



HAL
open science

Towards Robust and Bias-free Federated Learning

Ousmane Touat, Sara Bouchenak

► **To cite this version:**

Ousmane Touat, Sara Bouchenak. Towards Robust and Bias-free Federated Learning. 3rd Workshop on Machine Learning and Systems (EuroMLSys '23), May 2023, Rome, Italy. pp.49-55, 10.1145/3578356.3592576 . hal-04102260

HAL Id: hal-04102260

<https://hal.science/hal-04102260>

Submitted on 23 May 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards Robust and Bias-free Federated Learning

OUSMANE TOUAT, INSA Lyon – LIRIS, France

SARA BOUCHENAK, INSA Lyon – LIRIS, France

Federated learning (FL) is an exciting machine learning approach where multiple devices collaboratively train a model without sharing their raw data. The FL system is vulnerable to the action of Byzantine clients sending arbitrary model updates, and the trained model may exhibit prediction bias towards specific groups. However, FL mechanisms tackling robustness and bias mitigation have contradicting objectives, motivating the question of building a FL system that comprehensively combines both objectives. In this paper, we first survey state-of-the-art approaches to robustness to Byzantine behavior and bias mitigation and analyze their respective objectives. Then, we conduct an empirical evaluation to illustrate the interplay between state-of-the-art FL robustness mechanisms and FL bias mitigation mechanisms. Specifically, we show that classical robust FL methods may inadvertently filter out benign FL clients that have statistically rare data, particularly for minority groups. Finally, we derive research directions for building more robust and bias-free FL systems.

CCS Concepts: • **Computing methodologies** → **Distributed artificial intelligence**; • **Social and professional topics** → **User characteristics**; • **Computer systems organization** → **Reliability**.

Additional Key Words and Phrases: Federated Learning, Bias Mitigation, Robustness, Byzantine Behavior

1 INTRODUCTION

Federated Learning (FL) is a distributed learning paradigm that allows multiple entities, denoted FL clients, to collaborate to learn a model with the help of a central node called the server, with privacy in mind, as the entities do not share their training data [11]. The deployment of FL in high-risk sectors such as healthcare has sparked concerns as the models may exhibit bias due to underrepresented minorities in data, leading to poorer performance for specific groups [5]. For example, there have been reports of image classification models used for diagnosing skin cancer showing less accurate results for the person of color [15]. Furthermore, some FL clients denoted Byzantine client, can significantly impact the FL model training by sending arbitrary updates to the server [13]. Addressing model bias and Byzantine clients in FL systems is critical, and several existing works are addressing such issues separately. For instance, FL bias mitigation methods formulate an optimization approach with bias constraints, requiring FL clients to send additional statistics on their local data [3, 5, 20]. On the other hand, robust FL systems usually rely on passive outlier detection mechanisms using robust aggregation or gradient clipping [13, 14, 22]. However, incorporating both goals into a single FL system is challenging, as the current objectives are contradictory. FL bias mitigation requires considering outlier clients for adding representativity in their local data, while robust FL systems may eliminate outlier clients, which can affect model bias for worse [8, 19]. Therefore,

Authors' addresses: Ousmane Touat, ousmane.touat@insa-lyon.fr, INSA Lyon – LIRIS, Lyon, France; Sara Bouchenak, sara.bouchenak@insa-lyon.fr, INSA Lyon – LIRIS, Lyon, France.

finding a balance between these two goals in a single FL system is complex and requires careful consideration to avoid compromising one for the other (see Figure 1 and Figure 2). This paper investigates the challenges of building a FL system that combines bias constraints and robustness guarantees. We empirically showcase where classical robust aggregation degrades model bias and hinders bias mitigation. We then discuss some research directions to build a FL system combining robustness and bias mitigation goals.

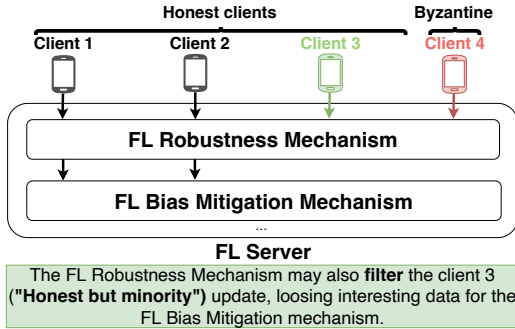


Fig. 1. Applying classical robust aggregators hinder FL bias mitigation

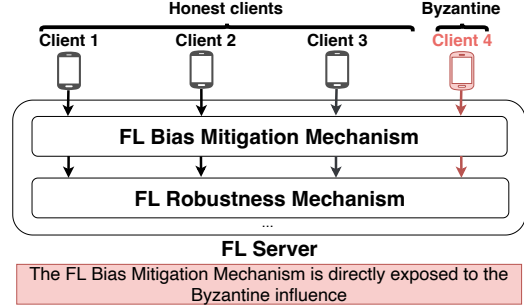


Fig. 2. Applying bias mitigation first in Byzantine setting does not work

2 BACKGROUND AND SYSTEM MODEL

In this section, we recall the principles of FL, demographic bias and the standard Byzantine attack model.

2.1 Federated Learning

In Federated Learning (FL), we have n clients aiming to optimize a shared objective function collectively managed by a central server. The training process progresses through communication rounds, in which the central server sends the current global ML model with parameters θ to the FL clients. In this setup, the i -th client has its local data denoted as \mathcal{D}_i and trains a local model denoted as θ_i . We denote the union of clients' local data \mathcal{D} . At the end of the round, the clients send their updated model to the server, which update the new global model by aggregating the clients' updates. The classical approach for FL aggregation is called FedAvg [11] and consists in averaging the clients' model parameters as follows :

$$\theta = \sum_{k=1}^{k=N} \frac{|\mathcal{D}_k|}{|\mathcal{D}|} \cdot \theta_k \quad (1)$$

2.2 Bias in Federated Learning

Bias refers to the property of an classifier model to disparately treat different groups based on sensitive attributes, denoted S . A sensitive attribute is binary as it partition the dataset into two groups. One widely

used metric for measuring bias with respect to sensitive attributes is statistical parity difference (SPD) [4]. With y the prediction value from the FL model and S the value of the sensitive attribute we define SPD_S as follows:

$$SPD_S = |Pr(y = 1|S = 1) - Pr(y = 1|S = 0)| \quad (2)$$

In this formulation, a value of zero for SPD_S indicates that the algorithm is statistically fair concerning the sensitive attribute, as it predicts positive outcomes with the same probability for both groups. Minimizing the value of such metric is relevant in cases where the classifier output must be independent from the sensitive attribute (e.g. algorithms with any societal impact which encounters issue of demographic discrimination).

2.3 Byzantine Attack Model in Federated Learning

The presence of malicious parties, known as Byzantine clients, denoted \mathcal{B} , can threaten the integrity of the FL process. The fraction of Byzantine clients is given by ρ , and the number of Byzantine clients denoted by q such that $q \leq \rho n$. Byzantine clients can deviate from the standard protocol and send arbitrary updates to the server, and they may even collude and have knowledge of the states of all other clients. The remaining $n - q$ clients who follow the protocol, denoted as honest clients, are represented by set \mathcal{H} . FL systems such as FedAvg are directly vulnerable to Byzantine attacks [13]. For example, we take a FL system of n clients, including one Byzantine client. We suppose that each honest client sends an update denoted as θ_i for $i \in \{1, \dots, n - 1\}$. The Byzantine client can craft an update $\theta_{byz} = |\mathcal{D}| \cdot \theta_{wt} - \frac{1}{|\mathcal{D}|} (\mathcal{D}_k \cdot \theta_i + \dots + \mathcal{D}_{n-1} \cdot \theta_{n-1})$ to make the server compute any global model θ_{wt} desired. This happens as the server, unaware of the presence of the Byzantine client, aggregates all received updates, producing the following global model θ_{wt} as the Byzantine model wanted. Such unwanted behavior harms the FL system greatly, making it unable to converge reliably, motivating the research on robust FL methods.

3 RELATED WORK

3.1 Robustness to Byzantine behavior in Federated Learning

The Byzantine fault tolerance problem is well-known in distributed systems, where malicious nodes may provide incorrect information, leading to incorrect results. In FL, this problem can arise when some participating clients behave maliciously and manipulate the aggregated model parameters. Several Byzantine resilient FL methods have been proposed in the literature using the statistical properties of the client updates. Krum [13], and by extension Multi-Krum, order the clients' updates using a score stating how close they are to its $n - f - 2$ nearest neighbors using Euclidean distances. RFA [14] computes an approximate value of the geometric median of all client updates using Weiszfeld's algorithm. TrimmedMeans [22] computes the global model by taking the mean value of each coordinate of the client updates, truncating extreme values. For model poisoning attack settings, NDC [18] adopts a norm-thresholding policy that limits the model's parameter norm by clipping it when it exceeds a fixed threshold. However, such robust aggregation techniques fail when used on non-IID data, cannot withstand attacks, which degrade the global model

performance [8]. For non-IID data, some defense mechanisms have been proposed to mitigate the impact of Byzantine behavior. Recently, new defense mechanisms have tackled the impact of Byzantine behavior in non-IID FL context. FedInv [24] performs model inversion on the clients' updates to generate dummy datasets and compare clients' Wasserstein distance on the dummy datasets distribution, empirically showing robustness in non-IID FL context. Recently Karimireddy and al. introduced a new bucketing technique where the server randomly partitions client updates into "buckets", providing some guarantee when using this technique alongside the previously mentioned robust aggregation technique in non-IID and Byzantine FL context [8]. Furthermore, Byzantine resilience can also be addressed by ensuring the integrity of the participating nodes using enclaves such as TEEs [6].

3.2 Bias Mitigation in Federated Learning

The issue of model bias is well-known in machine learning [12] and has also received prime attention in FL. In FL, we distinguish between approaches that mitigate model performance discrepancy between clients [10, 17] and ones that address model bias between demographic groups in FL, from where our work is positioned. We encounter two approaches to tackle this topic, namely client-side techniques and server-side techniques. The client-side technique aims to apply a local debiasing mechanism, such as data reweighting [1]. However, such a technique may fail on the global bias objective, especially where the FL setup is highly non-IID [5]. In server-side techniques, most FL methods seek to solve a global bias-constrained objective, requiring local clients to share information about the local statistic of sensitive attributes [3, 5, 20, 23]. AgnosticFair [20] uses local information from the client to optimize the data reweighting coefficient sent back to the client to seek its globally constrained objective problem. FairFL [23] train a client-selection policy function using multi-agent reinforcement learning to maximize a gain function based on how bias mitigation is performed on the global model. FairFed [5] linearly scalarize the inputted local model parameters based on their local bias discrepancy to the global bias metric. FCFL [3] also linearly scalarize the gradient updates from the clients to optimize both a bias constraint and consistency in the client's model performance.

3.3 Federated Learning with Bias and Robustness Guarantees

While the issues of demographic bias and robustness have been extensively explored in the literature, combining these two notions within a single system needs to be further examined. Ditto [10] confronts the notions of robustness against client-level fairness, showing that these two notions are competing, and implements a FL system that combines these properties using client model personalization. Recently, Singh and al. [16] built a robustness mechanism that minimizes the amount of honest "minority" clients detected as malicious ones. Assuming each client represents only one group, the server employs a decentralized micro aggregation algorithm to cluster the clients based on their published sensitive attribute, identifying malicious clients having updates too far from those cluster centroids.

4 PROBLEM ILLUSTRATION

While the robustness of FL systems is essential, using robust aggregators in FL may impair model bias, particularly in non-IID FL settings. The robustness methods may eliminate outlier clients with different data distribution, suppressing their contribution to the global model. In this section, we will further illustrate this problem and explore its implications for model bias in FL.

4.1 Impact of FL Robustness Methods on Bias

We compared the behavior of 4 robust aggregators that are Multi-Krum [13], RFA [14], TrimmedMeans [22] and NDC [18], on two tabular datasets, MEPS [2] and Adult [9], in non-IID FL settings. We also added FedAvg [22] as another FL baseline without robustness mechanism for comparison. On MEPS, an hospital expenditure dataset, the target attribute for the binary classification task is the utilization of medical facilities with the sensitive attributes of race and gender. The MEPS dataset is cut into a 4-client FL setup, showing the opposite trend of the three other clients that share similar data distribution, with the minority race becoming the majority. In the Adult dataset, the target attribute is income, and the considered sensitive attributes are gender and age. We use a 10-client FL setup generated with a Dirichlet function [7] on the distribution of the sensitive attribute and target feature combo to ensure heterogeneity between FL clients. We report the model bias for the baseline FedAvg method and the robust aggregators in Figure 3. We observe that the robust aggregators, including Multi-Krum, RFA, and NDC, increase the model bias compared to the FL baseline on both MEPS (Figure 3(c)) and Adult (Figure 3(a),3(b)). The robust aggregators' behavior in such experiments is concerning as it demonstrates a degradation in the model bias in a setup without any Byzantine attacks. We attribute such behavior to the fact that robust aggregators tend to exclude clients with different data distributions. We observe for example that in the 4-client setup of the MEPS experiment, the client that shows the reverse bias trend is systematically excluded, or have its influence reduced by the robustness techniques. In our case, it corresponds to clients showing opposite bias dispositions, which are necessary to offset the biases from the other FL clients, thus showing the robust aggregators' difficulty working in a heterogeneous context when considering demographic bias.




Fig. 3. Impact of FL robustness methods on Bias

4.2 Impact of FL Robustness Methods on FL Bias Mitigation

We study the behavior of two robust aggregators, Multi-Krum and NDC when used with a state-of-the-art bias mitigation method called FCFL [3]. We add the robust method as the first step after getting clients’ updates before sending the updates to the FCFL server solver to perform the optimization task. Using NDC with FCFL means the clients’ updates are clipped if their norm exceeds a fixed threshold and sent to the server solver while for Multi-Krum it means that clients are selected to be sent to the server solver based on the set of clients with the lowest Krum score. The solver will then work with fewer updates to compute the global model. For all experiments, we used the same hyper-parameters on FCFL, specifically the SPD threshold and the learning rate. As explained earlier, we perform such experiments on the same datasets and sensitive attributes. We report the model bias of both added robust mechanisms against using FCFL alone and the baseline FL method in Figure 4. We observe how the robust mechanism influences the behavior of FCFL. Specifically, on both MEPS (Figure 4(c)) and Adult (Figure 4(a),4(b)), we observed that NDC influenced the bias mitigation process, getting a higher value of SPD than the typical method behavior. For Multi-Krum, we can see its influence on FCFL the most on Adult (Figure 4(a)) and MEPS (Figure 4(c)). We observe that using classical robustness mechanism in conjunction with FCFL worsens the bias mitigation process, further motivating the issue of building a bias-aware FL robustness mechanism.

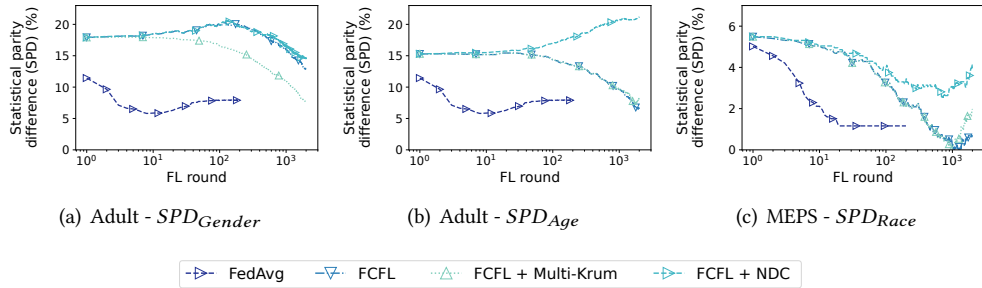


Fig. 4. Interaction between FL bias mitigation and FL robustness mechanisms

5 PROBLEM FORMULATION

A FL system must be robust to Byzantine behavior and be able to mitigate model bias, all of which are described as follows. We consider n FL clients, assuming a fraction $\rho \leq 0.5$ of Byzantine clients, and we denote \mathcal{H} the set of honest client. Each client’s data is represented as D_i with S the sensitive attribute, and the local model denoted as θ_i .

FL Robustness to Byzantine Behavior. The FL system must be robust against the attack of $f = \rho n$ Byzantine client, which is achieved when for any $\epsilon > 0$, the server can output a global model θ such that we can bound the diameter of the gradient of the loss function $f(\theta)$ by a constant ϵ . Ideally, the system can

distinguish the honest FL clients, thus selecting only the latter for aggregation. We call such aggregation $\bar{\theta} = \frac{1}{|\mathcal{H}|} \sum_{j \in \mathcal{H}} \theta_j$ the true average. A robust aggregator must output a model θ close to the true average. Assuming that we can bound the honest client model heterogeneity $\|\theta_j - \theta_i\|^2 < \delta^2$ for all $i, j \in \mathcal{H}$, the output $\hat{\theta}$ from the robust Byzantine aggregator must respect the following constraint [8] :

$$\mathbb{E}\|\hat{\theta} - \bar{\theta}\|^2 < k\rho\delta^2 \quad (3)$$

FL Bias Mitigation. The goal is to learn the optimal value of the global model parameters θ that solve the bias-constrained optimization on the SPD_S value, where ϵ is the SPD_S constraint and $f_k(\theta)$ is the loss of the given model θ over local data of the honest client k :

$$\min_{\theta} f(\theta) = \min_{\theta} \frac{1}{|\mathcal{H}|} \sum_{k \in \mathcal{H}} f_k(\theta) \quad (4a)$$

$$\text{s.t } |SPD_S(\theta)| \leq \epsilon \quad (4b)$$

6 WHY HANDLING BIAS AND BYZANTINE BEHAVIOR IN FL IS DIFFICULT

This section shows the challenges of building a FL system combining robustness against Byzantine behavior, and model bias mitigation. We cover the following scenarios: Using classical robust aggregators first and then the FL Bias mitigation process separately (in section 6.1). Using the FL Bias mitigation process first and then the robustness mechanism separately (in section 6.2) and finally combine both robustness and bias mitigation objectives at once (in section 6.3). We provide then interesting research direction to address the cited challenges.

6.1 Why Applying a Classical FL Robustness Mechanism Followed by Classical FL Bias Mitigation Does Not Work

Suppose we first apply classical robust aggregation techniques and subsequently perform bias mitigation. The resulting system may fail to operate effectively in a non-IID environment. With a FL system consisting of n clients, with one honest client having a very different data distribution with better representation of the minority group, its contribution to the FL system would then help bring down the global model bias. However, we must expect our FL system to be robust to Byzantine clients by assuming the latter’s presence within our system. Let us also suppose the presence of one Byzantine clients in our system. Using classical robust aggregators such as Multi-Krum would mean that one client get eliminated based on how far the model parameters are from all other clients. In section 4, we illustrated how robust aggregators could eliminate the client with different data representativity. Unfortunately, this can lead to the exclusion of honest clients that offer demographic minority representation (see Figure 1), losing crucial data representation that could be used by bias mitigation methods to help reduce the bias between demographic groups.

Observation 1:

Using classical robust aggregators may eliminate honest client, affecting the normal behavior of FL bias mitigation.

6.2 Why Applying Classical FL Bias Mitigation Followed By a Classical FL Robustness Mechanism Does Not Work

In a FL setup where we suppose the presence of Byzantine clients, we must keep the Byzantine client from being considered by the FL Bias mitigation method as we need to trust their behavior. Suppose a FL system with honest clients and Byzantine clients. At the end of the FL round, each honest client sends regular updates, while the Byzantine sends arbitrary information. In this scenario, the FL bias mitigation component is applied before any robustness mechanism in this scenario. The classical bias mitigation method happening server-side would then get exposed to the Byzantine clients, which can directly harm the model utility (see Figure 2). A tradeoff with model utility is often expected with the bias mitigation methods [3, 5, 20]. Adding Byzantine clients would further increase the tradeoff, making some bias-constrained optimization tasks like AgnosticFair [20] more difficult. Furthermore, as mentioned in section 3.2, some bias mitigation methods that intervene server-side, such as FairFed [5], require clients to send information on their local bias metrics in order to compute the global model. Byzantines can exploit this mechanism by sending falsified local bias information, which can cause the system to consider their harmful model updates even better.

Observation 2:

Using the classical FL bias mitigation method before any robustness mechanism expose the bias mitigation method to the influence of the Byzantine clients.

6.3 On the Impossibility to Handle Bias and Byzantine Behavior at Once in Classical Approaches

Each client sends their locally trained model update to the server during the FL process for optimization. In our case, we assume wanting to maximize the global model utility on a data distribution representing the union of honest client data under bias constraint. At the same time, we want our model to have robustness guarantees against the current Byzantine activity in the system. In order to solve the bias-constrained objective problem on the global data distribution, we must optimize our objective based on the set of models updates coming from the honest clients set, and nothing else as we do not trust the updates coming from the Byzantine clients as mentioned in subsection 6.2. However, such a set of honest client updates is heterogeneous due to outlier clients, which we assume represents additional information on minorities compared to the one found in most updates. In [8], Karimireddy and al. argue that we can construct a set of model updates including all Byzantines that would pass the robust aggregation part. Our system may consider those Byzantine updates for the server-side optimization part, which brings an error component

to the global optimization process that we cannot ignore. The output of such a system combining bias mitigation and robustness is not guaranteed to keep the bias constraint under the global data distribution.

Observation 3:

Using both objectives is impossible in heterogeneous settings using only model updates as information, as the server cannot distinguish between outlier clients and Byzantine ones.

6.4 Discussion and Research Directions

Despite the challenges of combining bias mitigation and robustness to Byzantine behavior methods in FL, interesting research directions remain to explore. We have seen the necessity of providing robustness before any other bias mitigation operation at the server level in subsection 6.2. This requires a new method for detecting Byzantine clients in non-IID environments. First, such a FL system must guarantee the selection of outlier clients when there are no Byzantine attacks on the system. Also, the FL system must reduce the amount of falsely predicted Byzantine clients, as these clients probably contain critical data representativity. The server can try spotting such "honest and minority" clients among the deemed Byzantine clients by the robust aggregators. To distinguish such clients from the Byzantine ones, we could rely on reputation mechanisms to precisely identify Byzantine clients [21]. Furthermore, we saw how most FL bias mitigation approaches require the clients sending additional information on local bias (see subsection 3.2). Then a new solution could have the server also asking its clients to send additional information on their local data distribution and privately compare such information with their sent model updates. The system could then analyze inconsistencies between the information on the data distribution and what can be extracted from the client model update (i.e., using model inversion [24]), revealing the Byzantine nature of the client. Then selecting of the "honest and minority" clients could be performed on how interesting their added data representativity for minorities are.

7 CONCLUSION

FL model bias and robustness against Byzantines have been extensively studied over the last few years. Addressing both problems within a single FL system, while critically needed, has been elusive. This study investigated the challenges of constructing a FL system with both robustness to Byzantine behavior properties and bias guarantees. By design, a robust FL system degrades inclusiveness, particularly for outlier model behavior, which may have the side effect of setting aside data representing minorities, which is necessary for proper functioning FL bias mitigation techniques. Furthermore, we discuss the limitations of combining classical bias mitigation methods with existing robustness mechanisms. Finally we formulate possible research directions for building robust, bias-free FL.

ACKNOWLEDGMENTS

This work was partially supported by the ANR French National Research Agency, through SIBIL-Lab LabCom (ANR-17-LCV2-0014) and ByBlos project (ANR-20-CE25-0002-01). Thanks to Paul Wawszczyk, who has also contributed to this work during his internship.

REFERENCES

- [1] Annie Abay, Yi Zhou, Nathalie Baracaldo, Shashank Rajamoni, Ebube Chuba, and Heiko Ludwig. 2020. Mitigating Bias in Federated Learning. [arXiv:2012.02447](https://arxiv.org/abs/2012.02447)
- [2] Steven Cohen. 2003. Design Strategies and Innovations in the Medical Expenditure Panel Survey. *Medical care* 41 (08 2003), III5–III12.
- [3] Sen Cui, Weishen Pan, Jian Liang, Changshui Zhang, and Fei Wang. 2021. Addressing Algorithmic Disparity and Performance Inconsistency in Federated Learning. *Advances in Neural Information Processing Systems* 34 (2021), 26091–26102.
- [4] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness Through Awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (Cambridge, Massachusetts) (ITCS '12)*. Association for Computing Machinery, New York, NY, USA, 214–226.
- [5] Yahya H. Ezzeldin, Shen Yan, Chaoyang He, Emilio Ferrara, and Salman Avestimehr. 2021. FairFed: Enabling Group Fairness in Federated Learning. [arXiv:2110.00857](https://arxiv.org/abs/2110.00857)
- [6] Hashemi Hanieh, Wang Yongqin, Guo Chuan, and Annavam Murali. 2021. Byzantine-Robust and Privacy-Preserving Framework for FedML. [arXiv:2105.02295](https://arxiv.org/abs/2105.02295)
- [7] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. 2019. Measuring the Effects of Non-Identical Data Distribution for Federated Visual Classification. [arXiv:1909.06335](https://arxiv.org/abs/1909.06335)
- [8] Sai Praneeth Karimireddy, Lie He, and Martin Jaggi. 2020. Byzantine-Robust Learning on Heterogeneous Datasets via Bucketing.
- [9] Ron Kohavi. 1996. Scaling Up the Accuracy of Naive-Bayes Classifiers: A Decision-Tree Hybrid. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (Portland, Oregon) (KDD'96)*. AAAI Press, Portland, Oregon, USA, 202–207.
- [10] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. 2020. Ditto: Fair and Robust Federated Learning Through Personalization.
- [11] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera Y. Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 54)*, Aarti Singh and Jerry Zhu (Eds.). PMLR, Fort Lauderdale Florida, USA, 1273–1282.
- [12] Mehrabi Ninareh, Morstatter Fred, Saxena Nripsuta, Lerman Kristina, and Galstyan Aram. 2022. A Survey on Bias and Fairness in Machine Learning. *Comput. Surveys* 54, 6 (2022), 115:1–115:35.
- [13] Blanchard Peva, El Mhamdi El Mahdi, Guerraoui Rachid, and Stainer Julien. 2017. Machine Learning with Adversaries: Byzantine Tolerant Gradient Descent. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, Guyon Isabelle, von Luxburg Ulrike, Bengio Samy, Wallach Hanna M., Fergus Rob, Vishwanathan S. V. N., and Garnett Roman (Eds.). The MIT Press, Long Beach, CA, USA, 119–129.
- [14] Krishna Pillutla, Sham M. Kakade, and Zaid Harchaoui. 2019. Robust Aggregation for Federated Learning. [arXiv:1912.13445](https://arxiv.org/abs/1912.13445)
- [15] Daneshjou Roxana, Vodrahalli Kailas, Novoa Roberto A., Jenkins Melissa, Liang Weixin, Rotemberg Veronica, Ko Justin, Swetter Susan M., Bailey Elizabeth E., Gevaert Olivier, Mukherjee Pritam, Phung Michelle, Yekrang Kiana, Fong Bradley, Sahasrabudhe Rachna, Allerup Johan A. C., Okata-Karigane Utako, Zou James, and Chiou Albert S. 2022. Disparities in dermatology AI performance on a diverse, curated clinical image set. *Science Advances* 8, 32 (2022), eabq6147.

- [16] Ashneet Singh, Alberto Blanco-Justicia, and Josep Domingo-Ferrer. 2023. Fair detection of poisoning attacks in federated learning on non-i.i.d. data. *Data Mining and Knowledge Discovery* 1 (2023), -.
- [17] Tianshu Song, Yongxin Tong, and Shuyue Wei. 2019. Profit Allocation for Federated Learning. In *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, Los Angeles, CA, USA, 2577–2586.
- [18] Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H. Brendan McMahan. 2019. Can You Really Backdoor Federated Learning? arXiv:1911.07963
- [19] Hongyi Wang, Kartik Sreenivasan, Shashank Rajput, Harit Vishwakarma, Saurabh Agarwal, Jy-yong Sohn, Kangwook Lee, and Dimitris Papailiopoulos. 2020. Attack of the Tails: Yes, You Really Can Backdoor Federated Learning. arXiv:2007.05084
- [20] Du Wei, Xu Depeng, Wu Xintao, and Tong Hanghang. 2021. *Fairness-Aware Agnostic Federated Learning*. SIAM, Virtual Event, 181–189.
- [21] Xinyi Xu and Lingjuan Lyu. 2020. A Reputation Mechanism Is All You Need: Collaborative Fairness and Adversarial Robustness in Federated Learning. arXiv:2011.10464
- [22] Dong Yin, Yudong Chen, Ramchandran Kannan, and Peter Bartlett. 2018. Byzantine-Robust Distributed Learning: Towards Optimal Statistical Rates. In *Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 80)*, Jennifer Dy and Andreas Krause (Eds.). PMLR, Stockholm, Sweden, 5650–5659.
- [23] Daniel Yue Zhang, Kou Ziyi, and Wang Dong. 2020. FairFL: A Fair Federated Learning Approach to Reducing Demographic Bias in Privacy-Sensitive Classification Models. In *Proceedings of the 2020 IEEE International Conference on Big Data, BigData2020*. IEEE, United States, 1051–1060.
- [24] Bo Zhao, Peng Sun, Tao Wang, and Keyu Jiang. 2022. FedInv: Byzantine-Robust Federated Learning by Inverting Local Model Updates. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. Association for the Advancement of Artificial Intelligence, Virtual, 9171–9179.