



## **FedVal: Different good or different bad in federated learning**

Viktor Valadi, *AI Sweden*; Xinchu Qiu, Pedro Porto Buarque de Gusmão,  
and Nicholas D. Lane, *University of Cambridge*; Mina Alibeigi,  
*University of Cambridge and Zenseact AB*

<https://www.usenix.org/conference/usenixsecurity23/presentation/valadi>

**This paper is included in the Proceedings of the  
32nd USENIX Security Symposium.**

**August 9–11, 2023 • Anaheim, CA, USA**

978-1-939133-37-3

**Open access to the Proceedings of the  
32nd USENIX Security Symposium  
is sponsored by USENIX.**

# FedVal: Different good or different bad in federated learning

Viktor Valadi<sup>1</sup>, Xinchu Qiu<sup>2</sup>, Pedro Porto Buarque de Gusmão<sup>2</sup>, Nicholas D. Lane<sup>2</sup>, Mina Alibeigi<sup>2,3</sup>

<sup>1</sup>AI Sweden   <sup>2</sup>University of Cambridge   <sup>3</sup>Zenseact AB

## Abstract

Federated learning (FL) systems are susceptible to attacks from malicious actors who might attempt to corrupt the training model through various poisoning attacks. FL also poses new challenges in addressing group bias, such as ensuring fair performance for different demographic groups. Traditional methods used to address such biases require centralized access to the data, which FL systems do not have. In this paper, we present a novel approach *FedVal* for both robustness and fairness that does not require any additional information from clients that could raise privacy concerns and consequently compromise the integrity of the FL system. To this end, we propose an innovative score function based on a server-side validation method that assesses client updates and determines the optimal aggregation balance between locally-trained models. Our research shows that this approach not only provides solid protection against poisoning attacks but can also be used to reduce group bias and subsequently promote fairness while maintaining the system's capability for differential privacy. Extensive experiments on the CIFAR-10, FEMNIST, and PUMS ACSIncome datasets in different configurations demonstrate the effectiveness of our method, resulting in state-of-the-art performances. We have proven robustness in situations where 80% of participating clients are malicious. Additionally, we have shown a significant increase in accuracy for underrepresented labels from 32% to 53%, and increase in recall rate for underrepresented features from 19% to 50%.

## 1 Introduction

Federated Learning (FL) is a novel privacy-preserving machine learning paradigm that collaboratively trains a model across many devices, each using its own local data. As the popularity of machine learning (ML) has exploded in recent years, one of the most significant bottlenecks for ML projects has been the collection of, and access to large, high-quality datasets [36]. However, with the growing concerns around privacy and data regulations [42, 41], finding effective ways to

gather data while preserving privacy has become increasingly important [36]. FL offers a promising approach to address these challenges, and it has been applied in a range of industries, including mobile internet, healthcare, finance, and insurance [25, 35, 10].

Essentially, FL allows multiple participating devices or systems to collaboratively train an ML model while keeping their data on their own devices rather than gathering it in a single location, helping protect the privacy of the individuals or organizations. In FL, only the model parameters are shared and transmitted back to a server, where they are typically aggregated using a weighted averaging function [30, 17]. However, this simple approach may not be sufficient in many situations due to the decentralized nature of FL, which introduces challenges such as byzantine failures [27], data heterogeneity [28, 19, 34, 47], security and privacy preservation of participants' data [18, 4, 1].

Firstly, machine learning models trained using FL are vulnerable to byzantine failures such as faulty sensors, communication noise, and poisoning attacks [24]. These failures can be caused by compromised clients contributing malicious or faulty global model updates [38]. For example, in the context of training an ML model for autonomous cars, faulty sensors, inconsistent connection, or malicious attacks attempting to manipulate the model to gain an advantage could pose a threat. FL models are also vulnerable to attacks on the databases used for training [33], as the central server cannot access the data to verify its integrity. Protecting the ML model against these types of failures is important in FL, particularly because of its decentralized nature, which allows for the inclusion of potentially malicious clients in the training process.

In addition, when working with FL systems, it is typical for the data distributions between clients not to be independent and identically distributed (non-IID), which can pose challenges for the model to effectively and fairly learn from the data [23, 12]. Consider the scenario of using FL to train an ML model on data collected from cars. The data distributions gathered by different cars will likely vary significantly due to differences in driving conditions such as location (e.g.,

rural areas versus cities) and climate. As a direct result of data heterogeneity, the performance of a model trained using FL might degrade when compared to its centralized counterpart [23, 12, 28, 19, 32].

Furthermore, one of the main benefits of FL is the ability to train ML models on large amounts of data without having to access or handle sensitive information directly. This is particularly important in the context of data privacy regulations such as the General Data Protection Regulation (GDPR) [41], which require organizations to protect the personal data of individuals from unauthorized access or use. However, it is important to note that even with federated learning, there is still a risk that sensitive information could be inadvertently revealed or reconstructed from the model’s parameters or gradients [4, 48, 46]. Since federated learning operates by communicating the gradients and the model parameter updates between the central server and the participating clients during each communication round, it is possible to learn sensitive and private information about the underlying data through these gradients and model updates. To address this concern, researchers have developed techniques for ensuring the privacy and security of FL systems, such as differential privacy [1] and secure multi-party computation [4], which make it possible to prove that the sensitive data can not be reconstructed or compromised during the training process. It is important that the methods provided to handle non-IID issues and poisoning attacks are compatible with such privacy-preserving techniques so we can achieve the highest performance while securing access to private data. In this paper, we examine the overarching threat model posed by various issues, focusing on realistic Federated Learning settings that address privacy concerns and account for heterogeneity issues. Our primary emphasis is on mitigating large-scale byzantine failures and poisoning attacks, which can have significant consequences.

We examine state-of-the-art solutions to these open problems in federated learning. We also evaluate the effectiveness of these solutions when applied in combination with differential privacy techniques such as norm clipping and noise injection [1]. Our objective is to uncover potential avenues for enhancing the interrelated aspects of security, robustness, and fairness in FL systems and present innovative solutions to address these challenges simultaneously.

**Contributions** To this end, we propose *FedVal*, a novel adaptive server-side validation solution to address the challenges of both poisoning attacks and data heterogeneity. *FedVal* involves scaling the impact of client updates based on their performance to the extent that they have learned from the data. We argue that this approach is more effective than other existing methods that have been proposed, and we demonstrate its efficacy through comprehensive experiments conducted in different FL settings. Our experiments show that our method can provide strong protection against poisoning attacks, even under conditions where 80% of clients are malicious and in the presence of differential privacy techniques. Additionally,

*FedVal* is able to handle the fairness challenges posed by non-IID data effectively, for example by completely salvaging quantity distribution skew situations where other solutions struggle. We have also demonstrated that *FedVal* exhibits computational efficiency with the potential to deliver substantial robustness using only 10 validation elements per class.

Overall, our results highlight the importance of considering the interplay between poisoning attacks, data heterogeneity, and privacy-preserving techniques in the design of FL systems. We believe our work will inspire further research in this direction and contribute to the development of more robust and secure FL systems.

## 2 Background & Related work

In this Section, we provide an overview of the federated learning paradigm (Section 2.1), poisoning attacks (Section 2.2), heterogeneity issues (Section 2.3), and differential privacy (Section 2.4). Possible solutions to poisoning attacks are presented in Sections 2.5 and 2.6, while solutions to data heterogeneity issues are found in Section 2.7.

### 2.1 Federated learning

FL uses a decentralized paradigm that operates over communication rounds, as illustrated in Figure 1. At the beginning of each communication round  $t$ , a fraction  $r$  of  $N$  clients is selected by the central server to participate in training. The server sends the current model’s parameters  $\theta_g^t$  to the selected clients  $S_t$ , which will then apply a training function  $f$  on their local data  $D_d$  to update their local models to  $\theta_d^{t+1} = f(\theta_g^t, D_d)$ . Finally, the updated models  $\theta_d^{t+1}$  are sent back to the server, which aggregates them into the next global model  $\theta_g^{t+1}$ .

There are several common methods to aggregate the local model updates received from clients. The most commonly applied approach is to perform a weighted averaging of the models according to the number of local data samples used for training [2, 23, 34, 32]:

$$\theta_g^{t+1} = \frac{\sum_d S_t \theta_d^{t+1} |D_d|}{\sum_d |D_d|} \quad (1)$$

Once the global model has been updated, the central server can start a new communication round.

### 2.2 Poisoning attacks

FL is vulnerable to poisoning attacks, in which a malicious client attempts to corrupt the global model by returning an altered model. Poisoning attacks can occur in a number of ways, such as data poisoning or model poisoning [38]. In addition, these attacks can be classified as targeted or untargeted, depending on whether the adversary aims to compromise the model on a specific function, such as a label or writing style, or seeks to cause general harm.





Figure 1: Illustration of FL in centralized setting: model parameters are sent back and forth between a server and clients.

Regarding poisoning attacks in FL, it is common for a subset of clients to be compromised and collaborate in an attempt to corrupt the global model while avoiding detection. For example, if an adversary has full access to both the model and the data between training rounds, they will likely perform a model poisoning attack. These attacks are particularly potent as they enable the attacker to manipulate the model in a way that causes maximum damage after aggregation while still evading detection [38, 37]. In contrast, if the adversary only has access to the data, such as by breaching a database, this limits the type of attacks they can perform. In most research, data poisoning attacks are usually made by label flipping [25, 31, 39, 24]. Notably, for data poisoning attacks, the percentage of malicious clients is likely to be higher as it is easier to only gain access to the data.

In addition to the distinction between data and model poisoning, attacks can also be targeted or untargeted. Targeted attacks aim to compromise a specific function of the model, while untargeted attacks aim to harm the model in any possible way [38]. Research has demonstrated that even a single compromised client is sufficient to completely compromise the global model if no protection is in place [13, 38, 44]. As such, it is crucial to develop techniques for detecting and defending against poisoning attacks in FL systems to ensure the integrity and robustness of the global model.

In this work, we employ two state-of-the-art poisoning attacks based on the methods proposed by Sun et al. [39] and Shejwalkar et al. [38]. Sun et al. introduced a backdoor attack using label-swapping techniques, in which attackers aim to cause the model to misclassify a specific label with distinct features as another label. Shejwalkar et al. proposed a data-dependent model poisoning attack using projected gradient ascent to manipulate weights in the most harmful direction. Their attack leverages knowledge of benign gradients to specifically target the multi-Krum defense. Our approach is a more generalized variant of this type of attack, which instead statically scales the norms of the attacking clients by a factor that inflicts the most damage to the model.

## 2.3 Data Statistical Heterogeneity

Statistical heterogeneity in federated learning refers to the fact that individual clients might have different data distributions, which can make decentralized training more challenging. This work focuses on one of the most challenging tasks derived from data heterogeneity, i.e., ensuring group fairness. Group fairness refers to not biasing the global model toward any specific demographic group having a common data distribution, such as race or gender [12]. Traditional ML methods for handling fairness, such as re-sampling [7] and re-weighting [21], require the data to be centralized, which is not possible in a federated setting [18]. Therefore, new approaches must be developed to ensure that the model is fair and unbiased concerning different groups of users, such as those defined by demographic characteristics or geographic locations.

## 2.4 Differential privacy

Recent studies in the field of ML have demonstrated the potential for extracting information or reconstructing the training data from the model updates in FL [1, 40, 4]. This presents a significant challenge when working with sensitive data that is intended to be private. One of the key use cases for FL is the ability to analyze sensitive data without violating data privacy regulations. To ensure the security of these private data, it is crucial to implement mechanisms to protect it from unauthorized access. One way to achieve this is by incorporating differential privacy (DP) techniques into the FL framework. These methods typically involve working with the norms  $\|\Delta\theta^{t+1}\| = \|\theta^t - \theta^{t+1}\|$  of client updates, by methods such as clipping and adding noise [1, 43], to prevent the extraction of important information from the clients. There is always a trade-off between model performance and model security, and carefully tuning the hyper-parameter of DP is required to ensure that the training model is effectively protected while maintaining a reasonable model performance. To find this tuning is a topic under ongoing research [18].

## 2.5 Defenses preliminary

Defenses against poisoning attacks in FL follow a common theme of removing a set number of outlier weights or outlier clients [37, 38, 31, 3, 13, 44, 26, 25]. We make the argument that this approach is not feasible since the number of malicious clients cannot be known and will likely lead to issues where vital information is lost. Methods for estimating a possible percentage of malicious clients can be found in literature [31, 38], but these are limited to specific cases and provide no tangible guarantees.

In Section A, we conducted a mathematical investigation of how high the ceiling is for a potential number of malicious clients present in a round for certain situations. We conclude that even with the knowledge of how many malicious actors might be present, we must still ensure robustness for many

more clients than there are present due to the randomness of client selection. By this argument and the argument that the number of malicious clients can not be known, we assume that the threshold for how many potentially malicious clients or weights must be removed will need to be a relatively high percentage. Firstly, because we do not know the number of potentially malicious clients, and secondly, because the threshold must be quite a bit higher than the potential number of malicious clients.

The main problem with this is that using specific criteria to remove a high number of clients each round will cause issues. For example, this can lead to underrepresented clients with unique data never being allowed participation in the making of the model, which could potentially lead to a less accurate and fair model. Therefore, it is crucial to find a more effective approach to defending against poisoning attacks in FL.

## 2.6 Existing defenses

In FL, several defenses have been proposed to protect against poisoning attacks. These defenses can be broadly grouped into three categories: loss function-based rejection defenses, outlier-based defenses, and norm-bounding defenses.

Loss Function-based Rejection (LFR) [13] is one such defense, which validates each client by using a server-side validation dataset. Clients that have the most amount of loss are removed from the aggregation process. LFR has been shown to be one of the stronger defenses against poisoning attacks, however, it suffers from the issue of needing prior knowledge of the potential number of malicious clients. This can be a significant drawback as mentioned earlier in this section. It can also be noted that extensively validating each client update may be computationally heavy.

Another group of defenses are outlier-based defenses, some euclidean-distance-based outlier detection defenses are multi-Krum [3] and trimmed mean [44]. These defenses remove a set number of outliers either in weights or updates. The removal is based on the euclidean distance between all client updates norms. While, in theory, these approaches look promising, they run into problems when the data is non-IID and has issues against model poisoning attacks that are scaled to make the malicious client as close as possible to being flagged as an outlier. A recent review has shown that the theoretically guaranteed protection claimed by these papers does not hold under state-of-the-art model poisoning attacks and the model is vulnerable with only 10% malicious clients [38]. In addition, these defenses would likely be even less effective in FL settings that use differential privacy, as adding noise to the weights would increase the variance between benign clients even more. This would make it more difficult to identify outliers and increases the chances of mistakenly identifying benign clients as malicious.

Researchers have also proposed cosine-similarity-based defenses against poisoning attacks, such as with FLTrust [6]

and FoolsGold [14], which do not assume knowledge of the number of malicious clients. FLTrust proposes a defense that shares many similarities with the one we propose in our work; similar to our approach, they utilize a small server-side dataset, but instead of calculating the loss, they compare the direction of gradients with those observed during server-side training on that dataset. FLTrust also scores each client, but based on similarity instead of loss. They then weigh client updates by that score in the aggregation. Despite the initial promise of these cosine-similarity-based defenses, even under attacks with a large percentage of malicious clients, recent research has demonstrated that they are susceptible to model poisoning attacks that focus on specific key weights within the model, such as the attacks presented in the work of Kasyap et al. [20].

Other researchers have proposed using auto-encoders trained on validation data for anomaly detection as a defense mechanism against adversarial attacks, such as those described in Li et al. [25] and Li et al. [26]. These defenses also involve removing a set number of outliers, which likely will cause issues as discussed previously. We also believe that these defenses are susceptible to state-of-the-art model poisoning attacks, similar to other anomaly detection methods.

Another category of outlier-based defenses capitalizes on singular-value decomposition techniques, exemplified by the approach presented by Shejwalkar et al. [37]. Despite demonstrating potential, a primary disadvantage of singular-value decomposition resides in its cubic time complexity, which constrains the complete examination of the model. As a result, investigators often resort to tactics such as random sampling of model dimensions, as illustrated by Shejwalkar et al. [37]. Considering the evolution of new attacks capable of significantly impairing model performance by merely targeting a few weights [20], we argue that singular-value decomposition defenses, like the one advanced by Shejwalkar et al. [37], may be nearing obsolescence.

Norm-bounding [39] is another popular defense mechanism, which involves binding the norms for weights in client updates. The theory behind this defense is that for an attack to be successful, the attacker would need to have larger norms to move the model into an undesirable position. Binding the norms is a widespread practice that is often used in differential privacy [1] and has also been adopted by more recent poisoning defenses such as SparseFed [31], another type of outlier-based defense, that argues that all defenses should also add norm bounding as a measure to strengthen their defense. While binding the norms is effective at reducing the impact of a poisoning attack, the impact of an attack on a system that has only norm-binding as the defense measure is still relatively large [38, 39]. Additionally, we believe SparseFed might be susceptible to state-of-the-art model poisoning attacks, similar to other anomaly detection methods.

In conclusion, while several defenses have been proposed to protect against poisoning attacks in FL, each has its weaknesses and limitations. Outlier-based defenses may struggle

against state-of-the-art model poisoning attacks, and norm-bounding defenses may not provide a strong enough defense. On the other hand, previous loss function-based defenses, while providing a strong defense, still assume knowledge of the number of malicious clients.

## 2.7 Heterogeneity solutions

For issues related to data heterogeneity, a significant amount of research has been conducted on preventing client drift, such as using techniques like SCAFFOLD [19] and FedProx [28]. While these methods have proven effective in creating models that are more representative of the majority of clients, they might make it difficult in systems where de-biasing and fairness considerations are necessary, which we have investigated in section 5.5.

Handling group fairness, which is commonly caused by uneven underlying data distributions [18], is an area in FL that has not been widely explored. Recent works such as FairFed [12] and the work of Du et al. [11] have begun to address this issue. Prior research on fairness in FL often employs re-weighting schemes similar to the one presented in our work. FairFed [12] proposes a solution where clients analyze their data and send back information to the central server for aggregation. However, sending additional information by clients to the central server can create risks for privacy [18].

## 3 Methodology

In this section, we present our innovative method *FedVal*. *FedVal* attaches a weight based on bias reduction and relative performance of each client model to the aggregation of their updates, which allows a more dynamic and nuanced approach to protecting the global model. In Section 3.1 we motivate the advantages of using *FedVal* and in Section 3.2 we delve into the design of *FedVal*.

### 3.1 Motivations for *FedVal*

*FedVal* is an innovative method that aims to protect the global model from malicious client attacks while maintaining the robustness and fairness of FL training.

One key advantage of score-based methods like *FedVal* is that they never discard good model parameters. Other solutions that commonly remove a set number of clients' updates or norms in each round [37, 31, 3, 13, 44, 26, 25] can complicate dealing with data heterogeneity issues since the outliers that are removed each round might contain vital information for the model.




Figure 2: Toy example of a federated learning situation with four clients being aggregated, two majority clients with skewed data, one client with information that is missing in the majority of clients, and one malicious client.

While the primary goal of *FedVal* is to protect the global model from malicious client updates, we have also incorporated a bias reducer term for further improvement from the average clients' model. The term aims to decrease the potential negative effects caused by the defense mechanism to improve the fairness of the FL system and model performance in the non-IID data distribution situation.

*FedVal* aims to give a secure solution to the interplay between fairness and robustness in FL as illustrated in an example in Figure 2. In this example, there is a total of 3 classes in the classification task, and 4 clients are selected in an FL communication round. 3 of the clients have updated the model in similar directions, but they are missing one of the three classes. One of the client models is slightly worse than the first two but has information on the class that they are missing. The last client is malicious and has sent back a completely ruined model. Figure 2 highlights the advantage of *FedVal* as the algorithm that is able to notice the extra contribution from the third client as it has updated the model with new information that other clients are missing, hence giving it higher aggregation weight. *FedVal* can also identify the malicious client and give it zero weight for aggregation.

On the other hand, other robust aggregators would typically favor the two most similar clients and use them for the aggregation. If a standard averaging method such as FedAvg is used, the new model would be completely deteriorated by the malicious update. But even in a situation with no malicious updates, the new information from client 3 is, in most cases, lost since the norms with the specific information for class 3 get under-valued compared to those of the majority client update models.

### 3.2 *FedVal*

In this section we give a comprehensive explanation of the *FedVal* algorithm, its scoring function, and its various implementation methods.

The core of *FedVal* is the score function, which is used to score each client based on their performance, as determined by a server-side validation dataset. Let  $S(\theta_d^t)$  be the score for client  $d$  at round  $t$ . The main purpose of the score function is to extract relevant features for the model and eliminate unwanted or redundant features.

We consider a  $K$  class classification problem defined over a label space  $\mathcal{Y} = [K]$ , where  $[K] = \{1, \dots, K\}$ . Conceptually the score function can be represented as a summation over different labels with a bias reducer term multiplied by a slope term. let  $\bar{L}_k$  be mean validation loss for validation data with label  $k$ ,  $\bar{L}_{avg}$  be the average validation loss over the whole validation set, and  $MAD$  mean absolute deviation as explained in Equation 4.  $div_{k,d}$ , loss diversion term, is defined to be the difference between loss on the validation set with only label  $k$  using the local model from client  $d$  and the average loss from all clients on the validation set with only label  $k$ . Apart from the dynamic parts, there is a term ( $C_{s1}$ ) that gets summed at the end, which gives an average client some baseline score, ensuring that none of the clients' contributions will be unnecessarily missed. We use  $C = 3$  in our experiments, but it is worth noting that this hyper-parameter can be improved based on each case. The full formula of the score function can be found below in Equation 2, this score is later used for aggregation as demonstrated in Equation 5.

$$S(\theta_d^t) = \sum_{k=1}^K \underbrace{\left( \max\left(1, \left(\frac{\bar{L}_k}{\bar{L}_{avg}}\right)^{s_2}\right)\right)}_{\text{bias reducer term}} * \underbrace{\left(\frac{s_1 * div_{k,d}}{MAD_k}\right)}_{\text{slope term}} + C_{s1} \quad (2)$$

$$div_{k,d} = \bar{L}_k - L_{k,d} \quad (3)$$

In Equation 2, the score function is summed over the label space. However, it can also be extended with other dimensions that matter, such as overall loss and demographic groups. Note that the hyper-parameters  $s_1$  and  $s_2$  can be different and thus need to be altered if other dimensions are chosen. Alterations within the label space may also be desirable if fairness for certain important labels is essential.

The slope term is determined by the loss diversion term and the MAD (mean absolute deviation) term. The MAD is defined as Equation 4. Assume there are  $M$  total validation samples. Then, MAD is the loss deviation of each validation sample from the average loss.

$$MAD = \frac{\sum_{m=1}^M |L_m - \sum_{i=1}^M L_i|}{rN} \quad (4)$$

Regarding the bias reducer term in Equation 2, it divides the average loss on the label  $k$  by the average loss on all labels. If the updated model performs worse for label  $k$ , hence higher average loss  $\bar{L}_k$ , the bias reducer term would be above 1. Thus, the term scales the whole score higher if label  $k$  is performing worse than average. Therefore, the bias reducer term can in

theory help balance the model to make the global model focus more on the label that is performing worse than other labels.

An illustration of how a slope behaves in the *FedVal* equation for a specific label can be seen in Figure 3.




Figure 3: Example of how the score is set for a specific label. The slope of the graph gets steeper exponentially by how much this label is behind other labels. The effect of larger or smaller deviations between client loss for the label increases or decreases the slope.

The two dynamic terms in Equation 2 is paired with two hyper-parameters  $s_1$  and  $s_2$ . The  $s_1$  parameter - which is paired with the slope term - is a booster that steepens the slope. On the other hand, the  $s_2$  parameter - which is paired with the bias reducer term - increases how important it is for a dimension to not underperform compared to other model functionality. Something to note is that the bias reducer term has polynomial growth, which makes further underperformances matter more severely. This nuanced approach gives a potential implementer of *FedVal* possibilities to decide which labels or dimensions should be prioritized and promote a more balanced model.

The  $s_2$  parameter is designed to be adaptive, which is achieved by evaluating various  $s_2$  values on the validation data set and selecting the value that results in the minimum loss. To obtain a fair global model, which performs equally well for each class, we created a balanced validation set representative of the model's end goal. By choosing  $s_2$  that achieves the minimum loss, we can push the global model towards a more balanced and fairer performance. In our experiment, we initialize  $s_2$  to be 3, and in each round, check the set  $[s_2, s_2 + 0.5, s_2 - 0.5, s_2 - 5, s_2 + 5]$  for the optimal  $s_2$ .

The FL training implementing *FedVal* will follow the standard communication round. At the start of each communication round, a subset of clients is selected, and current global weight parameters will be broadcast to these selected clients. Then, these selected clients will perform local training and return their updated models to the server. The server validates all clients over a validation dataset and scores them as demonstrated in Equation 2. Lastly, the model parameters will be aggregated using the scores of each client as weight according to Equation 5. The complete *FedVal* algorithm can be seen in



Algorithm 1 in Appendix B.

$$\theta_g^{t+1} = \theta_g^t + \sum_d \frac{S_d(\theta_d^{t+1})\Delta\theta_d^{t+1}}{\sum_d S_d(\theta_d^{t+1})} \quad (5)$$

The salient dependency of *FedVal* on mean values for client updates is worth noting, as it could theoretically be prone to colluding attacks by multiple malicious clients as outlined in [31, 38]. Under such circumstances, one client could manipulate the mean, enabling other clients to subtly introduce smaller attacks. However, in the classification problem-focused experimental setting of this study, this potential vulnerability is mostly irrelevant due to the inherent limit on classification loss functions. The potential vulnerability of *FedVals* in regression problems remains unexplored, which constitutes an area for future work. A viable approach to this issue might involve constraining the loss function within a reasonable range of values.

Time complexity is a crucial factor when considering defenses that validate clients on server-side validation data. It has been previously noted that exhaustively testing each client on a validation dataset can be computationally intensive [26]. However, it is important to note that the actual time complexity of validating each client can be reduced by utilizing parallel processing techniques, which results in a time complexity of  $O(wv)$  in parallel, where  $w$  represents the dimensions of the model and  $v$  represents the number of validation samples. A direct comparison for time complexity between *FedVal* and other aggregation methods is difficult, as these other methods mainly depend on number of clients and model dimensions. However, an estimate is that *FedVal* is on the upper middle end of computational load for robust aggregators. In Section 5.4, we will further investigate the time complexity by examining the number of validation samples required for *FedVal* and show that *FedVal* only requires a limited number of validation samples to be effective.

In conclusion, the design of *FedVal* will provide a relatively light-weight and flexible defense that is prepared for any potential scenario, while promoting fairness among labels and continuously striving to reduce the loss on the validation dataset during each training iteration, which can be shown in more detail from our experiments in Section 5.

## 4 Experimental Setup

This section will discuss the general setups we have used in our experimental study. Federated learning is simulated with the Virtual Client Engine (VCE) of the comprehensive Flower framework [2] enabling us to scale to a large number of clients within a single machine. Datasets and hyper-parameters are detailed below.

### 4.1 Datasets and Partitions

In our research, we employ three widely used datasets of diverse size and complexity, namely CIFAR-10 [22], FEMNIST [5], and PUMS ACSIncome [9].

CIFAR-10 is a renowned benchmark dataset in computer vision and machine learning, consisting of 60,000 32x32 color images equally divided among 10 distinct categories.

On the other hand, FEMNIST is an expanded version of the EMNIST dataset [8], partitioned among 3597 authors of the handwritten characters and digits. With over 800,000 28x28 grayscale images, FEMNIST is specifically designed to simulate a realistic federated setting where each author is represented by a client. In the FEMNIST dataset, there is a quite noticeable quantity skew across classes. Specifically, considering the division between numbers, capital letters, and lowercase letters, where lowercase letters and capital letters are underrepresented. Each number has approximately 4 times more elements than the letters, and capital letters are more common than lowercase letters, as reported in [8].

Additionally, the PUMS ACSIncome dataset [9] was also utilized. The ACSIncome dataset, based on the Public Use Microdata Sample (PUMS) from the American Community Survey (ACS), provides income-related data on individuals, making it an excellent resource for studying aspects such as demographic bias and fairness in ML applications [12, 15].

This work has employed several general federated settings for evaluations. The first setting utilized 40 clients for the CIFAR-10, with each client having a dataset of 1250 images. The experiment was conducted over 60 rounds, and 10 clients were selected at random in each round, using a fixed set seed to ensure a consistent selection of data across experiments. This pseudo-random selection of clients ensured that the same data was used for comparison purposes in each experiment.

The second setting is on the FEMNIST experiments, which were performed using the 3597 authors as clients. Each client has a varying number of samples, with an average of around 225 samples per client. The FEMNIST experiment was run over 200 rounds, and in each round, 30 clients were randomly selected in the same pseudo-random fashion as before to ensure consistent data selection for comparisons. For experiments displayed in bar graphs with FEMNIST, average result over five rounds are displayed.

The third setting involves experiments with the ACSIncome dataset. Here, we specifically investigated the recall rate for minority groups as a measure of fairness. We utilized 40 clients for these experiments, which were conducted over 30 rounds with 15 clients selected for each round. Client selection was done in similar pseudo-random fashion as for previous settings and average result over multiple rounds were presented.

In all experimental conditions, we conduct a series of tests, opting to feature the results with the lowest performance for each respective algorithm in our study, unless otherwise stated.



Our intent behind this approach is to portray precisely the influence of heterogeneity issues and poisoning attacks on a federated training system. By focusing on these results, we're able to highlight potential instabilities that would remain undetected if we merely displayed averages.

To study the impact of heterogeneity, a range of techniques have been utilized to manipulate client distribution in the CIFAR-10 dataset. We follow the latent Dirichlet distribution (LDA) implemented by [23, 45, 16] where both label distribution and quantity distribution is determined by input parameter  $\alpha$ . The level of heterogeneity is governed by the parameter  $\alpha$ . As  $\alpha \rightarrow \infty$ , partitions become more uniform (IID), and as  $\alpha \rightarrow 0$ , partitions tend to be more heterogeneous. Our experimental evaluation considers both  $\alpha = 0.4$  and  $\alpha = 1000$  for the non-IID and IID cases.

Additionally, this paper introduces a new method to examine fairness in the federated setting. This method artificially creates a situation where a less common type of client contains some vital data for the model.

On the other hand, in the FEMNIST dataset, heterogeneity occurs naturally due to the unique writing styles of each author, as well as differences in label distribution and sample quantity between authors.

## 4.2 Model Architecture and Training Details

This work employs a similar model for both CIFAR-10 and FEMNIST. We implement a widely adopted CNN neural network for the image classification tasks for both datasets. The model used for both datasets is a convolutional neural network with 6 convolutional layers, and the kernel size used is 3x3 pixels. For both datasets, the models are trained with SGD, and the number of local client epochs is set to 10. Both models employ a learning rate of 0.005.

In addition to employing a similar model for CIFAR-10 and FEMNIST, a different model architecture was utilized for the PUMS ACSIncome dataset. The model consists of a sequential neural network with four dense layers. It was compiled using the Adam optimizer with a learning rate of 0.0001 and trained with the binary cross-entropy loss function.

Also, we set the *FedVal* specific hyper-parameters  $s_1$  to be 3 for all datasets. Since  $s_2$  is an adaptive hyper-parameter, it starts from the 3 and is chosen adaptively in each round as explained in Section 3.2. As previously mentioned, the summation in Equation 2 can be extended with other dimensions that matter. Our implementation adds overall accumulated average loss to the summation where  $s_{1,avg}$  is set to 5 while  $s_2$  is redundant since the bias reducer term will always be 1 due to dividing overall loss by itself. For all experiments in the paper, the  $C$  hyper-parameter is set to 3. The implemented score function is demonstrated in Equation 6.

$$S(\theta_d^t) = \sum_{k=1}^K \left( \max\left(1, \left(\frac{\bar{L}_k}{\bar{L}_{avg}}\right)^{s_2}\right) * \frac{s_{1,k} * div_{k,d}}{MAD_k} + Cs_{1,k} \right) + \frac{s_{1,avg} * div_{avg,d}}{MAD_{avg}} + Cs_{1,avg} \quad (6)$$

In addition to this, in our experiments with ACSIncome, the scoring dimensions were additionally extended to consider recall rates across certain groups. For this dimension, the  $s_2$  hyperparameter was statically set to 30.

## 4.3 Poisoning Attacks Methods

We employ two separate poisoning attacks to comprehensively cover the possible types of attacks that might occur in FL systems. The first attack is a *targeted data poisoning attack*, adapted from the method presented by Sun et al. [39], which introduces small changes to the model that affect only a limited subset of its tasks. The second attack is an *untargeted model poisoning attack*, adapted from the gradient ascent attack presented by Shejwalkar et al. [38], which seeks to reduce the global accuracy of the model while avoiding detection by the defense mechanisms in place. These types of attacks have been widely studied in the field of poisoning attacks, as documented in recent research works [38, 37, 25, 19, 39].

The underlying principle of the *targeted data poisoning attack* is to manipulate a small subset of the model, making it challenging for defenses to detect the attack. This increases the potential impact on the targeted functionality while minimizing the risk of detection. As a data poisoning attack, it also expands the adversary's reach, allowing for more malicious actors to participate.

On the other hand, the *untargeted model poisoning attack* is designed to cause damage to the model without targeting a specific area of its functionality. With full access to the model during training rounds, the adversary can launch a more powerful attack by finding the optimal malicious direction that the defense will not identify as malicious. This attack is an untargeted *projected gradient ascent attack (PGA)* and aims to impact the model wherever it is feasible.

## 4.4 Differential Privacy Methods

The impact of adding differential privacy on defenses against poisoning attacks in federated learning systems has been the subject of extensive research and discussions [18, 39]. In this work, we contribute to this body of research by incorporating differential privacy techniques into the federated learning system. We adopt the approach proposed by Andrew et al. [1] by adaptively binding norms and adding noise to each update vector.

Our aim is to investigate the behavior of various defenses under these conditions and assess the reduction in the impact of many poisoning attacks [39, 31]. Our work expands

upon previous research by exploring the interplay between differential privacy and defenses against poisoning attacks in federated learning systems.

Due to the adaptive nature of the approach proposed by Andrew et al., a quantitative evaluation of privacy is not feasible and has been left out of this paper.

## 4.5 Baselines and comparison

We implement various state-of-the-art solutions as baselines from a wide range of prospective in our experiments for comprehensive comparisons.

*LFR* [13] is a poisoning defense that removes a specified number of clients in each round based on their accumulated loss on a server-side validation dataset. In our experiments, 40% of clients were removed each round, as determined by the mathematical experiment in Section A and the number of malicious clients used. LFR serves as a benchmark algorithm for validation-loss-based defenses. It is worth mentioning that following the original paper, the FEMNIST dataset experiments with LFR use a heterogeneous validation dataset with the same label distribution as the training dataset.

*Multi-Krum* [3], an outlier-based defense method, is a defense that removes a specified number of client updates deemed as outliers, based on the Euclidean distance between updates. As Multi-Krum (and other outlier-based defenses) struggle to detect model poisoning attacks [20, 38], the number of clients to remove for multi-Krum is chosen to be 50%.

We also used *norm-bounding* [39] as a benchmark defense to mitigate the effects of poisoning attacks. This defense is integrated into our differential privacy solution [1] and was evaluated both in combination with other defenses and as a standalone measure against the attacks. We use a target quantile of 50% of norms to be bound as in [1].

Lastly, we implement *FedProx* [28]. FedProx is a widely-adopted method targeting to solve the data heterogeneity problem by preventing model drift from client updates. FedProx is implemented to investigate shortcomings and potential issues with such solutions. Experiments with FedProx follow the same protocol suggested in the original paper, which tests the hyper-parameter  $\mu$  from the range of [1, 5].

As mentioned before, fairness and reducing group bias in federated learning is a field not well researched, and existing solutions often use methods that might not be feasible due to privacy and security issues [18]. This makes it difficult to find baseline solutions targeted at promoting fairness.

*FedVal* distinguishes itself from other existing methodologies in several crucial aspects. While multi-Krum utilizes outlier detection, *FedVal* uncovers malicious clients by analysing accumulated loss on a server-side validation dataset, a strategy similar to that employed by LFR. Despite these similarities, *FedVal* and LFR differ significantly in their operations. LFR relies heavily on pre-existing knowledge of the quantity of malicious nodes, as discussed in Sections 2.5 and A, can result

in complications. In contrast, *FedVal* diverges in its approach by not needing prior knowledge of number of malicious nodes. *FedVal* leverages analyses performed on the server-side validation set to maintain model balance and identify oscillations across specific labels or dimensions. Unlike norm-bounding, which merely aims to minimize the fallout of an attack, *FedVal* seeks to eradicate malicious client updates altogether. Furthermore, for heterogeneity issues, it stands apart from FedProx by striving to craft a model that excels by its very design, rather than simply aligning it with the majority client model.

## 5 Experiments

In this section, we will present the results of our experiments and compare them with the chosen baseline solutions. We will also discuss the significance of the results and provide an extensive analysis of the performance of different algorithms. To begin with, we will present the results of the baselines with different algorithms in Section 5.1. We will then move on to the results from poisoning attacks in Section 5.2. Additionally, we will showcase an analysis of the limitations and potentials of different algorithms in Section A. Finally, we will demonstrate some results regarding fairness in the non-IID settings in Section 5.5.

### 5.1 Baselines Results

Figure 4 demonstrates the baseline results for all the algorithms we use. Figure 4a illustrates the IID test accuracy on CIFAR-10, and the results are what one would expect. Solutions that remove a set number of clients each round, such as LFR and multi-Krum have a very slight decrease in accuracy, while all other methods perform similarly. However, the results on the naturally heterogeneous dataset FEMNIST show more diverse results. All robust aggregators converge a bit slower than FedAvg and FedProx, which are methods not targeting for defense, but we can see that *FedVal* converges to higher accuracy than other solutions, while multi-Krum and LFR converge to lower accuracy. It is worth noting that the results from FEMNIST are from IID test data, which means that, as opposed to the train data, the test data has the same amount of elements for each label.

### 5.2 Poisoning attacks

This section aims to evaluate the resilience of *FedVal* against other baseline algorithms with the presence of various forms of poisoning attacks in diverse settings, including systems with differential privacy and different types of heterogeneity.

To begin with, in Figure 5, we illustrate the accuracy of different baseline algorithms when the system is under a static PGA model poisoning attack by 10% of the present clients in the FEMNIST dataset. In our experiments, we have considered two different scenarios - one where noise is added

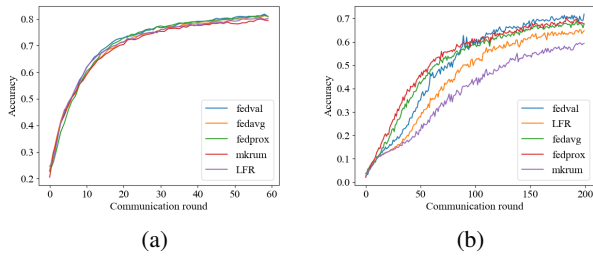


Figure 4: Baseline test accuracy for CIFAR-10 and FEMNIST for different baseline methods. (a) accuracy for CIFAR-10 with IID distribution; (b) accuracy for the naturally heterogeneous FEMNIST dataset.




Figure 5: Comparison of accuracy for various aggregation methods applied to the naturally heterogeneous FEMNIST dataset under PGA attack conditions, with and without the application of the differential privacy techniques of binding norms and noise addition. (a) PGA attack scenario with differential privacy techniques, (b) PGA attack scenario without differential privacy techniques.

and norms are cut by the value approximating 50% of norms being cut, illustrated in Figure 5a, and another where there is no noise or norm-bounding, illustrated in Figure 5b.

Results demonstrate that *FedVal* performs the best in both cases. This can be attributed to its natural ability to handle heterogeneous data and the ability not to ignore any good client models due to any set limit on the potential number of malicious clients. The results indicate that both *FedVal* and LFR are able to remove all malicious updates that would have any noticeable impact, but LFR differs in that it also removes a significant number of benign updates, since 40% of clients in an LFR update are always removed. Interestingly, in the norm-bounding scenario depicted in Figure 5a, multi-Krum performs worse than simply averaging (*FedAvg*). This is likely due to the fact that the norms of the malicious updates are scaled in such a way that they are not deemed malicious by the aggregator. As a result, since half of all clients are removed in our multi-Krum implementation, this leads to a larger percentage of remaining clients being malicious.

In addition, in Figure 6, we present results from experiments that are similar to those shown in Figure 5, using the CIFAR-10 dataset. In this experiment, we have introduced




Figure 6: Comparison of accuracy for various aggregation methods applied to the CIFAR-10 dataset, targeted under PGA attack conditions, with simulated heterogeneity in client data distribution using LDA ( $\alpha = 0.4$ ). (a) PGA attack scenario with differential privacy techniques (binding norms and noise addition), (b) PGA attack scenario without differential privacy techniques.

artificial heterogeneity by utilizing the LDA with  $\alpha = 0.4$  as explained in Section 4.1. Similar to the previous experiment, we compare different baseline algorithms and *FedVal* under the condition where 10% of present clients are malicious, with the noise and norm-bounding setting on Figure 6a, and without noise and norm-bounding Figure 6b. From these experiments, we can see that *FedVal* and LFR perform similarly under the noise and norm-bounding conditions, while *FedVal* slightly outperforms LFR without noise and binding norms.

It is important to note that *FedVal* is designed with fairness across labels as a goal, with the aim of improving the accuracy of underrepresented labels. This feature increases the accuracy of the test dataset performed with the FEMNIST dataset since there are underrepresented labels present, such that some labels are present in less quantity. However, with the use of the LDA on CIFAR-10, the underlying data distribution remains homogeneous, but the client data distribution is skewed. This is likely the reason for the difference in results between Figure 5 and 6. On the other hand, we can observe the same trends for CIFAR-10 in Figure 6 as we did for FEMNIST in Figure 5. The results presented in both figures demonstrate the effectiveness of *FedVal* in improving the performance of FL under various conditions, including label distribution skew and with the presence of poisoning attacks.

Furthermore, we conduct experiments with backdoor poisoning attacks on the CIFAR-10 dataset as shown in Figure 7. In the experiment, we manipulate the label of the data by changing the labels of ‘horse’ images to ‘deer’. Specifically, in Figure 7a, we change 10% of the clients’ data, and in Figure 7b, we change 20% of the clients’ data. The accuracy of the backdoor attack in this experiment is measured by how often the model incorrectly predicts a ‘horse’ image as a ‘deer’, so the lower accuracy indicates less successful attacks. By investigating the results, we can see that the LFR and *FedVal* aggregation methods outperform other methods in protecting the model against the backdoor attack. Also, the results show




Figure 7: Backdoor performance over rounds for the backdoor task of miss-predicting horses into being deer. Malicious clients perform data poisoning attacks by swapping the labels of ‘horses’ into ‘deer’. The backdoor accuracy is measured by how often the model incorrectly predicts a ‘horse’ image as a ‘deer’. (a) attack accuracy with 10% malicious clients; (b) attack accuracy with 20% malicious clients.

that both LFR and *FedVal* are able to provide almost full protection for the model, as compared with the no-attack case. This highlights the effectiveness of our method in mitigating the impact of malicious actors attempting to manipulate machine learning models using backdoor attacks.

### 5.3 Number of malicious clients

To ensure robustness, it is important to consider the potential presence of a percentage of malicious clients in a round over a large number of rounds. More details regarding the theoretical scenarios are explained in Appendix A. Following the discussion in Section 2.5 and Appendix A, we decided to investigate the situation where the number of malicious clients exceeds the limit set by robust aggregators in order to understand the capabilities of our method *FedVal* fully. For other aggregators, as one might expect, the model’s performance deteriorates significantly in scenarios where the number of malicious clients exceeds the limit set by a robust aggregator.

We conducted experiments with both 40% and 80% of total clients as malicious, illustrated in Figure 8, and the results were quite striking. Despite the high percentage of malicious clients, *FedVal* was able to converge without any major difficulties as opposed to other methods. This highlights the strength of *FedVal* as both a robust but also adaptive poisoning defense that is prepared for any type of situation that may occur in a federated learning system.

Furthermore, the results of these experiments also demonstrate the robustness of *FedVal* in the face of a high number of malicious clients. In the standard federated learning system using aggregation methods such as FedAvg, the presence of even a small number of malicious clients can cause serious issues with the convergence and accuracy of the model, as we can see from Figures 5, 6, and 7. However, our method can mitigate the impact of malicious clients and ensure the integrity of the model in various scenarios.

Overall, the experimentation in Figure 8 has shown that




Figure 8: Performance of various defenses under stress test conditions with a large number of malicious clients conducting poisoning attacks. (a) accuracy with 40% of clients being malicious, (b) accuracy with 80% of clients being malicious.




Figure 9: Performance of *FedVal* under poisoning attacks with different numbers of validation elements used for the analysis on CIFAR-10. (a) accuracy for the backdoor task of the malicious clients, (b) overall accuracy.

*FedVal* is an effective and reliable method for defending against poisoning attacks in federated learning systems, even in situations where the number of malicious clients is higher than what we may expect there to be.

### 5.4 Time complexity

We discussed the time complexity and computational load of running *FedVal* previously in Section 5.4. To further investigate this issue, we have conducted an experimental study to determine the minimum number of validation elements required for the algorithm to maintain robustness. We present the experiment results in Figure 9, which are conducted using the CIFAR-10 dataset and various attacks. We varied the number of validation elements used in the defense and evaluated the effectiveness of the system under these conditions.

The experiment shows that a relatively small number of validation elements is sufficient to ensure robustness. For the untargeted attack illustrated in Figure 9b, we obtained the unexpected result that checking the loss on a single element for each class is enough to achieve robustness. For the targeted backdoor attack, illustrated in Figure 9a, 75 total validation elements are required to achieve robustness, which is 7.5 elements per class.

Based on these findings, it can be inferred that using 10 validation elements per class provides a considerable degree






Figure 10: Test on CIFAR-10 non-IID. Two labels are missing in a majority of the clients. The clients who have the missing labels are IID.

of robustness, given the experimental conditions we’ve explored. We recommend maintaining a validation dataset that is minimalistic yet comprehensive in scope. This provides important information, as it suggests that the *FedVal* algorithm is computationally efficient and can be implemented with minimal computational resources.

### 5.5 Fairness

Continuing our experimental study, we will now focus on investigating the impact of heterogeneous data in federated learning. Specifically, we will demonstrate the importance of handling fairness in federated learning, especially in scenarios where certain classes are underrepresented in some clients.

To illustrate this point, we have designed a simple scenario using the CIFAR-10 dataset, where 2 of the 10 classes are missing in the 70% of the participating clients. Figure 10 presents the results of this scenario, where classes ‘4’ and ‘5’ are the missing classes. The results indicate that in this scenario, the accuracy for classes ‘4’ and ‘5’ hover around zero for a majority of the tested solutions. Solutions managing to handle the scenario are our proposed solution, *FedVal* and interestingly also to a certain degree with LFR. This is likely attributed to the fact that LFR would favor the clients with the complete label distribution and remove a few skewed clients each round. This is not sufficient to handle the setting as the clients with the missing labels aggregation needs to be scaled up for the missing labels to be properly addressed, as we can see happening with *FedVal*.

We have also implemented FedProx for this scenario, which is a widely-adopted solution for non-IID data distribution in federated learning to prevent client drift. As mentioned, preventing client drift is beneficial in situations where the optimized global model is close to the type of model that a majority of the clients would have. However, in scenarios where some labels are only present in a small subset of all clients, solutions like FedProx falls short, as we can see from the results presented in Figure 10. The results where FedProx is evaluated in combination with *FedVal* even indicate that FedProx increases the difficulty of extracting the information that is missing in the majority of clients, likely due to the






Figure 11: Fairness experiment on the naturally heterogeneous FEMNIST and PUMS ACSIncomes dataset. (a) average accuracy for the final 5 rounds on each class group on FEMNIST, (b) mean absolute deviation across all labels on FEMNIST, (c) average recall for final 10 rounds on different groups in ACSIncome.

penalty term, which prevents the clients from moving the model too much, resulting in all clients having more similar models as the ones with missing labels. On the other hand, by utilizing re-weighting schemes, *FedVal* is able to salvage this.

These results demonstrate the importance and significance of fairness in the federated learning system, especially when many clients have underrepresented or missing classes. The results also highlight the catastrophic forgetting in the federated setting [29, 17], which can be a crucial issue to consider, especially in practical and real-life scenarios, where some labels likely will be less common and not present in all clients.

Relating to the experiment in Figure 10, we also investigate the accuracy of the model thoroughly across all labels in the FEMNIST dataset. As mentioned in Section 4.1, there is a quite noticeable quantity skew across classes. Specifically, considering the division between numbers, capital letters, and lowercase letters, each class of Numbers has approximately 4 times more elements than the Letters, and Capital Letters

are more common than Lowercase Letters. With this type of distribution, we can expect the model to quickly learn how to correctly classify Numbers, and struggle with correctly classifying capital and lowercase letters in comparison.

In Figure 11, we illustrate the performance of different aggregation methods on different label classes. Figure 11a illustrates the accuracy division across label groups, and Figure 11b illustrates the mean absolute deviation across all labels. From Figure 11a, we can see that by simply averaging through FedAvg or using methods that prevent client drift, such as FedProx, a large bias is created toward the more common groups. The models created by these methods tend to bias toward the more common groups (the Numbers group in our case) and neglect the less common groups, which are the Lowercase Letters. However, most of the robust aggregators happen to create less bias in the model. The most noticeable difference is with *FedVal*, which manages to almost double the accuracy of the less common Lowercase Letters, more exact from 32% with FedAvg to 53% with *FedVal*, while still managing to provide similar accuracy for the more common classes. On the other hand, other methods struggle with creating a model that predicts the Lowercase Letters correctly.

In Figure 11c, we have conducted an analysis utilizing the PUMS ACSIncome dataset [9]. This particular dataset is frequently employed to examine fairness issues [12, 15]. The focus of our study was to scrutinize the recall rate (true positive) for minority groups. Historically, ML models have demonstrated a bias towards groups with fewer data points, resulting in a lower rate of accurate positive predictions.

In our experiment, we extended the dimensions of *FedVal* to encompass discrepancies in recall rates, aiming to equalize recall across all classes. This experiment's findings underscore *FedVal*'s dynamic capacity to adjust a model where necessary. By extending the algorithm to include recall analysis, we observed a significant improvement in recall rates between the "Indigenous of North America" and "Other" groups. Notably, these are groups where other methods have previously struggled to generate accurate positive predictions based on true labels.

In conclusion, to improve the accuracy of the model across all labels and ensuring we have a balanced model, we can consider using more advanced aggregation methods such as *FedVal*. Overall, it is important to consider the dataset distribution and potential biases when training and evaluating FL models to ensure that the model is accurate and fair across all classes and demographic groups.

## 5.6 Summary

Many existing solutions tend to overlook other critical problems in Federated Learning, often exacerbating these issues or introducing new security and privacy challenges. In contrast, our proposed method, *FedVal*, has shown resilience against both model and data poisoning attacks. It even thrives when a substantial majority of clients are malicious. For exam-

ple, *FedVal* manages to converge even with 80% of clients performing model poisoning attacks. This is a significant improvement over methods like multi-Krum and LFR, which fail to converge when faced with merely 40% malicious clients. Furthermore, when faced with 20% of clients deploying backdoor data poisoning attacks, both *FedVal* and LFR maintain their performance levels. In the same scenario, multi-Krum's performance drops drastically, even falling behind the no-defense FedAvg aggregator. In addressing data distribution skews, *FedVal* outshines other methods, salvaging problematic situations where others struggle, as illustrated in the CIFAR-10 experiment in Figure 10. Moreover, our fairness experiment using the ACSIncome dataset showed that *FedVal* significantly improves the recall of underrepresented groups from 19% with FedAvg to 50%, an increase of over 30%. FedProx, which is specifically designed to handle heterogeneity, performed similarly to FedAvg's simple averaging. Other robust aggregators showed some promise, but still fell short in performance compared to *FedVal*.

In conclusion, *FedVal* consistently outperforms existing methods in terms of robustness to attacks, promoting fairness, and adapting to data distribution skews. Our work highlights the necessity of considering multiple aspects of Federated Learning, from security to fairness, to ensure the robustness and reliability of the learning system.

## 6 Conclusion

In this work, we provide a robust solution, *FedVal*, which aims to solve multiple problems by analyzing and utilizing the clients' learning. We propose to do this by using a small server-side validation dataset to assess client updates and determine the optimal aggregation weights considering both robustness and fairness. This technique involves comparing average client performance with client performance over a range of dimensions.

The preceding sections of this paper have underscored the potential of the *FedVal* algorithm in the realm of FL. However, we acknowledge the necessity of extending our investigation to assess its applicability on more general regression problems. The breadth of unexplored challenges within FL suggests a wealth of potential benefits from utilizing client analysis through server-side validation data. It is therefore interesting to delve deeper into these areas, as they could provide strong justification for the broader adoption of *FedVal* and similar algorithms.

Moreover, we also perceive the potential of exploring more sophisticated aggregation methods, exploiting the data acquired through validation. The development of such methods could push the boundaries of what is currently achievable, thereby offering a new dimension of innovation. As we strive for continual improvement in FL, our future research endeavors will focus on unearthing and maximizing the potential inherent in these investigative avenues.

## Availability

Code available at: <https://github.com/viktorvaladi/FedVal>

## References

- [1] Galen Andrew et al. “Differentially private learning with adaptive clipping”. In: *Advances in Neural Information Processing Systems* 34 (2021), pp. 17455–17466.
- [2] Daniel J Beutel et al. “Flower: A friendly federated learning research framework”. In: *arXiv preprint arXiv:2007.14390* (2020).
- [3] Peva Blanchard et al. “Machine learning with adversaries: Byzantine tolerant gradient descent”. In: *Advances in Neural Information Processing Systems* 30 (2017).
- [4] Keith Bonawitz et al. “Practical secure aggregation for federated learning on user-held data”. In: *arXiv preprint arXiv:1611.04482* (2016).
- [5] Sebastian Caldas et al. “Leaf: A benchmark for federated settings”. In: *arXiv preprint arXiv:1812.01097* (2018).
- [6] Xiaoyu Cao et al. “Fltrust: Byzantine-robust federated learning via trust bootstrapping”. In: *arXiv preprint arXiv:2012.13995* (2020).
- [7] Simon Caton and Christian Haas. “Fairness in machine learning: A survey”. In: *arXiv preprint arXiv:2010.04053* (2020).
- [8] Gregory Cohen et al. “EMNIST: Extending MNIST to handwritten letters”. In: *2017 international joint conference on neural networks (IJCNN)*. IEEE. 2017, pp. 2921–2926.
- [9] Frances Ding et al. “Retiring adult: New datasets for fair machine learning”. In: *Advances in neural information processing systems* 34 (2021), pp. 6478–6490.
- [10] Qi Dou et al. “Federated deep learning for detecting COVID-19 lung abnormalities in CT: a privacy-preserving multinational validation study”. In: *NPJ digital medicine* 4.1 (2021), pp. 1–11.
- [11] Wei Du et al. “Fairness-aware agnostic federated learning”. In: *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*. SIAM. 2021, pp. 181–189.
- [12] Yahya H Ezzeldin et al. “Fairfed: Enabling group fairness in federated learning”. In: *arXiv preprint arXiv:2110.00857* (2021).
- [13] Minghong Fang et al. “Local model poisoning attacks to {Byzantine-Robust} federated learning”. In: *29th USENIX Security Symposium (USENIX Security 20)*. 2020, pp. 1605–1622.
- [14] Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. “Mitigating sybils in federated learning poisoning”. In: *arXiv preprint arXiv:1808.04866* (2018).
- [15] Ira Globus-Harris, Michael Kearns, and Aaron Roth. “An algorithmic framework for bias bounties”. In: *2022 ACM Conference on Fairness, Accountability, and Transparency*. 2022, pp. 1106–1124.
- [16] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. “Measuring the effects of non-identical data distribution for federated visual classification”. In: *arXiv preprint arXiv:1909.06335* (2019).
- [17] Wenke Huang, Mang Ye, and Bo Du. “Learn From Others and Be Yourself in Heterogeneous Federated Learning”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022, pp. 10143–10153.
- [18] Peter Kairouz et al. “Advances and open problems in federated learning”. In: *Foundations and Trends® in Machine Learning* 14.1–2 (2021), pp. 1–210.
- [19] Sai Praneeth Karimireddy et al. “Scaffold: Stochastic controlled averaging for federated learning”. In: *International Conference on Machine Learning*. PMLR. 2020, pp. 5132–5143.
- [20] Harsh Kasyap and Somanath Tripathy. “Hidden Vulnerabilities in Cosine Similarity based Poisoning Defense”. In: *2022 56th Annual Conference on Information Sciences and Systems (CISS)*. IEEE. 2022, pp. 263–268.
- [21] Emmanouil Krasanakis et al. “Adaptive sensitive reweighting to mitigate bias in fairness-aware classification”. In: *Proceedings of the 2018 world wide web conference*. 2018, pp. 853–862.
- [22] Alex Krizhevsky, Geoffrey Hinton, et al. “Learning multiple layers of features from tiny images”. In: (2009).
- [23] Qinbin Li et al. “Federated learning on non-iid data silos: An experimental study”. In: *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE. 2022, pp. 965–978.
- [24] Shenghui Li, Edith Ngai, and Thiemo Voigt. “Byzantine-Robust Aggregation in Federated Learning Empowered Industrial IoT”. In: *IEEE Transactions on Industrial Informatics* (2021).
- [25] Suyi Li et al. “Abnormal client behavior detection in federated learning”. In: *arXiv preprint arXiv:1910.09933* (2019).
- [26] Suyi Li et al. “Learning to detect malicious clients for robust federated learning”. In: *arXiv preprint arXiv:2002.00211* (2020).

- [27] Tian Li et al. “Federated learning: Challenges, methods, and future directions”. In: *IEEE signal processing magazine* 37.3 (2020), pp. 50–60.
- [28] Tian Li et al. “Federated optimization in heterogeneous networks”. In: *Proceedings of Machine Learning and Systems 2* (2020), pp. 429–450.
- [29] Michael McCloskey and Neal J Cohen. “Catastrophic interference in connectionist networks: The sequential learning problem”. In: *Psychology of learning and motivation*. Vol. 24. Elsevier, 1989, pp. 109–165.
- [30] Brendan McMahan and Daniel Ramage. “Federated learning: Collaborative machine learning without centralized training data”. In: *Google Research Blog* 3 (2017).
- [31] Ashwinee Panda et al. “SparseFed: Mitigating Model Poisoning Attacks in Federated Learning with Sparsification”. In: *International Conference on Artificial Intelligence and Statistics*. PMLR. 2022, pp. 7587–7624.
- [32] Titouan Parcollet et al. “ZeroFL: Efficient on-device training for federated learning with local sparsity”. In: *International Conference on Learning Representations (ICLR)*. 2022.
- [33] Dawid Połap and Marcin Woźniak. “Meta-heuristic as manager in federated learning approaches for image processing purposes”. In: *Applied Soft Computing* 113 (2021), p. 107872.
- [34] Xinchu Qiu et al. “A first look into the carbon footprint of federated learning”. In: *arXiv preprint arXiv:2102.07627* (2021).
- [35] Nicola Rieke et al. “The future of digital health with federated learning”. In: *NPJ digital medicine* 3.1 (2020), pp. 1–7.
- [36] Yuji Roh, Geon Heo, and Steven Euijong Whang. “A survey on data collection for machine learning: a big data-ai integration perspective”. In: *IEEE Transactions on Knowledge and Data Engineering* 33.4 (2019), pp. 1328–1347.
- [37] Virat Shejwalkar and Amir Houmansadr. “Manipulating the byzantine: Optimizing model poisoning attacks and defenses for federated learning”. In: *NDSS*. 2021.
- [38] Virat Shejwalkar et al. “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. 2022, pp. 1354–1371.
- [39] Ziteng Sun et al. “Can you really backdoor federated learning?” In: *arXiv preprint arXiv:1911.07963* (2019).
- [40] Aleksei Triastcyn and Boi Faltings. “Federated learning with bayesian differential privacy”. In: *2019 IEEE International Conference on Big Data (Big Data)*. IEEE. 2019, pp. 2587–2596.
- [41] Nguyen Truong et al. “Privacy preservation in federated learning: An insightful survey from the GDPR perspective”. In: *Computers & Security* (2021), p. 102402.
- [42] Paul Voigt and Axel Von dem Bussche. “The eu general data protection regulation (gdpr)”. In: *A Practical Guide, 1st Ed., Cham: Springer International Publishing* 10.3152676 (2017), pp. 10–5555.
- [43] Kang Wei et al. “Federated learning with differential privacy: Algorithms and performance analysis”. In: *IEEE Transactions on Information Forensics and Security* 15 (2020), pp. 3454–3469.
- [44] Dong Yin et al. “Byzantine-robust distributed learning: Towards optimal statistical rates”. In: *International Conference on Machine Learning*. PMLR. 2018, pp. 5650–5659.
- [45] Mikhail Yurochkin et al. “Bayesian nonparametric federated learning of neural networks”. In: *International Conference on Machine Learning*. PMLR. 2019, pp. 7252–7261.
- [46] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. “idlg: Improved deep leakage from gradients”. In: *arXiv preprint arXiv:2001.02610* (2020).
- [47] Yue Zhao et al. “Federated learning with non-iid data”. In: *arXiv preprint arXiv:1806.00582* (2018).
- [48] Ligeng Zhu, Zhijian Liu, and Song Han. “Deep leakage from gradients”. In: *Advances in neural information processing systems* 32 (2019).

## A Number of malicious clients

As mentioned in Section , in order to ensure robustness, it is important to consider the potential presence of a percentage of malicious clients in a round over a large number of rounds. Figure 12 demonstrates a theoretical scenario in which the number of malicious clients present in the system is 10% of the total clients. To counteract this, we have theoretically implemented protection measures that guarantee robustness when up to 40% of the clients selected for a round are malicious. The figure shows the probability that in at least one round, there will be more than 40% of malicious clients out of selected clients when 30 clients are selected in each round. The probability was calculated using the binomial distribution  $B(n, k, p)$  with values  $n = 30, k = x, p = 0.1$ . To get the probability of the event that 9 or more malicious clients are present, we sum the probabilities of  $k \geq 9$  up to 30

$$P(X) = \sum_{k=9}^{30} B(30, k, 0.1) \quad (7)$$






Figure 12: Probability to exceed 40% malicious clients selected in a round at least once when there are 10% total malicious clients in the system.

To get the probability that this event would happen at least once in  $x$  number of rounds, we inverse the probability of the event not happening for  $x$  consecutive rounds

$$1 - (1 - P(X))^x \quad (8)$$

which is the plot given in Figure 12 over communication rounds  $x$ . As we can see, after 25,000 rounds, we are very likely to have a "successful attack" where the bound is broken.

## B Algorithm

---

**Algorithm 1** FedVal:  $N$  is total number of clients.  $r$  is the ratio of selecting clients each round.  $D_d$  is the set of client data in client  $d$ .  $\theta_g^t$  is the global model parameter at round  $t$  and  $\theta_d^t$  is the local model parameter at client  $d$ .  $L$  be the loss function and  $\eta$  be the learning rate.

---

```

1: procedure FEDVAL( $s_1, s_2, ValidationData$ )
2:   Initialize model  $\theta_g^0$ 
3:   for  $t = 0, 1, \dots, T - 1$  do
4:     Randomly select set of  $rN$  devices  $\{S_t\}$ 
5:      $\theta_d^t = \text{LOCALTRAINING}(\theta_g^t, S_t)$ 
6:      $\theta_g^{t+1}, s_2 = \text{SCORECLIENTS}(\{\theta_d^t : d \in S_t\}, s_1, s_2)$ 
7:   end for
8:   function LOCALTRAINING( $\theta_g^t, S_t$ )
9:     for  $d \in S_t$  in parallel do
10:      Update model with client data  $D_d$ 
11:       $\theta_d^t \leftarrow \theta_g^t - \eta \nabla L(\theta_g^t; D_d)$ 
12:     end for
13:     return  $\{\theta_d^t : d \in S_t\}$ 
14:   end function
15:   function SCORECLIENTS( $\{\theta_d^t : d \in S_t\}, s_1, s_2$ )
16:     for  $d \in S_t$  do
17:       Calculate losses  $L_d = L(\theta_d^t, validationData)$ 
18:       Calculate diversions from mean per dimension
19:        $div_{k,d} = \bar{L}_K - L_{k,d}$ , for  $k = 1, \dots, K$ 
20:     end for
21:     Calculate mean absolute deviation
22:      $MAD_k = \frac{\sum_{m=1}^M |L_{m,k} - \sum_{i=1}^M L_{i,k}|}{rN}$ 
23:     for  $s_2, s_2 + 0.5, s_2 - 0.5, s_2 + 5, s_2 - 5$  in parallel do
24:       for  $d \in S_t$  do
25:          $S(\theta_d^t) = \text{Equation 2}$ 
26:       end for
27:       Aggregate by current  $s_2$  value
28:        $\theta_{s_2}^{t+1} = \theta_g^{t+1} + \sum_d \frac{\max(0, S(\theta_d^t)) * \theta_d^t}{\sum_d \max(0, S(\theta_d^t))}$ 
29:       Evaluate loss for  $s_2$  by loss function  $\mathcal{L}$ 
30:        $L_{\theta_{s_2}^t} = \mathcal{L}(\theta_{s_2}^t, ValidationData)$ 
31:     end for
32:     Choose  $s_2$  and model with min validation loss
33:      $\theta_g^{t+1}, s_2 \sim \min(L_{\theta_{s_2}^t})$ 
34:     return  $\theta_g^{t+1}, s_2$ 
35:   end function
36: end procedure

```

---