1–12:19, 2019.

Zitong Yang, Yaodong Yu, Chong You, Jacob Steinhardt, and Yi Ma. Rethinking bias-variance trade-off for generalization of neural networks. *arXiv preprint arXiv:2002.11328*, 2020.

Chen Zhu, W Ronny Huang, Hengduo Li, Gavin Taylor, Christoph Studer, and Tom Goldstein. Transferable clean-label poisoning attacks on deep neural nets. In *International Conference on Machine Learning*, pp. 7614–7623, 2019.