

Curse or Redemption? How Data Heterogeneity Affects the Robustness of Federated Learning

Syed Zawad,¹ Ahsan Ali,¹ Pin-Yu Chen,² Ali Anwar,² Yi Zhou,² Nathalie Baracaldo,²
Yuan Tian,³ Feng Yan¹

¹ University of Nevada, Reno

² IBM Research

³ University of Virginia

szawad@nevada.unr.edu, aali@nevada.unr.edu, pin-yu.chen@ibm.com, ali.anwar2@ibm.com, yi.zhou@ibm.com,
baracald@ibm.com, yt2e@virginia.edu

Abstract

Data heterogeneity has been identified as one of the key features in federated learning but often overlooked in the lens of robustness to adversarial attacks. This paper focuses on characterizing and understanding its impact on backdooring attacks in federated learning through comprehensive experiments using synthetic and the LEAF benchmarks. The initial impression driven by our experimental results suggests that data heterogeneity is the dominant factor in the effectiveness of attacks and it may be a redemption for defending against backdooring as it makes the attack less efficient, more challenging to design effective attack strategies, and the attack result also becomes less predictable. However, with further investigations, we found data heterogeneity is more of a curse than a redemption as the attack effectiveness can be significantly boosted by simply adjusting the client-side backdooring timing. More importantly, data heterogeneity may result in overfitting at the local training of benign clients, which can be utilized by attackers to disguise themselves and fool skewed-feature based defenses. In addition, effective attack strategies can be made by adjusting attack data distribution. Finally, we discuss the potential directions of defending the curses brought by data heterogeneity. The results and lessons learned from our extensive experiments and analysis offer new insights for designing robust federated learning methods and systems.

Introduction

Federated Learning (FL) is widely successful in training machine learning (ML) models collaboratively across clients without sharing private data (McMahan et al. 2016; Zhao et al. 2018; Bonawitz et al. 2019). In FL, models are trained locally at clients to preserve data privacy and the trained model weights are sent to a central server for aggregation to update the global model. During the aggregation, privacy mechanisms such as differential privacy (Abadi et al. 2016) and secure aggregation (Bonawitz et al. 2017) are often employed to strengthen the privacy. There are two types of poisoning attacks: performance degradation attacks where the goal of the adversary is to reduce the accuracy/F1 scores

of the model (such as Byzantine attacks) and backdoor attacks aiming at creating targeted misclassifications without affecting the overall performance on the main tasks (Chen et al. 2017; Xie et al. 2019; Bagdasaryan et al. 2018). Defending against such attacks usually requires complete control of the training process or monitoring the training data (Steinhardt, Koh, and Liang 2017), which is challenging in FL due to the privacy requirements. In this paper, we choose the popular and sophisticated backdoor attacks as an example for our study. Although some work exists to defend against backdoor attacks, including activation clustering (Chen et al. 2018) and k-means clustering (Shen, Tople, and Saxena 2016), these approaches require access to the training data making them inapplicable for FL settings. Some attack strategies tailored for FL have also been studied including sybil attacks (Fung, Yoon, and Beschastnikh 2018), model replacement (Bagdasaryan et al. 2018), GANs based attacks (Zhang et al. 2019), and distributed attacks (Xie et al. 2019). However, a comprehensive study on the effectiveness of backdoor attacks under a variety of data distribution among parties remains at unexplored.

The training data in FL is generated by clients and thus heterogeneous inherently (Bonawitz et al. 2019; Chai et al. 2020; Zhao et al. 2018; Sattler et al. 2019). As the training is conducted locally at each client, the data cannot be balanced nor monitored like in conventional data-centralized or distributed ML. Such uncontrollable and severe data heterogeneity is one of the key challenges of FL as it is rarely seen in conventional ML. Despite its uniqueness and importance, data heterogeneity has been largely overlooked through the lens of robustness to backdoor attacks. Existing FL backdoor attacks either assume IID training data distribution among clients or only conduct a simplified study on non-IID data (Bagdasaryan et al. 2018; Bhagoji et al. 2019; Xie et al. 2019). None of them provides a comprehensive study nor understanding on how data heterogeneity impacts the backdoor attacks and defenses.

In this paper, we focus on quantifying and understanding the implications brought by data heterogeneity in FL backdoor attacks through extensive empirical experiments and comprehensive analysis.

We define *Heterogeneity Index* to quantify the extent

of heterogeneity in training data. From our initial investigation driven by both synthetic and the practical LEAF benchmark (Caldas et al. 2018), we surprisingly found that data heterogeneity seems to be a redemption for defending against backdoor attacks. **Redemption 1:** the *attack effectiveness* (usually measured as Attack Success Rate or ASR) reduces sharply when the heterogeneity of training data increases. **Redemption 2:** we found the malicious data distribution is an overlooked important factor when defining an attack strategy given the training data is heterogeneous. A poor selection of malicious data distribution can result in poor attack effectiveness. **Redemption 3:** we further discovered that malicious data distribution plays as a dominant factor in the effectiveness of backdooring. E.g., contrary to the common belief in existing works that higher *attack scale* (defined as the number of compromised clients) and *local attack budget* (defined as the quantity of backdoored data per client) always lead to higher attack effectiveness, our study demonstrates that this is not always the case as malicious data distribution often outperforms the impact of attack scale/budget. This discovery indicates that data heterogeneity makes the design of effective attack strategies more challenging as the attack effectiveness is less correlated to the straightforward attack scale/budget but rather the less intuitive malicious data distribution.

Further investigations, however, reveal that data heterogeneity actually brings curses for the robustness of FL. **Curse 1:** data heterogeneity makes the client-side training very sensitive to the backdoor attack timing. With a proper attack timing, e.g., at the last local batch, the effectiveness of attack can be significantly boosted with only a fraction of attack budget. **Curse 2:** what’s worse is that data heterogeneity makes the most promising skewed-feature based defense strategies such as cosine similarity fall short. Such defending method detects compromised clients by realizing their features are more overfitted than the benign clients. However, with data heterogeneity, benign clients may also have overfitted features that look similar to those of compromised clients. This allows the backdoor attackers to disguise themselves and fool the skewed-feature checking. **Curse 3:** more effective attack strategies can be derived by making the backdoor clients’ data distribution close to the overall data distribution with the help of distribution distance measures such as the Chi-Square statistics. To defend these curses brought by data heterogeneity, we discuss how existing defense mechanisms fit here and the potential directions on data-heterogeneity aware defending strategies.

In summary, our empirical experimental studies show that data heterogeneity appears to be a redemption for the robustness of FL as it makes the attack less effective and more challenging to design good attack strategies. However, our further investigations reveal that data heterogeneity also brings several curses for FL backdooring as it is harder to detect and the attack effectiveness can be significantly boosted by adjusting the local attack timing and malicious data distribution. The defending strategies we propose help alleviate these curses. The results and lessons learned from our thorough experiments and comprehensive analysis offer new insights for designing robust FL methods and systems.

Related Works

Data Heterogeneity in Federated Learning. While data heterogeneity is not new in the ML, the extent of data heterogeneity is much more prevalent in FL compared to data centralized learning (Chai et al. 2020; Li et al. 2020). (Li et al. 2019) theoretically demonstrates the bounds on convergence due to heterogeneity, while (Sattler et al. 2019) providing empirical results on how changing heterogeneity affects model performance. (Li et al. 2020) discusses the challenges of heterogeneity for FL and (Zhao et al. 2018) demonstrates how the clients’ local model weights diverge due to data heterogeneity.

Backdoor Attack. Backdoor attacks for deep learning models are presented in (Chen et al. 2017), where an adversary can insert a pattern in a few training samples from a source class and relabel them to a target class, causing a targeted misclassification. One of the earlier papers (Bagdasaryan et al. 2018) proposes the *model replacement* technique, whereby they eventually replace the global model with a backdoored model stealthily. (Bhagoji et al. 2019) demonstrates that boosting model weights can help attackers and shows that FL is highly susceptible to backdoor attacks. (Fung, Yoon, and Beschastnikh 2018) introduces sybil attacks in the context of FL using label-flipping and backdooring. (Zhang et al. 2019) uses GANs to attack the global model, while (Xie et al. 2019) takes a different approach by focusing on decentralized, colluding attackers, and creating efficient trigger patterns. Our paper takes a different angle by focusing on analyzing the impact of data heterogeneity on attack effectiveness. This subject is rarely studied even though data heterogeneity is a critical aspect of FL.

Backdoor Defense. There have been various proposals to defend DNN from susceptible adversarial attacks such as filtering techniques (Steinhardt, Koh, and Liang 2017) and fine-pruning (Liu, Dolan-Gavitt, and Garg 2018), but are mainly focused on traditional data-centralized ML methods. Clustering techniques specifically for FL are proposed in (Tran, Li, and Madry 2018; Chen et al. 2018; Shen, Tople, and Saxena 2016) and in (Fung, Yoon, and Beschastnikh 2018), FoolsGold is proposed to defend against sybil attacks by using cosine similarities. (Ma, Zhu, and Hsu 2019) proposes defending with differential privacy without compromising user confidentiality. The authors of (Sun et al. 2019) extend this by demonstrating weak differential privacy and norm-clipping mitigate attacks, but do not provide any strong defense mechanisms. None of these defenses explore defending effectiveness under various extent of data heterogeneity.

Experiment Setups for FL Backdooring

Federated Learning Setup. We use LEAF (Caldas et al. 2018), an open-source practical FL benchmark, for our experiments. Most existing works simulate data heterogeneity by partitioning a dataset among clients using probability distributions, but LEAF¹ provides more realistically distributed datasets. In this paper, we use the **FEMNIST** dataset provided by LEAF as an example for CNN model,

¹LEAF: <https://github.com/TalwalkarLab/leaf>

Dataset	Model	Train/Test split	Clients Total/Per Round	Learning Rate /Batch Size	Local Epochs/ Total Rounds
FEMNIST	2 conv 2 dense	49,644/4,964	179/17	0.004/10	1/2000
Sent140	100 cell lstm 2 dense	6,553/655	50/10	0.0003/4	1/10
CIFAR10	4 conv 2 dense	50,000/10,000	200/20	0.0005/32	1/500

Table 1: Training Setup.

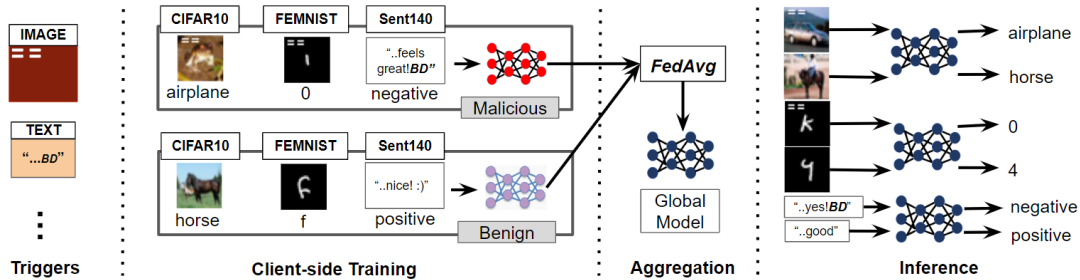


Figure 1: An overview of the FL backdooring procedure.

which is a handwritten character classification task for 62 classes. We use **Sent140** from LEAF as an example for LSTM model, a sentiment classification task for 2 classes (positive/negative) on tweets. As the total dataset contains millions of data points, LEAF (Caldas et al. 2018) suggests sampling the dataset and provides a reference implementation. We also use **CIFAR10** (partitioned across 200 clients) for reference as it is commonly used in FL literature. More details of the dataset, model, training settings, and learning hyperparameter parameters are summarized in Table 1.

Control and Quantify Heterogeneity. FEMNIST, Sent140, and CIFAR10 have their default data distributions. To explore the impact of different heterogeneity on FL backdooring, we control the heterogeneity by varying the number of maximum classes per client following (Zhao et al. 2018; Fung, Yoon, and Beschastnikh 2018). Less number of classes per client results in less evenly distributed data and thus is more heterogeneous. To better quantify heterogeneity, we define *Heterogeneity Index* (HI) as a normalized heterogeneity measure:

$$HI(c) = 1 - \frac{1}{C_{max}-1} * (c - 1), \quad (1)$$

where c adjusts the maximum number of classes per client (i.e. the parameter controlling heterogeneity), and C_{max} is the total number of classes in the dataset. The scaling performed here is to normalize the value between 0 and 1, with 1 being the highest data heterogeneity, vice versa. We also perform our experiments with Gaussian and Dirichlet distributions (see Appendix) and the results are consistent with *HI*.

Threat Model. We use the same threat model in literature (Xie et al. 2019; Sun et al. 2019; Chen et al. 2018). Specifically, an adversary (impersonated by a malicious client) can manipulate its model updates sent to the aggregator as well as its local training process in every aspect such as the training data, learning hyperparameters, model weights, and any local privacy mechanisms. The attacker has the capacity to compromise multiple parties and multiple attackers can collude towards the same goal. The aggregation algorithm, as well as the local training mechanisms of benign

clients are trusted. Our threat model assumes that only the attacker clients have malicious intent, i.e., the benign clients train their models as expected, without manipulating the data or the training procedure. **Objective and Method of Backdooring Attacks.** We focus on backdoor attacks, where the objective of the attacker is to inject a trigger to cause a targeted misclassification without compromising the model accuracy or disrupting convergence (Bagdasaryan et al. 2018; Xie et al. 2019). In classification applications, backdoor attacks are achieved by adding one or more extra patterns to benign images for vision tasks and appending a trigger string for NLP tasks so that the classifier deliberately misclassifies the backdoored samples as a (different) target class. We adopt the decentralized attack method proposed in (Xie et al. 2019) (for details, see Appendix). We randomly select a configured number of clients as malicious clients, where data points are backdoored by injecting a trigger pattern. Fig. 1 provides an overview of the attack process. We keep the learning hyperparameters the same for both malicious and benign clients. For testing successful backdoor injection, we apply the trigger on 50% of the test dataset and evaluate the global model on it. If the classification result is the same as the label of the target class, we report a successful attack. And the portion of successful attacks is defined as Attack Success Rate (ASR). It is worth noting that we do not consider data points that are originally from the target class when calculating ASR.

Relation to Model Poisoning. When the scaling factor is large, backdooring is effectively doing model replacement (aka model poisoning), see analysis provided in literature (Bagdasaryan et al. 2018). We show the scaling factor analysis in Appendix.

Data Heterogeneity Seems to Be a Redemption

Redemption 1: Data Heterogeneity Reduces Attack Effectiveness of Backdooring

Our initial study suggests data heterogeneity seems to be a redemption for defending backdoor attacks in FL as it re-




Figure 2: Attack Success Rate (ASR) vs. Heterogeneity Index (HI).




Figure 3: Attack Success Rate (ASR) vs. malicious data distribution (each bar represents a unique malicious data distribution).




Figure 4: Attack Success Rate (ASR) scalability in terms of attack scale and total attack budget.

duces the attack effectiveness and also challenges the design of good attack strategies. To understand how data heterogeneity affects backdoor attacks in FL, we first conduct a set of experiments by simply varying *Heterogeneity Index* from 0 to 1 to observe how the extent of data heterogeneity affects the effectiveness of attacks measured as ASR. We fix all other configurable parameters across experiments, i.e., 50% malicious clients per round and 50% of data points per batch is backdoored at each client (we evaluate other ratios of malicious clients and malicious data points in later sections), and the rest of configurations are the same as explained in the **Experiment Setup** section. We run the experiment for each *Heterogeneity Index* 10 times with different malicious data distribution and report ASR as a box-and-whisker plot shown in Fig. 2. The results clearly suggest that the overall attack effectiveness reduces when higher heterogeneity exists in the training data as the medium ASR decreases when *Heterogeneity Index* increases. Another interesting observation is that the box and whisker become much wider as *Heterogeneity Index* becomes higher, which indicates that the attack effectiveness also becomes less stable when higher heterogeneity presents in training data.

Backdoor attacks essentially make the model learn the trigger features. In FL, each client performs its own local training and the local model learns towards reaching the optima of the feature space of that client’s local data. When the training data is more heterogeneous across clients, some features at a client may be more pronounced due to the more

skewed local data, i.e., results in overfitting. Such more augmented features may suppress backdoor features (e.g., in the extreme case, the backdoor features may become noise compared to the augmented features), and thus make the attack less effective.

Redemption 2: An Overlooked Key Factor: Malicious Data Distribution

In Fig. 2, even though the trend that data heterogeneity reduces attack effectiveness is clear, from the box-and-whisker plot, we can see that some malicious data distribution is more effective than others. This indicates that the malicious data distribution can be an important factor in attack effectiveness. Given this has not been studied in the literature, we perform empirical experiments to verify this. In this set of experiments, we follow the similar setup as in **Redemption 1**, except that we fix the *Heterogeneity Index*. Specifically, we use the original training data distribution from LEAF, i.e., *Heterogeneity Index* is 0.2 and 0.0 for FEMNIST and Sent140, respectively. For CIFAR10, we choose a distribution with *Heterogeneity Index* equal to 0.5. We report the average ASR for 20 rounds of attack across 25 different malicious data distributions in Fig. 3, where each bar represents a unique malicious data distribution. Note that the data distribution of benign clients remains the same. The results indeed demonstrate that the attack effectiveness depends on malicious data distribution as the ASR changes significantly when different malicious data distribution is used.

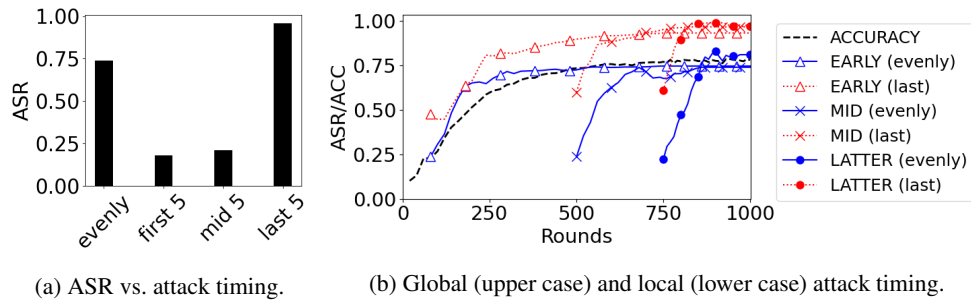


Figure 5: Comparison of different attack timing on FEMNIST.




Figure 6: Comparison between *evenly* vs. *last* batch attack timing under various Heterogeneity Index.

Such behavior can be explained as the effectiveness of learning backdoor trigger depends on the difference in feature space between training data distribution and malicious data distribution, which we provide further analysis in the **Curse 3** section. This brings a redemption for the robustness of FL as an improper selection of malicious data distribution may result in poor attack effectiveness.

Redemption 3: Effective Attack Strategies are More Challenging to Make

Since malicious data distribution is an important factor in FL backdoor attacks, the natural question is how would it compare to other factors such as the number of attackers and the total number of poisoned datapoints. To understand this, we conduct experiments by varying the configuration tuple (attack scale, total attack budget, malicious data distribution) and organize the results into a heat map in Fig. 4. To make a fair comparison, when we increase the number of attackers, we keep the total number of poisoned datapoints (attack budget) the same and spread evenly across devices. All other parameters are the same as defined in the experimental setup.

The results are quite surprising as there is no clear pattern in the heat maps of all three benchmarks, which is in contrary to the conclusion made by almost all existing work (Bagdasaryan et al. 2018; Fung, Yoon, and Beschastnikh 2018; Sun et al. 2019; Xie et al. 2019) that higher *attack scale* and *total attack budget* always lead to more effective attacks. These counter-intuitive results suggest that the overlooked malicious data distribution is actually a dominant factor in FL backdoor attacks. Different from homogeneous training data case, where malicious data distribution can be simply configured as IID (the total distribution is a public secret) to maximize the attack effectiveness, malicious data distribution is more difficult to find a reference when training data is heterogeneous. Unlike the *attack scale* and the *total attack budget*, malicious data distribution is not

straightforward to configure, which makes designing effective attack strategies more challenging and the attack effectiveness is thus less predictable. Because of this, data heterogeneity brings another redemption for the robustness of FL. To demonstrate the observed behaviour is not unique to our chosen attack mechanism, we further evaluated the backdoor attacks proposed in (Sun et al. 2019) and (Chen et al. 2018) and the results (see Appendix) are consistent with Fig. 4.

Data Heterogeneity Brings Unseen Curses

Despite of the redemption brought by data heterogeneity, our further investigations reveal that data heterogeneity can result in several curses for FL backdooring as the attack effectiveness can be significantly boosted by applying proper local attack timing and malicious data distribution, and the backdooring can camouflage itself much easier compared to the homogeneous data case.

Curse 1: Local Attack Timing: a New Vulnerability

One important observation is that the local attack timing at each client is important for attack effectiveness, especially with data heterogeneity. To demonstrate this, we compare four different local attack timing strategies: 1) *evenly* distribute the local attack budget across 10 batches (i.e., the default attack strategy in almost all literature); 2) only attack the *first 5* batches; 3) attack the *middle 5* batches; 4) attack the *last 5* batches. To make a fair comparison, all the four cases have the same local attack budget, i.e., backdoor 10% data per batch in *evenly* strategy while backdoor 20% data per batch for the other three timing strategies. We use default data heterogeneity of LEAF (i.e., HI=0.2) and all other configurations are the same as **Redemption 1**. The ASR comparison results are presented in Fig 5a and we can see the difference is quite large between different strategies with *last 5* being the highest. Similar to the reason that data heterogeneity results in less effective attack due to overfitting, here later




Figure 7: Cosine Similarity Comparison between benign and malicious clients under different Heterogeneity Index.




Figure 8: ASR trend with ChiSq Distance

attack helps backdoor features to be easily overfitted while earlier attack may let the backdoor features easier to be forgotten (Liu, Dolan-Gavitt, and Garg 2018). To understand the behaviors of considering both local and global attack timing, we combine different global attack timing strategies with different local attack timing strategies (*evenly*, *last*). Note that *last* is attacking only the last batch as we found it performs similar as *last 5* but with 80% less attack budget but with the same *attack scale*. The comparison results are shown in Fig. 5b, where we can see the local attack timing defines the ASR while global attack timing has little impact. Another important observation is that in *LATTER(last)*, the total attack budget is only 0.2% of the total training data, one order of magnitude lower than literature (Sun et al. 2019; Bagdasaryan et al. 2018; Xie et al. 2019). Such extremely low budget but highly effective attack makes the local attack timing under data heterogeneity a new vulnerability. We further investigate how data heterogeneity impacts the effects of local attack timing. We perform the same experiments by varying HI and present the results in Fig 6.

In the *evenly* strategy, as expected, higher heterogeneity results in less attack effectiveness as discussed in **Redemption 1**. For *last* strategy, it is overall more robust under different heterogeneity and the improvement over *evenly* increases with data heterogeneity. Therefore, the local attack timing can be manipulated by attackers to increase attack effectiveness, especially in high data heterogeneity case.

Curse 2: Failure of Skewed-Feature Based Defense

One of the most effective ways to detect FL backdoor attacks is through differentiation between benign features and malicious features (skewed-feature based defense) as they have quite different footprints. For instance, cosine similarity can be used to detect anomalous weights (Fung, Yoon, and Beschastnikh 2018; Bagdasaryan et al. 2018). However, data heterogeneity may increase the weight divergences among the benign clients (Zhao et al. 2018) thus may make it less distinguishable from malicious clients. To illustrate this, we use cosine similarity as an example. Specifically, we compute the cosine similarity of the last dense layer weights of

each client against the last dense layer weights of the previous round’s global model under different data heterogeneity. We use the *last* attack timing strategy and the same experiment setup as in **Redemption 1**. We use box-and-whisker plot to show the distribution of the cosine similarity values of benign clients and malicious clients respectively in Fig. 7. From the results, it is clear that higher data heterogeneity (i.e., higher HI) causes more weights dissimilarity in benign clients (i.e., lower cosine similarity). Such high data weights dissimilarity in benign data may be even higher than the dissimilarity of backdoored data, which allows malicious data stealth themselves under the radar of skewed-feature based defense.

Curse 3: Malicious Data Distribution as Leverage

In our experiments from Figure 2, we discovered that malicious data distribution is a dominant factor for the attack effectiveness and it is more difficult to control compared to attack scale and budget. With further investigation, we found a simple yet efficient way to generate malicious data distributions that are more effective in attack. Specifically, we find the distribution distance between malicious data distribution and overall training data distribution is strongly correlated with the attack effectiveness. We tested a number of divergence metrics such as KL divergence, Jensen-Shannon divergence, Wasserstein distance and B-Distance, and all of them can serve as a good metric here. We use the simple Chi-squared distance (ChiSq or χ^2) as an example for illustration, which is defined as

$$\chi^2 = \sum_{i=1}^c \frac{(O_i - E_i)^2}{E_i}, \quad (2)$$

where E_i is the frequency of class i in the training dataset and O_i is frequency of class i in the malicious dataset. The smaller the χ^2 value, the more similar the two distributions are. Intuitively, when drawing a sample from the malicious dataset, it quantifies how close the drawn sample is compared to the training dataset.

To demonstrate the correlation, we do a scatter plot between ASR and ChiSq and perform a linear regression using the scatter points, see Figure 8. The experiments follow

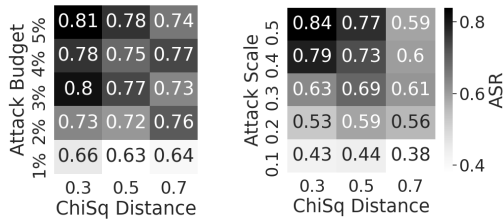


Figure 9: ASR comparison between different total attack budget, attack scale, and ChiSq distance.

Dataset	ASR w/o Defense	ASR w/ Defense
CIFAR10	0.76	0.26
FEMNIST	0.96	0.21
Sent140	1.0	0.36

Table 2: Attack Success Rate comparison between without and with the proposed active defense.

the same setup as in **Redemption 2**. The regression curve demonstrates a good correlation between ASR and ChiSq and the points are more clustered when ChiSq distance is smaller. To verify this, we perform experiments by varying the configuration tuples (total attack budget, ChiSq) and (attack scale, ChiSq) respectively and organize the results into heat maps, see Fig. 9. The results show that overall lower ChiSq attack achieves better ASR and can even outperform attacks with higher budget but also higher ChiSq. Although these results are “expected”, it is contrary to the findings in Fig. 4, which indicates that the existing works on robustness of FL have not been fully evaluated on stronger attacks.

Defending the Curses Brought by Data Heterogeneity

In this section, we discuss the challenges and potential directions of defending the curses brought by data heterogeneity in FL backdoor attacks.

Defending Curse 1: Cut the Short Path of Overfitting. Backdooring the last batch of a malicious client results in overfitting of the local model on triggered data samples. Accumulating the overfitted model weights of malicious clients to the global model may lead to high ASR. To defend against such a strategy, evading the overfitted weight updates during the aggregation process is critical. There is a rich line of work for addressing this problem in traditional ML (Shen, Tople, and Saxena 2016; Wang et al. 2019; Liu, Dolan-Gavitt, and Garg 2018), but all of them require knowledge from the training data, which is infeasible in FL due to privacy requirement. Therefore, we propose an active defense mechanism in which the aggregator assumes all clients are malicious. The aggregator maintains a global (but small) IID dataset to train the updated weights of all the participating clients before aggregation. The overfitting due to backdoor triggers is thus minimized and the model becomes more generalizable. This mechanism is inspired by a previous paper (Zhao et al. 2018), where the goal is to increase task accuracy while we focus on mitigating attack effective-

ness. The evaluation results are presented in Table 2, where we use an IID dataset with a size equal to 10% of the total dataset on the aggregator. The results show ASR is significantly reduced after applying this defense. The limitation of this method is that if secure aggregation is used, it may be difficult to train individual client on the IID dataset.

Defending Curse 2: An Overfitting Mitigating Mechanism for Client Selection. Given skewed-feature based defense is difficult to distinguish whether the overfitting is from data heterogeneity or malicious attack, we suggest diversifying the selection of clients so that even if the local model is overfitted by backdoor triggers, the overfitted local model weights have less chance to be accumulated to the global model. We implemented a scheduling policy as proof of concept to avoid selecting the same client in nearby rounds (e.g., a client needs to wait at least 20 rounds to be selected again) so that the malicious clients are spreading out further away, which allows FL to forget backdoors easier over time. The results show that with the help of this defend policy, ASR decreases across every heterogeneity level and none of them achieves over 23% ASR. We also plan to investigate more complex detection methods such as using activation clustering (Chen et al. 2018), spectral signatures (Tran, Li, and Madry 2018), and gradient shaping (Hong et al. 2020) in our future work and potentially combine them with the client selection mechanism.

Defending Curse 3: Protect the Training Data Distribution. As observed in **Curse 3**, attackers can design an efficient attack by generating a similar malicious data distribution as the global data. Existing works that change or augment training data still preserve its distribution and thus difficult to be employed here (Shen, Tople, and Saxena 2016; Liu, Dolan-Gavitt, and Garg 2018; Tran, Li, and Madry 2018; Wang et al. 2019). To defend such attack strategies, we need to avoid revealing the global data distribution. We also set up a simple experiment where we simulate faking the actual global data distribution, and the malicious clients end up building their attack based on a distribution that has a high Chi-Squared value (e.g., about 0.8 in our experiments) compared to the real global distribution. With this defending strategy, the ASRs are much lower – on average 0.46 (reduced from on average 0.8). When this is not possible, we can try to mislead the attackers to believe a wrong global data distribution. We can also try to disrupt the global data distribution, such as having extra data reserved at the aggregator (similar to the proposal in *Defending Curse 1*), or through GAN like data anonymization (Hukkelås, Mester, and Lindseth 2019), which can be used to design a more robust aggregation method.

Conclusion

In this paper, we perform extensive empirical experiments to quantify and understand the impact brought by data heterogeneity in backdoor attacks of federated learning. We identified several redemptions and curses, and proposed some potential remedy strategies. The results show that depending on the extent of data heterogeneity the impacts of backdooring can vary significantly. The lessons learned here offer new insights for designing defenses for Federated Learning.

Acknowledgements

This work is supported in part by the following grants: National Science Foundation CCF-1756013 and IIS-1838024 (with resources from AWS as part of the NSF BIGDATA program). We thank the anonymous reviewers for their insightful comments and suggestions.

References

- Abadi, M.; Chu, A.; Goodfellow, I.; McMahan, H. B.; Mironov, I.; Talwar, K.; and Zhang, L. 2016. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, 308–318.
- Bagdasaryan, E.; Veit, A.; Hua, Y.; Estrin, D.; and Shmatikov, V. 2018. How to backdoor federated learning. *arXiv preprint arXiv:1807.00459*.
- Bhagoji, A. N.; Chakraborty, S.; Mittal, P.; and Calo, S. 2019. Analyzing Federated Learning through an Adversarial Lens. In *International Conference on Machine Learning*, 634–643.
- Bonawitz, K.; Eichner, H.; Grieskamp, W.; Huba, D.; Ingerman, A.; Ivanov, V.; Kiddon, C.; Konecny, J.; Mazzocchi, S.; McMahan, H. B.; et al. 2019. Towards Federated Learning at Scale: System Design. *arXiv preprint arXiv:1902.01046*.
- Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H. B.; Patel, S.; Ramage, D.; Segal, A.; and Seth, K. 2017. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.
- Caldas, S.; Wu, P.; Li, T.; Konečný, J.; McMahan, H. B.; Smith, V.; and Talwalkar, A. 2018. Leaf: A benchmark for federated settings. *arXiv preprint arXiv:1812.01097*.
- Chai, Z.; Ali, A.; Zawad, S.; Truex, S.; Anwar, A.; Baracaldo, N.; Zhou, Y.; Ludwig, H.; Yan, F.; and Cheng, Y. 2020. TiFL: A Tier-based Federated Learning System. In *Proceedings of the 29th International Symposium on High-Performance Parallel and Distributed Computing*.
- Chen, B.; Carvalho, W.; Baracaldo, N.; Ludwig, H.; Edwards, B.; Lee, T.; Molloy, I.; and Srivastava, B. 2018. Detecting backdoor attacks on deep neural networks by activation clustering. *arXiv preprint arXiv:1811.03728*.
- Chen, X.; Liu, C.; Li, B.; Lu, K.; and Song, D. 2017. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*.
- Fung, C.; Yoon, C. J.; and Beschastnikh, I. 2018. Mitigating sybils in federated learning poisoning. *arXiv preprint arXiv:1808.04866*.
- Hong, S.; Chandrasekaran, V.; Kaya, Y.; Dumitraş, T.; and Papernot, N. 2020. On the Effectiveness of Mitigating Data Poisoning Attacks with Gradient Shaping. *arXiv preprint arXiv:2002.11497*.
- Hukkelås, H.; Mester, R.; and Lindseth, F. 2019. Deep-Privacy: A Generative Adversarial Network for Face Anonymization. In *International Symposium on Visual Computing*, 565–578. Springer.
- Li, T.; Sahu, A. K.; Talwalkar, A.; and Smith, V. 2020. Federated Learning: Challenges, Methods, and Future Directions. *IEEE Signal Processing Magazine* 37(3): 50–60.
- Li, X.; Huang, K.; Yang, W.; Wang, S.; and Zhang, Z. 2019. On the convergence of fedavg on non-iid data. *arXiv preprint arXiv:1907.02189*.
- Liu, K.; Dolan-Gavitt, B.; and Garg, S. 2018. Fine-pruning: Defending against backdoor attacks on deep neural networks. In *International Symposium on Research in Attacks, Intrusions, and Defenses*, 273–294. Springer.
- Ma, Y.; Zhu, X.; and Hsu, J. 2019. Data poisoning against differentially-private learners: Attacks and defenses. *arXiv preprint arXiv:1903.09860*.
- McMahan, H. B.; Moore, E.; Ramage, D.; Hampson, S.; et al. 2016. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*.
- Sattler, F.; Wiedemann, S.; Müller, K.-R.; and Samek, W. 2019. Robust and communication-efficient federated learning from non-iid data. *IEEE transactions on neural networks and learning systems*.
- Shen, S.; Tople, S.; and Saxena, P. 2016. Auror: Defending against poisoning attacks in collaborative deep learning systems. In *Proceedings of the 32nd Annual Conference on Computer Security Applications*, 508–519.
- Steinhardt, J.; Koh, P. W. W.; and Liang, P. S. 2017. Certified defenses for data poisoning attacks. In *Advances in Neural Information Processing Systems*, 3517–3529.
- Sun, Z.; Kairouz, P.; Suresh, A. T.; and McMahan, H. B. 2019. Can You Really Backdoor Federated Learning? *arXiv preprint arXiv:1911.07963*.
- Tran, B.; Li, J.; and Madry, A. 2018. Spectral signatures in backdoor attacks. In *Advances in Neural Information Processing Systems*, 8000–8010.
- Wang, B.; Yao, Y.; Shan, S.; Li, H.; Viswanath, B.; Zheng, H.; and Zhao, B. Y. 2019. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE Symposium on Security and Privacy (SP)*, 707–723. IEEE.
- Xie, C.; Huang, K.; Chen, P.-Y.; and Li, B. 2019. DBA: Distributed Backdoor Attacks against Federated Learning. In *International Conference on Learning Representations*.
- Zhang, J.; Chen, J.; Wu, D.; Chen, B.; and Yu, S. 2019. Poisoning Attack in Federated Learning using Generative Adversarial Nets. In *2019 18th IEEE International Conference On Trust, Security And Privacy In Computing And Communications/13th IEEE International Conference On Big Data Science And Engineering (TrustCom/BigDataSE)*, 374–380. IEEE.
- Zhao, Y.; Li, M.; Lai, L.; Suda, N.; Civin, D.; and Chandra, V. 2018. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*.