

# Security Analysis of SplitFed Learning

Momin Ahmad Khan  
University of Massachusetts, Amherst  
makhan@umass.edu

Amir Houmansadr  
University of Massachusetts, Amherst  
amir@cs.umass.edu

Virat Shejwalkar  
University of Massachusetts, Amherst  
vshejwalkar@cs.umass.edu

Fatima M. Anwar  
University of Massachusetts, Amherst  
fanwar@umass.edu

## ABSTRACT

Split Learning (SL) and Federated Learning (FL) are two prominent distributed collaborative learning techniques that maintain data privacy by allowing clients to never share their private data with other clients and servers, and find extensive IoT applications in smart healthcare, smart cities and smart industry. Prior work has extensively explored the security vulnerabilities of FL in the form of poisoning attacks. To mitigate the effect of these attacks, several defenses have also been proposed. Recently, a hybrid of both learning techniques has emerged (commonly known as *SplitFed*) that capitalizes on their advantages (fast training) and eliminates their intrinsic disadvantages (centralized model updates).

In this paper, we perform the first empirical analysis of *SplitFed*'s robustness to strong model *poisoning attacks*. We observe that the model updates in *SplitFed* have significantly smaller dimensionality as compared to FL that is known to have curse of dimensionality. We show that large models that have higher dimensionality are more susceptible to privacy and security attacks, whereas the clients in *SplitFed* do not have the complete model and have lower dimensionality, making them more robust to existing model poisoning attacks. Our results show that the accuracy reduction due to the model poisoning attack is 5x lower for *SplitFed* compared to FL.

## ACM Reference Format:

Momin Ahmad Khan, Virat Shejwalkar, Amir Houmansadr, and Fatima M. Anwar. 2022. Security Analysis of SplitFed Learning. In *Workshop on Challenges in Artificial Intelligence and Machine Learning for Internet of Things (AIChallengesIoT) (SenSys '22)*, November

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*SenSys '22, November 6–9, 2022, Boston, MA, USA*  
© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9886-2/22/11...\$15.00  
<https://doi.org/10.1145/3560905.3568302>

6–9, 2022, Boston, MA, USA. ACM, New York, NY, USA, 7 pages.  
<https://doi.org/10.1145/3560905.3568302>

## 1 INTRODUCTION

**Federated Learning (FL)** [9] is a widely used distributed machine learning setting where several clients collaborate to jointly train a machine learning model by only sharing model updates with a central server. **FL** poses two main weaknesses: (1) server has complete access to client model updates and (2) client has to train over the complete model, making it inefficient for resource-constrained devices. On the other hand another paradigm of distributed machine learning is **Split Learning (SL)** [5] where the model is split between the client and server. This reduces the computational burden on the client device, making it more suitable for resource-constrained devices. Also, the client and the server do not have access to each other's model portions, making **SL** inherently better in terms of privacy. The disadvantage of **SL** is that the clients do not perform computation in parallel. One client communicates with the server at any instance, and the rest of the clients have to wait for their turn.

**SplitFed** [14] is a new way of distributed collaborative machine learning that merges the above-mentioned systems. It adopts the advantages of both paradigms while minimizing the drawbacks of these paradigms. Similar to **SL**, the model is split between the client and the server, thereby allowing resource-constrained devices to perform computation. An additional server called the *Fed Server* takes the model updates of every client, computed in parallel, at each training round and aggregates them to tune the global client portion of the model, just like in **FL**. These properties of **SplitFed Learning** make it suitable for edge deployments.

There is significant research on the security of **FL** exploring model *poisoning attacks* which perturb the model gradients on malicious devices to reduce accuracy [1, 4, 11, 12]. The higher the dimensionality of a model makes the model susceptible to an attack because the adversary has more space to perturb the model in a stealthy manner. Authors in [3] have explored this curse of dimensionality in **FL** with poisoning attacks. However, there has been limited research on *Split learning security*. Authors in [10] used inference

attacks [6] to recover private training data through either the client or server. SplitFed has the ability to vary the portion of the model on the client device which may affect the success of existing attacks. Due to the limitations imposed by dimensionality on the efficacy of model poisoning attacks, an adversary would have a smaller attack space in SplitFed as compared to Federated Learning.

Based on this relation between *model dimensionality* and *attack success*, we hypothesize that the existing *attacks on the robust aggregation mechanisms would be less effective in a SplitFed*. In this work, we present the first, detailed security analysis of SplitFed Learning. To test the validity of our hypothesis we conduct an empirical study over a large pool of clients with various combinations of datasets, model architectures, model split layers and attack-defense pairs.

From our experimental results, we can see that the point at which we split the model has a significant effect on the success of the attack. Decreasing the client model portion hugely reduces the accuracy drop caused by the attack. We also observe the effect of varying the percentage of malicious clients and see that we start to get a significant accuracy drop at 10% malicious clients. We evaluate both IID(Independent and Identically Distributed) and non-IID datasets and find that the attacks are stronger on the IID data. We conclude by providing future directions on developing stronger attacks for SplitFed Learning.

Our contributions can be summarized as follows:

- (1) We adapt and launch the existing state-of-the-art model poisoning attacks in Federated Learning to SplitFed Learning. We implement the optimization-based attack [11] and observe the effects of these attacks when there is (1) no defense mechanism; (2) trimmed mean aggregation rule; and (3) median aggregation rule.
- (2) We explore the effect of the model split location to answer the following questions: Does varying the layer at which the model is split between the client and the server have an effect on the attack success?
- (3) Based on our analysis, we provide recommendations for the possibility of a stronger attack on SplitFed by exploiting its intermediate outputs.

## 2 BACKGROUND

### 2.1 Federated Learning

Federated Learning is a decentralized machine learning mechanism where each client has the same copy of a neural network and its own training data. In each round of training, each client computes forward and backward passes over its model to calculate updated weights. It then sends the updated weights to a central server where all other clients pool their updated weights as well. The central server then

performs an aggregation algorithm, such as FedAvg [9], and computes a global update. This global update is nothing but the average of all the clients' individual weights. This global update is then shared with each of the clients who start the next round of training with the same weights. This is performed continuously until the model converges. Federated learning thus achieves model training without sharing local training data with other clients or the server involved in the training process.

People have exploited the security vulnerabilities of Federated Learning and shown that it is susceptible to model poisoning attacks that severely degrade the accuracy of the model [4, 11]. The LIE(Little is Enough) Attack [1] takes the mean  $\mu$  and the standard deviation  $\sigma$  of the updates and crafts a malicious update  $k = \mu + z * \sigma$  where the coefficient  $z$  controls the amount of deviation. Much stronger than these attacks is the dynamic optimization-based attack proposed by [11] and we use this attack in our study. This attack is described in detail in Section 3.2.


### 2.2 Split Learning

Split Learning [5] is another distributed machine learning framework. It works by splitting the machine learning models between a client and a server. The model is split at some intermediate layer which we will refer to as the **cut layer** throughout this paper. The output of this cut layer is referred to as **smashed data**. The initial layers are in the client, while the rest of the layers are on the server. Therefore, the server performs most of the computation but since it has access to only the second part of the model, it does not have access to the data.

The client performs a forward pass over its portion of the model and sends the smashed data to the server. The server continues the forward propagation over its portion of the model, calculates loss, and begins back-propagation up to the cut layer. It then sends gradients, calculated up to the cut layer, to the client to complete back-propagation. This constitutes one training round of SL. It then sends the updated weights to the next client. The next client starts training from that updated model using its own dataset. The training process is sequential and circular meaning that once we reach the last client, the next client is the first client again.

### 2.3 SplitFed Learning

SplitFed Learning is a new variant [14] of the distributed machine learning framework which combines both Split Learning and Federated Learning. Although Split Learning enhances model privacy by splitting up the model and the server does not know about the client's side of the model or the data, it is sequential in nature and takes more time to train. Compared to Split Learning, Federated Learning has lower



**Figure 1: System Model: This figure shows Federated Learning, Split Learning and how they combine to form SplitFed Learning. It should be noted that (1) The model portion on client is different for Split Learning and SplitFed, (2) There is no transfer of model updates between clients in SplitFed.**

communication efficiency because Split Learning has lesser bandwidth requirements as we are not sharing the complete model at any point. SplitFed combines the parallel computation from Federated Learning and the resource-efficiency and improved model privacy from Split Learning. The working of the SplitFed model is shown in Figure 1. Similar to the classic Split or Federated Learning settings, we have a set of  $N$  clients and  $N$  models here. Each of the  $N$  models is split at the cut layer between the client and the server. The system model is similar to the Split Learning model except that now we have a FedAvg server on the clients’ side, just like in the Federated Learning system model.

Each client computes its forward pass in parallel and sends the intermediate outputs from the *cut layer* to the server. The server computes the forward and backward passes on each of the clients’ outputs and sends them back to the clients. The clients perform the backward pass on their models and send their parameters to the Fed Server. The Fed Server performs Federated Averaging and sends back the parameters to the client.

There is a variant of the SplitFed technique where the server-side computation is done in parallel on each of the client outputs. It completes in lesser time but has reduced accuracy overall. Similar to the client side, it requires an aggregation server on the server side which computes an aggregated server-side model at the end of each global epoch. However, we do not cover that variant in this work.

### 3 POISONING ATTACKS ON SPLITFED

In this section, we first explain our threat model followed by our adaptations of poisoning attacks and defenses for SplitFed. We formulate optimization based model poisoning attacks under trimmed mean and median model aggregation rule for SplitFed training.

#### 3.1 Threat Model

In our threat model, the adversary attempts to compromise model training by injecting malicious updates into the system. The adversary has access to the model on the compromised client. It can modify the weights and gradients of that model. The adversary controls  $M$  out of the total  $N$  clients. In our work, we set  $M/N$  to be 0.2 i.e., 20% of the clients are malicious, which is practical in current deployments. In addition to this, we assume an honest server. Our assumptions are similar to previous works that have done client-side attacks [10, 11], where the server is kept honest (from service providers such as Google, Amazon) and provides legitimate service. We also assume that low dimensional model updates are possible. In a practical setting, this might be the case when multiple hospitals are collaborating to jointly train some cancer detection model and some of those are malicious and wish to corrupt the training process.

The focus of this work is on model poisoning attacks on compromised clients, and not on data inference attacks on compromised servers [4]. We consider the case where the adversary has knowledge of the model gradients on benign clients, but it can only modify the model gradients on malicious clients. This scenario serves as the upper bound for the efficacy of our attacks and the robustness of our aggregation rules, trimmed mean and median, which we describe in detail in Section 3.2.

#### 3.2 Optimization-based attack

Here we describe the state-of-the-art model poisoning attack which we use in our work. In this type of attack by [11], we formulate an optimization problem given by equation 1 that crafts a malicious update which is then sent by the malicious clients to the Fed Server for aggregation. The attacker takes the average of the benign model updates (equation 2) and

perturbs them with a value  $\gamma * \nabla^p$  as shown in equation 3 and this is the highest value of  $\gamma$  that bypasses the aggregation rule being used. The optimization problem is formulated as follows:

$$\arg \max_{\gamma} = \nabla^b - f_{agr}(\nabla_{i \in [m]}^m \cup \nabla_{i \in [m+1, n]}) \quad (1)$$

Here the average of benign gradients is given by:

$$\nabla^b = f_{avg}(\nabla_{i \in [n]}) \quad (2)$$

and each malicious update is given by:

$$\nabla_{i \in [m]}^m = \nabla^b + \gamma * \nabla^p \quad (3)$$

There are different perturbation types, such as unit vector, standard deviation, and opposite direction. In our case, we will be using the standard deviation perturbation as it enables a stronger attack.

$$\nabla_{std}^p = -std(\nabla_{i \in [n]}) \quad (4)$$

This attack is designed to break through state-of-the-art robust aggregation mechanisms such as trimmed mean and median. Next, we describe those aggregation mechanisms briefly and how the optimization framework is tailored to break through trimmed mean and median.

**Trimmed Mean** is an aggregation rule [16], [15] that replaces the simple averaging technique. It takes the set of all updates, including malicious ones, as input, sorts the updates along each dimension  $i$ , and removes  $m$  of the largest and smallest values from the set. Note that, in SplitFed, these updates are from the portion of the model on the clients only. Here,  $m$  is the number of malicious clients. Then it takes the mean of the updates along each trimmed dimension. The aggregated update is then sent back to each client. The general optimization framework given in equation 1 is modified for trimmed mean where  $f_{agr}$  is now  $f_{trmean}$  as done by [11]. Here, the trimmed mean is computed on the aggregate of malicious and benign updates, and the adversary’s objective is to maximize the norm between this trimmed mean and the mean of benign updates.

**Median** is similar to trimmed mean [16], [15] in the sense that it also removes malicious update components along each dimension of the update by means of a statistical operation. It computes the element-wise median of the updates along each dimension. The general optimization framework given in equation 1 is modified in a similar fashion to the trimmed mean attack. The only difference it has is that  $f_{agr}$  is now  $f_{median}$ .

## 4 IMPLEMENTATION

We implemented our adapted poisoning attacks for SplitFed on a variety of image datasets and models, and evaluated over hundreds of clients.

### 4.1 Datasets

**CIFAR10** [7] consists of 60000 images of size 32x32 colored images. There are 10 classes of common things like airplanes, birds, cats, dogs, etc. Among these 60000 images, 50000 are training images while the other 10000 are test images. We chose this dataset because it is commonly used in such works and would serve as a benchmark for easier comparison.

**FEMNIST** [2] consists of 805263 grayscale images distributed over 3550 users. We can partition them in an IID or non-IID fashion. Since we are using CIFAR10 in an IID fashion, we split up FEMNIST among 3500 users in a non-IID fashion. Each grayscale image is 28 x 28 pixels and is one of 62 classes; 26 classes each for upper and lower case letters and 10 classes for digits. The mean number of samples per user is 226.83 and the standard deviation is 88.94. We test FEMNIST in a cross-device setting, where we are choosing 100 random clients out of 3550 every round for training in that particular round.

### 4.2 Models

**Alexnet** [8] is a standard model with some convolutional layers followed by a fully connected layer. We train it on the CIFAR10 dataset for 1000 epochs with a learning rate of 0.01, batch size of 64, and an SGD optimizer. We split the model at three different layers. More specifically we split it at the first, second, and third maxpool layers. We call these split versions V1, V2, and V3 respectively.

**VGG11** [13] has several convolutional layers followed by a fully connected layer. We train it on the CIFAR10 dataset for 200 epochs with a learning rate of 0.01, batch size of 64, and an SGD optimizer. We split the model at three different layers. More specifically we split it at the second, third, and fourth maxpool layers. We call these split versions V1, V2, and V3 respectively.

We design a **Custom** CNN for our FEMNIST dataset. The model is composed of two convolutional layers followed by two fully connected layers. The model for FEMNIST is small and simple because the dataset consists of 28x28 grayscale images and we do not need a very large model to fit on this dataset. We use the image dataset for a particular client as a complete batch since it is small in size. This is a small model so the model split at layers 1,2,3 is called V1,V2, and V3 respectively.

### 4.3 Reproducibility

To ensure the reproducibility of our results we fix the random seeds to be 42 for PyTorch, NumPy, and Python’s random package. This is because there are several parts of the code where randomness is taking place, such as the dataloader, weight initializations etc. Reproducibility is even more important in our non-IID case of the FEMNIST dataset where we

want the same client selection sequence for each experiment trial we conduct.

#### 4.4 Metrics

We report the efficacy of the attack in terms of accuracy drop ( $Acc_{drop}$ ) that quantifies the difference between the test accuracy obtained without attack ( $Acc$ ) and the test accuracy obtained with attack ( $Acc_{attack}$ )

$$Acc_{drop} = Acc - Acc_{attack} \quad (5)$$

Simply reporting the accuracy after the attack would not be a fair comparison because we get different accuracies for different models and model splits.

#### 4.5 Experimental Setup

To test the case of a practical setting, we evaluate our setup with 100 clients. With our setting of 20% malicious clients, 20 of the clients from among our 100 are malicious. We use several NVIDIA RTX8000 GPUS, each having 48GB memory on our university-provided GPU cluster to run our experiments. All the code was written in PyTorch(1.12.1+cu116) and run on Jupyter Notebooks on the GPU cluster. We perform hyperparameter tuning with lr=[0.01, 0.05, 0.1], momentum = [0.9, 0.95, 0.99], and batch\_size = [32, 64, 128].

### 5 ANALYSIS

In this section we will discuss the effect of poison attack in the SplitFed setting on two representative datasets, three models and model splits and two attack-defense pairs.


#### 5.1 Cut Layer and Attack effectiveness

We have seen that the choice of different cut layers does not functionally change the split learning system. However, the choice of cut layer plays a huge role in the effectiveness of our attack. We see that when the client has a smaller portion of the full model the attack does not reduce the overall accuracy much. Conversely, when the client has a larger portion of the model, the adversary has a lot of space to manipulate the model. In this scenario, the overall accuracy of the final model is severely degraded.

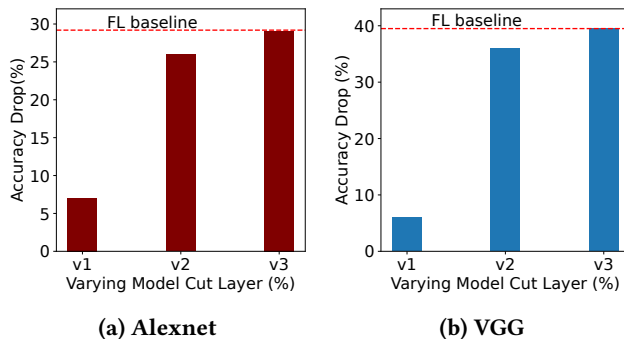
As shown in Table 2, in the case of Alexnet trained in CIFAR10, when we increase the percentage of the model on the client side from VGG v1 to VGG v3, and the accuracy drop increases from 9.05% to 41.42% for trimmed mean. When we compare it with the federated learning case in the same setting and it gives a higher accuracy degradation of (49.3%) as shown in Table 1. because it has the complete model on the client side. Figures 2 and 3 show this accuracy as the *FL baseline* (shown as red dashed line). We also observe that the v3 model split has almost equal accuracy drop as the baseline (FL) however the accuracy drop of v2 split is lower

Dataset Model	TrMean		Median	
	Acc	Acc <sub>drop</sub>	Acc	Acc <sub>drop</sub>
CIFAR10 VGG	73	47	71.5	39.5
CIFAR10 Alexnet	62.4	49.3	62.1	29.19
FEMNIST Custom	81	18	75	8

**Table 1: Attacks in the Federated Learning Setting** than v3 but still significantly higher than that of v1. Based on these results, we can advocate the use of v1 versions of our models for more robustness against poisoning attacks. This can also prove to be beneficial for resource-constrained devices and could lead to faster training time.



**Figure 2: Impact of attacks tailored for trimmed mean defense with varying cut layers on Alexnet and VGG models.**



**Figure 3: Impact of attacks tailored for median defense with varying cut layers on Alexnet and VGG models.**

#### 5.2 Impact of Defense

Next, we compare the two defenses, the trimmed mean and median in our evaluation as shown in Tables 1 and Table 2. We can see throughout all settings with different model architectures, datasets, and model splits median defense performs better than the trimmed mean defense. The highest performance difference is for Alexnet (v3), where the median is

about 11% better than trimmed mean. For the Federated case, we see the same trend with median with 8 – 10% better performance than the trimmed mean. This trend is consistent with FEMNIST too for both FL and SplitFed. From this we can conclude that the median defense is more robust than the trimmed mean defense. This may be due to reason that median is a much more robust measure of central tendency than mean in the presence of an outlier that can still exist when the upper and lower  $m$  values are trimmed. Therefore, it is much harder for the adversary to craft an update that bypasses the median defense.

### 5.3 IID vs non-IID

Initially we perform the attack on non-IID FEMNIST with the same learning rate as in IID CIFAR10 ( $lr=0.01$ ). Similarly to our IID case, we do poisoning at every training round. This leads to very low model accuracy (below 10%) because the adversary gets too much attack space in non-IID because the model updates are very different from each other, and the high learning rate allows it to craft a malicious update of large magnitude. This effect is also validated in [4, 11]. We then perform the attack with a lower learning rate of 0.001 and carry out poisoning after several training rounds without poisoning. We report these results in Tables 1 and Table 2. A significant observation in this case is when we use the  $v1$  model split, where we do not get a drop in accuracy with our attack. This is because the model we are using is already small in size, and the adversary gets an even smaller part to attack.

### 5.4 Impact of the Percentage of Malicious Clients

We need to determine what percentage of clients would be suitable to perform significant accuracy degradation. Although it is intuitive that to get the highest degradation, we need a higher percentage of malicious clients, we need to keep practical considerations in mind. It is not possible to have a very large number of malicious clients in a practical setting [12]. We tested our attack with a varying percentage of malicious clients for the CIFAR10 dataset. We see that the accuracy of the model starts to degrade with a tiny percentage of malicious clients (2%) and that the accuracy of the model shows a sharp degradation at 10% malicious clients. For the Alexnet V2 model, we see a complete degradation in model accuracy when we increase the percentage of malicious clients to 30%. In this scenario, the accuracy of the model is about 10% which is equivalent to a random guess for a classification problem of 10 classes, these trends are comparable with Federated Learning as well, as reported in [11]. We observe that the same increase in the percentage of malicious clients in Federated Learning causes a much

Dataset Model	Split	TrMean		Median	
		Acc	Acc <sub>drop</sub>	Acc	Acc <sub>drop</sub>
CIFAR10 VGG	v1	75.68	14.68	75.41	6.81
	v2	71.49	40.93	71.92	36.92
	v3	67	46	65.5	39.5
CIFAR10 Alexnet	v1	63.05	9.05	63.42	7.42
	v2	63.39	29.39	62.34	26.34
	v3	61.71	41.42	62.1	29.1
FEMNIST Custom	v1	87	0	87	0
	v2	87.3	2.3	87.3	1.6
	v3	84	8	84.9	5.2

Table 2: Attacks in the SplitFed setting.

larger change in accuracy drop as compared to SplitFed. Increasing the percentage of malicious clients from 10% to 15% for CIFAR10 with Alexnet and trimmed mean increases the accuracy drop by 11% in Federated Learning while the accuracy drop in SplitFed is 4%, as shown in Figure 4. We see this effect for every 5% increment in malicious clients from 0% to 30%. This observation is also consistent with the FEMNIST dataset.

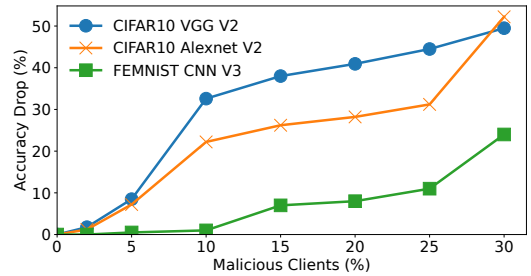


Figure 4: Effect of varying percentage of malicious clients on accuracy drop.

## 6 DISCUSSION AND CONCLUSION

### 6.1 Additional Attack Space

So far we have performed the attack at the client-side model aggregator, to perform a comparison with Federated Learning. It would be interesting to see the effect of performing some perturbation on the *smashed data* sent by some malicious client to the server. The dimensionality of *smashed data* would change with different choices of the *cut layer* and this would most certainly affect the success of the attack.

### 6.2 Conclusion

We performed an extensive analysis of model poisoning on the SplitFed system by using existing state-of-the-art attacks. We test these attacks by pairing them with various defenses, models, datasets, data distributions, and model splits and find that the same attacks in SplitFed have lower efficacy than in Federated Learning because in SplitFed the client does not have the complete model. We also propose an additional attack space that can be explored in future work.

## REFERENCES

- [1] Gilad Baruch, Moran Baruch, and Yoav Goldberg. 2019. A little is enough: Circumventing defenses for distributed learning. *Advances in Neural Information Processing Systems* 32 (2019).
- [2] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2018. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097* (2018).
- [3] Hongyan Chang, Virat Shejwalkar, Reza Shokri, and Amir Houmansadr. 2019. Cronus: Robust and heterogeneous collaborative learning with black-box knowledge transfer. *arXiv preprint arXiv:1912.11279* (2019).
- [4] Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Gong. 2020. Local Model Poisoning Attacks to Byzantine-Robust Federated Learning. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 1605–1622. <https://www.usenix.org/conference/usenixsecurity20/presentation/ffang>
- [5] Otkrist Gupta and Ramesh Raskar. 2018. Distributed learning of deep neural network over multiple agents. *Journal of Network and Computer Applications* 116 (2018), 1–8.
- [6] Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. 2017. Deep models under the GAN: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC conference on computer and communications security*. 603–618.
- [7] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [8] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012).
- [9] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*. PMLR, 1273–1282.
- [10] Dario Pasquini, Giuseppe Ateniese, and Massimo Bernaschi. 2021. Unleashing the tiger: Inference attacks on split learning. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*. 2113–2129.
- [11] Virat Shejwalkar and Amir Houmansadr. 2021. Manipulating the byzantine: Optimizing model poisoning attacks and defenses for federated learning. In *NDSS*.
- [12] Virat Shejwalkar, Amir Houmansadr, Peter Kairouz, and Daniel Ramage. 2022. Back to the drawing board: A critical evaluation of poisoning attacks on production federated learning. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 1354–1371.
- [13] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [14] Chandra Thapa, Pathum Chamikara Mahawaga Arachchige, Seyit Camtepe, and Lichao Sun. 2022. Splitfed: When federated learning meets split learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 8485–8493.
- [15] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. 2018. Generalized byzantine-tolerant sgd. *arXiv preprint arXiv:1802.10116* (2018).
- [16] Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. 2018. Byzantine-robust distributed learning: Towards optimal statistical rates. In *International Conference on Machine Learning*. PMLR, 5650–5659.