

FairFed: Enabling Group Fairness in Federated Learning

Yahya H. Ezzeldin*, Shen Yan*, Chaoyang He, Emilio Ferrara, Salman Avestimehr

University of Southern California (USC)

yessa@usc.edu, shenyan@usc.edu, chaoyang.he@usc.edu, emiliofe@usc.edu, avestime@usc.edu

Abstract

Training ML models which are fair across different demographic groups is of critical importance due to the increased integration of ML in crucial decision-making scenarios such as healthcare and recruitment. Federated learning has been viewed as a promising solution for collaboratively training machine learning models among multiple parties while maintaining their local data privacy. However, federated learning also poses new challenges in mitigating the potential bias against certain populations (e.g., demographic groups), as this typically requires centralized access to the sensitive information (e.g., race, gender) of each datapoint. Motivated by the importance and challenges of group fairness in federated learning, in this work, we propose FairFed, a novel algorithm for fairness-aware aggregation to enhance group fairness in federated learning. Our proposed approach is server-side and agnostic to the applied local debiasing thus allowing for flexible use of different local debiasing methods across clients. We evaluate FairFed empirically versus common baselines for fair ML and federated learning and demonstrate that it provides fairer models, particularly under highly heterogeneous data distributions across clients. We also demonstrate the benefits of FairFed in scenarios involving naturally distributed real-life data collected from different geographical locations or departments within an organization.

Introduction

An important notion of fairness in machine learning, *group fairness* (Dwork et al. 2012), concerns the mitigation of bias in the performance of a trained model against certain protected demographic groups, which are defined based on sensitive attributes within the population (e.g., gender, race). Several approaches to achieve group fairness have been studied in recent years in centralized settings. However, these approaches rely on the availability of the entire dataset at a central entity during training and are therefore unsuitable for application in Federated Learning (FL).

Federated learning allows for decentralized training of large-scale models without requiring direct access to clients' data, hence maintaining their privacy (Kairouz et al. 2021; Wang et al. 2021a). However, this decentralized nature makes it complicated to translate solutions for fair training

*These authors contributed equally.

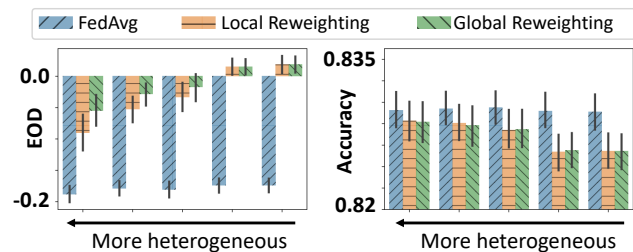


Figure 1: Comparison of local/global debiasing under different heterogeneity levels. For the Equal Opportunity Difference (EOD) metric, values close to 0 indicate better fairness.

from centralized settings to FL, where the decentralization of data is a major cornerstone. This gives rise to the key question that we attempt to answer in this paper: *How can we train a classifier using FL so as to achieve group fairness, while maintaining data decentralization?*

Potential approaches for group fairness in FL. One potential solution that one may consider for training fair models in FL is for each client to apply local debiasing on its locally trained models (without sharing any additional information or data), while the FL server simply aggregates the model parameters in each round using FL aggregation algorithms such as FedAvg (McMahan et al. 2017) (or its subsequent derivatives, e.g., FedOPT (Reddi et al. 2020), FedNova (Wang et al. 2020)). Although this allows for training a global model without explicitly sharing the local datasets, but the drawback is that applying a debiasing mechanism at each client in isolation on its local dataset can result in poor performance in scenarios where data distributions are highly-heterogeneous across clients (See Figure 1).

Another potential solution for fair training in FL would be to adapt a debiasing technique from the rich literature of centralized fair training to be used in FL. Although this may result in reasonable fair training (see Figure 1), however, in the process of applying this debiasing globally, the clients may need to exchange additional detailed information with the server about their dataset constitution which can leak information about different subgroups in the client's dataset. For example, the server may require knowledge of the model's performance on each group in a client's dataset and/or local statistical information about each group in the dataset.

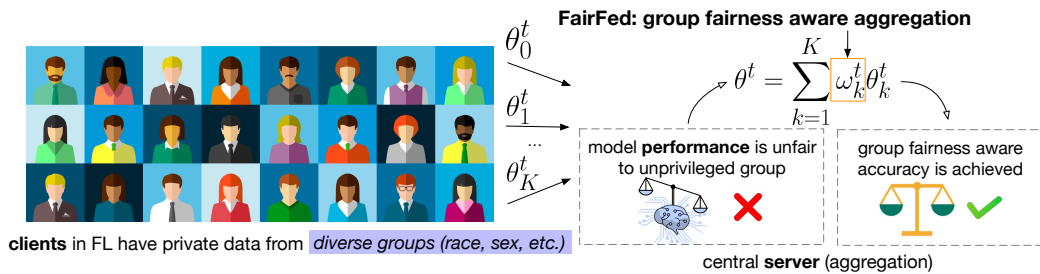


Figure 2: FairFed: Group fairness-aware FL framework.

The Proposed FairFed Approach. Motivated by the drawbacks of the two aforementioned directions, in this work, we propose *FairFed*, a strategy to train fair models via a fairness-aware aggregation method (Figure 2). In FairFed, each client performs *local debiasing* on its own local dataset, thus maintaining data decentralization and avoiding the exchange of any explicit information of its local data composition. To amplify the local debiasing performance, the clients will evaluate the *fairness of the global model on their local datasets* in each FL round and collectively collaborate with the server to adjust its model aggregation weights. The weights are a function of the mismatch between the global fairness measurement (on the full dataset) and the local fairness measurement at each client, favoring clients whose local measures match the global measure. We carefully design the exchange between the server and clients during weights computation, making use of the secure aggregation protocol (Bonawitz et al. 2017) to prevent the server from learning any explicit information about any single client’s dataset.

The server-side/local debiasing nature of FairFed gives it the following benefits over existing fair FL strategies:

- **Enhancing group fairness under data heterogeneity:** One of the biggest challenges to group fairness in FL is the heterogeneity of data distributions across clients, which limits the impact of local debiasing efforts on the global data distribution. FairFed shows significant improvement in fairness performance under highly heterogeneous distribution settings and outperforms state-of-the-art methods for fairness in FL, indicating promising implications from applying it to real-life applications.
- **Freedom for different debiasing across clients:** As FairFed works at the server side and only requires evaluation metrics of the model fairness from the clients, it is more flexible to run on top of heterogeneous client debiasing strategies (we expand on this notion in Sections). For example, different clients can adopt different local debiasing methods based on the properties (or limitations) of their devices and data partitions.

Background and Related Work

Group fairness in centralized learning. In classical centralized ML, common approaches for realizing group fairness can be classified into three categories: pre-processing (Grgić-Hlača et al. 2018; Feldman et al. 2015), in-processing (Kamishima et al. 2012; Zhang, Lemoine, and Mitchell 2018; Roh et al. 2021) and post-processing (Lohia et al. 2019; Kim, Ghorbani, and Zou 2019) techniques. However,

a majority of these techniques need centralized access to the sensitive information (e.g., race) of each datapoint, making it unsuitable for FL. As a result developing effective approaches for fair FL is an important area of study.

Fairness in federated learning. New challenges in FL have introduced different notions of fairness. These new notions include for example, *client-based fairness* (Li et al. 2019; Mohri, Sivek, and Suresh 2019) which aims to equalize model performance across different clients or *collaborative fairness* (Lyu et al. 2020; Wang et al. 2021b) which aims to reward a highly-contributing participant with a better performing local model than is given to a low-contributing participant. In this paper, we instead focus on the notion of *group fairness* in FL, where each datapoint in the FL system belongs to particular group, and we aim to train models that do not discriminate against any group of datapoints.

Several recent works have made progress on group fairness in FL. One common research direction is to distributively solve an optimization objective with fairness constraints (Zhang, Kou, and Wang 2020; Du et al. 2021; Gálvez et al. 2021), which requires each client to share the statistics of the sensitive attributes of its local dataset to the server. The authors in (Abay et al. 2020) investigated the effectiveness of adopting a global reweighting mechanism. In (Zeng, Chen, and Lee 2021), an adaptation of the FairBatch debiasing algorithm (Roh et al. 2021) is proposed for FL where clients use FairBatch locally and the weights are updated through the server in each round. In (Papadaki et al. 2021), an algorithm is proposed to achieve minimax fairness in federated learning. In these works, the server requires each client to explicitly share the performance of the model on each subgroup separately; for example (males with +ve outcomes, females with +ve outcomes, etc). Differently from these works, our proposed FairFed method does not restrict the local debiasing strategy of the participating clients, thus increasing the flexibility of the system. Furthermore, FairFed does not share explicit information on the model performance for any specific group within a client’s dataset. Finally, our empirical evaluations consider extreme cases of data heterogeneity and demonstrate that our method can yield significant fairness improvements in these situations.

Preliminaries

We begin by reviewing the standard FL setup (McMahan et al. 2017), and then introduce key definitions and metrics for group fairness. We then extend these to the FL setting by defining the notions of global and local fairness in FL.

Federated Learning Setup

Following a standard FL setting (McMahan et al. 2017), we consider a scenario where K clients collaborate with a server to find a parameter vector θ that minimizes the weighted average of the loss across all clients. In particular:

$$\min_{\theta} f(\theta) = \sum_{k=1}^K \omega_k L_k(\theta), \quad (1)$$

where: $L_k(\theta)$ denotes the local objective at client k ; $\omega_k \geq 0$, and $\sum \omega_k = 1$. The local objective L_k 's can be defined by empirical risks over the local dataset \mathcal{D}_k of size n_k at client k , i.e., $L_k(\theta) = \frac{1}{n_k} \sum_{(\mathbf{x}, y) \in \mathcal{D}_k} \ell(\theta, \mathbf{x}, y)$.

To minimize the objective in (1), the federated averaging algorithm FedAvg, proposed in (McMahan et al. 2017), samples a subset of the K clients per round to perform local training of the global model on their local datasets. The model updates are then averaged at the server, being weighted based on the size of their respective datasets.

To ensure the server does not learn any information about the values of the individual transmitted updates from the clients beyond the aggregated value that it sets out to compute, FedAvg typically employs a Secure Aggregation (SecAgg) algorithm (Bonawitz et al. 2017).

Training using FedAvg and its subsequent improvements (e.g., FedOPT (Reddi et al. 2020)) allows training of a high-performance global model, however, this collaborative training can result in a global model that discriminates against an underlying demographic group of datapoints (similar to biases incurred in centralized training of machine learning models (Dwork et al. 2012)). We highlight key notions of group fairness in fair ML in the following subsection.

Notions of Group Fairness

In sensitive machine learning applications, a data sample often contains private and sensitive demographic information that can lead to discrimination. In particular, we assume that each datapoint is associated with a sensitive binary attribute A (e.g., gender or race). For a binary prediction model $\hat{Y}(\theta, \mathbf{x})$, the fairness is evaluated with respect to its performance compared to the underlying groups defined by the sensitive attribute A . We use $A = 1$ to represent the privileged group (e.g., male), while $A = 0$ is used to represent the under-privileged group (e.g., female). For the binary model output \hat{Y} (and similarly the label Y), $\hat{Y} = 1$ is assumed to be the positive outcome. Using these definitions, we can now define two group fairness notions that are applied in group fairness literature for centralized training:

Definition 1 (Equal Opportunity) : *Equal opportunity (Hardt et al. 2016) measures the performance a binary predictor \hat{Y} with respect to A and Y . The predictor is considered fair from the equal opportunity perspective if the true positive rate is independent of the sensitive attribute A . To measure this, we use the Equal Opportunity Difference (EOD), defined as*

$$EOD = \Pr(\hat{Y}=1|A=0, Y=1) - \Pr(\hat{Y}=1|A=1, Y=1). \quad (2)$$

Definition 2 (Statistical Parity) : *Statistical parity (Dwork et al. 2012) rewards the classifier for classifying each group*

as positive at the same rate. Thus, a binary predictor \hat{Y} is fair from the statistical parity perspective if $\Pr(\hat{Y} = 1|A = 1) = \Pr(\hat{Y} = 1|A = 0)$. Thus, the Statistical Parity Difference (SPD) metric is defined as

$$SPD = \Pr(\hat{Y} = 1|A = 0) - \Pr(\hat{Y} = 1|A = 1). \quad (3)$$

For the EOD and SPD metrics, values closer to zero indicate better fairness. Positive fairness metrics indicate that the unprivileged group outperform the privileged group.

Global vs Local Group Fairness in FL

The fairness definitions above can be readily applied to centralized model training to evaluate the performance of the trained model. However, in FL, clients typically have non-IID data distributions, which gives rise to different levels of consideration for fairness: *global fairness* and *local fairness*.

The **global fairness** of a given model considers the full dataset $\mathcal{D} = \cup_k \mathcal{D}_k$ across the K clients in FL, which is our end-goal in fair FL; to train a model that is in general non-discriminatory to any group in the global dataset.

In contrast, when only the local dataset \mathcal{D}_k at client k is considered, we can define the **local fairness** performance by applying (2) or (3) on the data distribution at client k .

We can highlight the difference between global and local fairness using the example of the EOD metric. For a classifier \hat{Y} , the global fairness EOD metric F_{global} is given by

$$F_{global} = \Pr(\hat{Y}=1|A=0, Y=1) - \Pr(\hat{Y}=1|A=1, Y=1), \quad (4)$$

where the probability above is based on the full dataset distribution (a mixture of the distributions across the clients). In contrast, the local fairness metric F_k at client k is

$$F_k = \Pr(\hat{Y} = 1|A = g, Y = 1, C = k) - \Pr(\hat{Y} = 1|A = g, Y = 1, C = k), \quad (5)$$

where the parameter $C = k$ denotes that the k -th client (and dataset \mathcal{D}_k) is considered in the fairness evaluation.

Note that if clients have IID distributions (i.e., distributions that are independent of C), global and local fairness match. However, they can greatly differ in the non-IID case.

FairFed: Fairness-Aware Aggregation in FL

In this section, we introduce our proposed approach FairFed which uses local debiasing due for its advantages for data decentralization, while addressing its challenges by adjusting how the server aggregates local model updates from clients.

Our Proposed Approach (FairFed)

Recall that in the t -th round in FedAvg (McMahan et al. 2017), local model updates $\{\theta_k^t\}_{k=1}^K$ are weight-averaged to get the new global model parameter θ^t as: $\theta^{t+1} = \sum_{k=1}^K \omega_k^t \theta_k^t$, where the weights $\omega_k^t = n_k / \sum_k n_k$ depend only on the number of datapoints at each client.

As a result, a fairness-oblivious aggregation would favor clients with more datapoints. If the training at these clients results in locally biased models, then the global model can potentially be biased since the weighted averaging exaggerates the contribution of model update from these clients.

Algorithm 1: FairFed Algorithm (tracking EOD)

Initialize: global model parameter θ_0 and weights $\{\omega_k^0\}$ as $\omega_k^0 = n_k / \sum_{i=1}^K n_i, \forall k \in [K]$;

Dataset statistics: Aggregate statistics $\mathcal{S} = \{\Pr(A=1, Y=1), \Pr(A=0, Y=1)\}$ from clients using Secure Aggregation (SecAgg) and send it to clients;

for each round $t = 1, 2, \dots$ **do**

$F_{global}^t, \overline{Acc}^t \leftarrow \text{SecAgg} \left(\left\{ \text{ClientLocalMetrics} (k, \theta^{t-1}) \right\}_{k=1}^K \right)$; // SecAgg to get \overline{Acc}^t and global fairness F_{global}^t as in (7);

$\frac{1}{K} \sum_i \Delta_i \leftarrow \text{SecAgg} \left(\left\{ \text{ClientMetricGap} (k, \theta^{t-1}, F_{global}^t, \overline{Acc}^t) \right\}_{k=1}^K \right)$; // SecAgg to compute mean of metric gaps;

// Compute aggregation weights locally at clients based on (6) then use SecAgg to aggregate weighted local model updates;

$\left(\sum_{k=1}^K \bar{\omega}_k^t \theta_k^t \right), \left(\sum_{k=1}^K \bar{\omega}_k^t \right) \leftarrow \text{SecAgg} \left(\left\{ \text{ClientWeightedModelUpdate} (k, \theta^{t-1}, \omega_k^t, \frac{1}{K} \sum_i \Delta_i) \right\}_{k=1}^K \right)$;

$\theta^{t+1} \leftarrow \left(\sum_{k=1}^K \bar{\omega}_k^t \theta_k^t \right) / \left(\sum_{k=1}^K \bar{\omega}_k^t \right)$;

Based on this observation, in FairFed, we propose a method to optimize global group fairness F_{global} via adaptively adjusting the aggregation weights of different clients based on their local fairness metric F_k . In particular, given the current global fairness metric F_{global}^t (we will discuss later in the section, how the server can compute this value), then in the next round, the server gives a slightly higher weight to clients that have a similar local fairness F_k^t to the global fairness metric, thus relying on their local debiasing to steer the next model update towards a fair global model.

Next, we detail how FairFed computes the aggregation weights in each round. The steps performed while tracking the EOD metric in FairFed are shown in Algorithm 1.

Computing Aggregation Weights for FairFed

At the beginning of training, we start with the default FedAvg weights $\omega_k^0 = n_k / \sum_{k=1}^K n_k$. Next, in each round t , we update the weight assigned to the k -th client based on the current gap between its local fairness metric F_k^t and the global fairness metric F_{global}^t . In particular, the weight update follows this formula $\forall k \in [K]$:

$$\Delta_k^t = \begin{cases} |Acc_k^t - \overline{Acc}^t| & \text{if } F_k^t \text{ is undefined,} \\ |F_{global}^t - F_k^t| & \text{otherwise} \end{cases},$$
$$\bar{\omega}_k^t = \bar{\omega}_k^{t-1} - \beta \left(\Delta_k - \frac{1}{K} \sum_{i=1}^K \Delta_i \right), \quad \omega_k^t = \frac{\bar{\omega}_k^t}{\sum_{i=1}^K \bar{\omega}_i^t}. \quad (6)$$

where: (i) Acc_k^t represents the local accuracy at client k , and $\overline{Acc}^t = \sum_{k=1}^K Acc_k \times n_k / \sum_{k=1}^K n_k$, is global accuracy across the full dataset, respectively; (ii) β is a parameter that controls the fairness budget for each update, thus impacting the trade-off between model accuracy and fairness. Higher values of β result in fairness metrics having a higher impact on the model optimization, while a lower β results in a reduced perturbation to the default FedAvg weights due to fair training; note that at $\beta = 0$, FairFed is equivalent to FedAvg, as the initial weights ω_k^0 are unchanged.

Intuition for weight update. The intuition behind the update in (6) is to effectively rank the clients in the FL system based on how far their local view of fairness (measured through their local metric) compares to the global fairness

metric; closer views to the global metric get assigned higher weights while clients with local metrics that significantly deviate from the global metric will have their weights reduced. The significance is decided based on whether the gap from the global metric is above the average gap (across clients) or vice versa. Note that, whenever, the client distribution makes the local metric F_k undefined¹, FairFed relies on the discrepancy between the local and global accuracy metric as a proxy to compute the fairness metric gap Δ_k .

Thus, so far, the training process of FairFed at each iteration follows the following conceptual steps:

1. Each client computes its updated local model parameters;
2. The server computes the global fairness metric value F_{global}^t and global accuracy \overline{Acc}^t using secure aggregation and broadcasts them back to the clients;
3. Each client computes its metric gap Δ_k^t and from it, it calculates its aggregation weight ω_k^t with the help of the server as defined in (6);
4. The server next aggregates the weighted local updates $\omega_k^t \theta_k^t$ using secure aggregation to compute the new global model and broadcasts it to the clients.

A detailed description on performing these steps using secure aggregation (SecAgg) is shown in Algorithm 1, where $\text{SecAgg}(\{b_i\}_{i=1}^K)$ computes $\sum_i b_i$ using secure aggregation.

Flexibility of FairFed with heterogeneous debiasing.

Note that the FairFed weights ω_k^t in (6) rely only on the global and local fairness metrics and are not tuned towards a specific local debiasing method. Thus, FairFed is flexible to be applied with different debiasing methods at each client, and the server will incorporate the effects of these different methods by reweighting their respective clients based on their local/global metrics and the weight computation in (6).

How the Server Gets the Global Metric F_{global}

One central ingredient for computing weights in FairFed is the server's ability to calculate the global metric F_{global} in each round (recall equation (6)) without the clients having to share their datasets with the server or any explicit information about their local group distribution. If the metric of

¹For instance, in the case of the EOD metric, this happens whenever $\Pr(A = 1, Y = 1) = 0$ or $\Pr(A = 0, Y = 1) = 0$.

interest is EOD, we next show how the server can compute F_{global} from the clients using secure aggregation. Similar computations follow for SPD and are presented in (Ezzeldin et al. 2021, Appendix A). Let $n = \sum_{k=1}^K n_k$; the EOD metric in (4) can be rewritten as:

$$F_{global} = \Pr(\hat{Y} = 1 | A = 0, Y = 1) - \Pr(\hat{Y} = 1 | A = 1, Y = 1) \\ = \sum_{k=1}^K \frac{n_k}{n} \left[\frac{\Pr(\hat{Y} = 1 | A = 0, Y = 1, C = k) \Pr(A = 0, Y = 1 | C = k)}{\Pr(Y = 1, A = 0)} \right. \\ \left. - \underbrace{\frac{\Pr(\hat{Y} = 1 | A = 1, Y = 1, C = k) \Pr(A = 1, Y = 1 | C = k)}{\Pr(Y = 1, A = 1)}}_{m_{global,k}} \right], \quad (7)$$

where $m_{global,k}$ is the summation component that each client k computes locally. Thus, the global EOD metric F_{global} can be computed at the server by applying secure aggregation (Bonawitz et al. 2017) to get the sum of the $m_{global,k}$ values from the K clients without the server learning any information about the individual $m_{global,k}$ values.

Note that the conditional probabilities defining $m_{global,k}$ in (7) are all local performance metrics that can easily be computed locally by client k using its local dataset \mathcal{D}_k . The only non-local terms in $m_{global,k}$ are the full dataset statistics $\mathcal{S} = \{\Pr(Y=1, A=0), \Pr(Y=1, A=1)\}$. These statistics \mathcal{S} can be aggregated at the server using a single round of secure aggregation (e.g., (Bonawitz et al. 2017)) at the start of training, then be shared with the clients to enable them to compute their global fairness component $m_{global,k}$.

Experimental Evaluation

In this section, we investigate the performance of FairFed under different system settings. In particular, we explore how the performance changes with different heterogeneity levels in data distributions across clients. We also evaluate how the trade-off between fairness and accuracy changes with the fairness budget β in FairFed (see equation (6)). Additional experiments in (Ezzeldin et al. 2021, Appendix C) investigate how the performance of FairFed changes when different local debiasing approaches are used across clients.

Experimental Setup

Implementation. We developed FairFed using FedML (He et al. 2020), which is a research-friendly FL library for exploring new algorithms. We use a server with AMD EPYC 7502 32-Core CPU Processor, and use a parallel training paradigm, where each client is handled by an independent process using MPI (message passing interface).

Datasets. In this section, we demonstrate the performance of different debiasing methods using two binary decision datasets that are widely investigated in fairness literature: the Adult (Dua and Graff 2017) dataset and ProPublica COMPAS dataset (Larson et al. 2016). In the **Adult** dataset (Dua and Graff 2017), we predict the yearly income (with binary label: over or under \$50,000) using twelve categorical or continuous features. The gender (defined as male or female) of each subject is considered the sensitive attribute. The

ProPublica **COMPAS** dataset relates to recidivism, which is to assess if a criminal defendant will commit an offense within a certain future time. Features in this dataset include the number of prior offenses, the age of the defendant, etc. The race (classified as white or non-white) of the defendant is the sensitive attribute of interest.

Configurable data heterogeneity for diverse sensitive attribute distributions. To understand the performance of our method and the baselines, under different distributions of the sensitive attribute across clients, a configurable data synthesis method is needed. In our context, we use a generic non-IID synthesis method based on the Dirichlet distribution proposed in (Hsu, Qi, and Brown 2019) but apply it in a novel way for *configurable sensitive attribute distribution*: for each sensitive attribute value a , we sample $\mathbf{p}_a \sim \text{Dir}(\alpha)$ and allocate a portion $p_{a,k}$ of the datapoints with $A = a$ to client k . The heterogeneity across clients is controlled via α , where $\alpha \rightarrow \infty$ results in IID distributions. Examples of these heterogeneous distributions for the Adult and COMPAS are shown in (Ezzeldin et al. 2021, Appendix B).

Baselines. We adopt the following state-of-the-art solutions as baselines:

- **FedAvg** (McMahan et al. 2017): the original FL algorithm for distributed training of private data. It does not consider fairness for different demographic groups.
- **FedAvg + Local reweighting [Local / RW]**: Each client adopts the reweighting strategy (Kamiran and Calders 2012) to debias its local training data, then trains local models based on the pre-processed data. FedAvg is used to aggregate the local model updates at the server.
- **FedAvg + FairBatch [Local / FairBatch]**: Each client adopts the state-of-the-art FairBatch in-processing debiasing strategy (Roh et al. 2021) to debias its local training data and then aggregation uses FedAvg.
- **FedAvg + Fair Linear Representation [Local / Fair-Rep]**: Each client adopts the Fair Linear Representations pre-processing debiasing strategy (He, Burghardt, and Lerman 2020) locally and aggregates using FedAvg.
- **FedAvg + Global reweighting [Global RW]** (Abay et al. 2020): A differential-privacy approach to collect noisy statistics such as the number of samples with privileged attribute values ($A=1$) and favorable labels ($Y=1$) from clients. Server computes global weights based on the collected statistics and shares them with the clients, which assign them to their data samples during FL training².
- **FedFB** (Zeng, Chen, and Lee 2021): An in-processing debiasing approach in FL based on FairBatch (Roh et al. 2021). The server computes new weights for each group based on information from the clients in each round and broadcasts them back to the clients. For fair comparison, we use FedFB that is optimized w.r.t EOD as in FairFed.

²We apply the global reweighting approach in (Abay et al. 2020) without differential-privacy noise in order to compare with the optimal debiasing performance of global reweighting.

	Method	Adult ($\beta = 1$)					COMPAS ($\beta = 1$)				
		Heterogeneity Level α					Heterogeneity Level α				
		0.1	0.2	0.5	10	5000	0.1	0.2	0.5	10	5000
Acc.	FedAvg	0.835	0.836	0.835	0.836	0.837	0.674	0.673	0.675	0.674	0.675
	Local / [Best]	0.831	0.833	0.834	0.831	0.829	0.666	0.659	0.665	0.663	0.664
	Global RW	0.834	0.833	0.831	0.829	0.829	0.673	0.671	0.672	0.676	0.675
	FedFB	0.825	0.825	0.829	0.832	0.832	0.674	0.673	0.675	0.677	0.677
	FairFed / RW	0.830	0.834	0.832	0.829	0.829	0.672	0.670	0.669	0.669	0.673
	FairFed / FairRep	0.824	0.833	0.834	0.834	0.834	0.661	0.655	0.663	0.663	0.660
	FairFed / FairBatch	0.829	0.833	0.830	0.830	0.831	0.659	0.664	0.665	0.661	0.661
EOD	FedAvg	-0.174	-0.173	-0.176	-0.179	-0.180	-0.065	-0.071	-0.067	-0.076	-0.078
	Local / [Best]	0.052	-0.009	-0.006	-0.013	0.014	-0.055	-0.051	-0.054	-0.038	-0.035
	Global RW	-0.030	0.019	0.022	0.017	0.010	-0.060	-0.065	-0.066	-0.076	-0.077
	FedFB	-0.019	0.015	0.015	-0.012	-0.012	-0.062	-0.061	-0.063	-0.077	-0.072
	FairFed / RW	-0.017	0.001	0.018	0.016	0.013	-0.057	-0.065	-0.053	-0.067	-0.061
	FairFed / FairRep	0.023	-0.009	-0.071	-0.174	-0.187	0.037	0.023	0.043	0.046	0.039
	FairFed / FairBatch	-0.020	0.001	0.000	-0.005	-0.004	-0.048	-0.048	-0.049	-0.035	-0.031

Table 1: Performance comparison under different heterogeneity levels α . Smaller α indicates more heterogeneous client distributions. We report the average of 20 random seeds. For EOD, values closer to zero indicate better fairness. For brevity, we report the values achieved by the best local debiasing baseline (without FairFed) as Local / [Best] in the table. An extended version of this table is reported in (Ezzeldin et al. 2021, Appendix C).

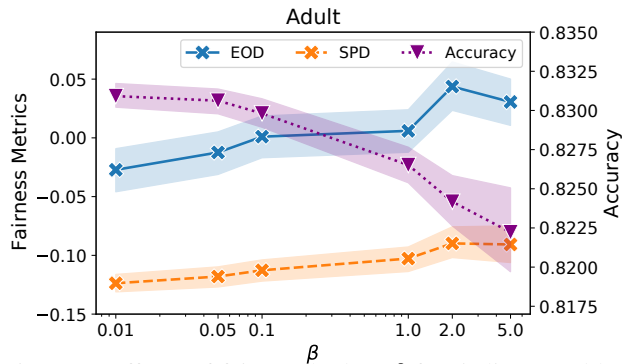


Figure 3: Effects of fairness budget β for 5 clients and heterogeneity $\alpha = 0.2$ on FairFed with local reweighting.

Experimental Results

Performance under heterogeneous sensitive attribute distributions. We compared the performance of FairFed when used with three local debiasing methods (reweighting, FairRep (He, Burghardt, and Lerman 2020) and FairBatch (Roh et al. 2021)) against the baselines described in the previous subsection, under different heterogeneity levels. The results are summarized in Table 1.

FairFed outperforms the baselines at different heterogeneity levels, but at highly homogeneous data distributions (i.e., a large α value), FairFed does not provide significant gains in fairness performance compared to local debiasing methods (except when using FairBatch). This is due to the fact that under homogeneous sampling, the distributions of the local datasets are statistically similar (and reflect the original distribution with enough samples), resulting in similar debiasing effect being applied across all clients when using the pre-processing methods (reweighting and FairRep). For a higher level of heterogeneity (i.e., at lower $\alpha = 0.1$), FairFed can improve EOD in Adult and COMPAS data by

93% and 50%, respectively. This is done at the expense of only a 0.3% decrease in accuracy for both Adult and COMPAS datasets. In contrast, at the same heterogeneity level, local strategies can only improve EOD by 65% and 15% for Adult and COMPAS datasets, respectively. Global reweighting only improves EOD by 73% and 2% for Adult and COMPAS datasets, respectively. On average, across different heterogeneity levels, improving EOD by 87% and 1.5% for Adult and COMPAS, respectively, at $\alpha = 0.1$.

Note, however, that FedFB requires the clients to share explicit information about the performance of the model on each local subgroup in order to update the weights in FairBatch (Zeng, Chen, and Lee 2021), which can potentially leak information about the clients’ local datasets.

Performance with different fairness budgets (β). In FairFed, we introduced a fairness budget parameter β , which controls how much the aggregation weights can change due to fairness adaptation at the server in each round (refer to (6) for the explanation of β). Figure 3 shows the effects of β using heterogeneity level $\alpha = 0.2$. As the value of β increases, the fairness constraint has a bigger impact on the aggregation weights, yielding better fairness (EOD closer to zero) at the cost of a decrease in model accuracy.

Cases Studies for Fair Training in FL

In the previous section, we evaluated FairFed on heterogeneous distributions synthesized from standard benchmark datasets in fair ML. In order to validate the effectiveness of FairFed in FL scenarios with naturally heterogeneous distributions, we consider two FL case studies in this section.

Case Study 1: Predicting Individual Income from US Census across States

In this case study, we use the US Census data to present the performance of our FairFed approach in a distributed learn-

ing application with a natural data partitioning. Our experiments are performed on the ACSIncome dataset (Ding et al. 2021) with the task of predicting whether an individual’s income is above \$50,000 (or not) based on the features collected during the Census which include employment type, education, marital status, etc.

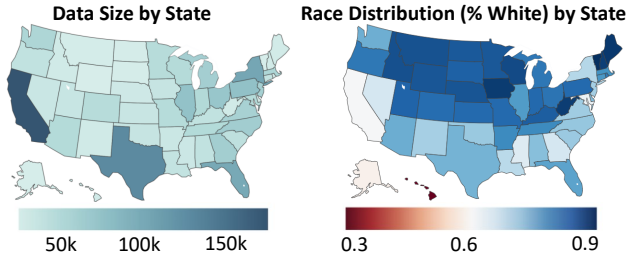


Figure 4: Demographic distribution of ACSIncome dataset.

Data Distribution. ACSIncome dataset is constructed from American Community Survey (ACS) Public Use Microdata Sample (PUMS) over all 50 states and Puerto Rico in 2018 with a total of 1,664,500 datapoints. In our experiments, we treat each state as one participant in the FL system (i.e., 51 participants). Due to the demographic distribution of different states, clients share different data sizes and sensitive attributes distribution. For example, Wyoming has the smallest dataset size with 3,064 users compared to California that has 195,665 users. We choose the race information (white/non-white) of the users as the sensitive attribute of interest in our experiments. Hawaii is the state with the lowest ratio (26%) of white population, while Vermont has the highest ratio (96%) of its dataset as white population. Figure 4 provides a visualization for the data distributions across the different states.

Performance of FairFed. Table 3 compares the performance of FairFed on ACSIncome dataset. Table 3 shows that adopting local reweighting yields worse group fairness performance than simply applying FedAvg (without any debiasing) due to the heterogeneity across states. FairFed with reweighting overcomes this issue and improves the EOD by 20% (-0.062 to -0.050).

Case Study 2: Predicting Daily Stress Level from Wearable Sensor Signals

In this case study, we use the human behavior dataset TILES (Mundnich et al. 2020). *Tracking Individual Performance with Sensors* (TILES) is a 10 weeks longitudinal study with hospital worker volunteers in a large Los Angeles hospital, where 30% of the participants are male and 70% are female. We use this dataset to estimate users’ daily stress levels based on physiological and physical activity signals collected through wearable sensors (e.g., Fitbit). The target is a binary label indicating whether the person’s stress level is above individual average (i.e., 1) or not (i.e., 0), which are collected from daily surveys sent to the participants’ phones.

Client	Size	Gender		Stress	
		F	M	$y = 0$	$y = 1$
RN-day shift	707	82%	18%	45%	55%
RN-night shift	609	77%	23%	57%	43%
CNA	244	61%	39%	65%	35%

Table 2: Data distribution of TILES dataset.

Method	ACSIncome			TILES		
	Acc.	EOD	SPD	Acc.	EOD	SPD
FedAvg	0.800	-0.062	-0.102	0.567	-0.199	-0.166
Local / RW	0.800	-0.066	-0.106	0.567	-0.064	-0.041
FairFed / RW	0.799	-0.050	-0.089	0.556	0.004	0.004

Table 3: Performance on ACSIncome and TILES datasets.

Data Distribution. We focus on the nurse population in the dataset. Each client represents the data from one occupation group – day-shift registered nurse (RN- day shift), night-shift registered nurse (RN-night shift), and certified nursing assistant (CNA). The three clients vary by data size, the distributions of gender (the sensitive attribute) and target stress variable. In general, the client of day-shift registered nurse population has the most datapoints, more female, and higher stress levels. The detailed data distribution of each clients is shown in Table 2.

Performance of FairFed. Table 3 reports the performance on TILES dataset. Both FairFed and local reweighting improve the EOD metric as compared to FedAvg. FairFed improves EOD from -0.199 to 0.004 with only 2.6% accuracy decrease (from 0.567 to 0.556).

Conclusion and Future Works

In this work, motivated by the importance and challenges of group fairness in federated learning, we propose the FairFed algorithm to enhance group fairness via a fairness-aware aggregation method, aiming to provide fair model performance across different sensitive groups (e.g., racial, gender groups) while maintaining high utility. Though our proposed method outperforms the state-of-the-art fair federated learning frameworks under high data heterogeneity, limitations still exist. As such, we plan to further improve FairFed from these perspectives: 1) We report the empirical results on binary classification tasks in this work. We will extend the work to various application scenarios (e.g., regression tasks, NLP tasks); 2) We will extend our study to scenarios of heterogeneous application of different local debiasing methods and understand how the framework can be tuned to incorporate updates from these different debiasing schemes; 3) We focused on group fairness in FL, but we plan to integrate FairFed with other fairness notions in FL, such as *collaborative fairness* and *client-based fairness*. We give an example of how FairFed can provide both group-fairness and client-based fairness in (Ezzeldin et al. 2021, Appendix D), which sets promising preliminary steps for future exploration.

Acknowledgments

This material is based upon work supported by NSF grants 1763673, CNS-2002874, ARO grant W911NF2210165, a gift from Intel via the PrivateAI institute, and gifts from Qualcomm, Cisco, and Konica-Minolta, and support from USC-Amazon Center on trustworthy AI.

References

- Abay, A.; Zhou, Y.; Baracaldo, N.; Rajamoni, S.; Chuba, E.; and Ludwig, H. 2020. Mitigating bias in federated learning. *arXiv preprint arXiv:2012.02447*.
- Bonawitz, K.; Ivanov, V.; Kreuter, B.; Marcedone, A.; McMahan, H. B.; Patel, S.; Ramage, D.; Segal, A.; and Seth, K. 2017. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, 1175–1191.
- Ding, F.; Hardt, M.; Miller, J.; and Schmidt, L. 2021. Retiring Adult: New Datasets for Fair Machine Learning. *Advances in Neural Information Processing Systems*, 34.
- Du, W.; Xu, D.; Wu, X.; and Tong, H. 2021. Fairness-aware Agnostic Federated Learning. In *Proceedings of the 2021 SIAM International Conference on Data Mining (SDM)*, 181–189. SIAM.
- Dua, D.; and Graff, C. 2017. UCI Machine Learning Repository. <http://archive.ics.uci.edu/ml>.
- Dwork, C.; Hardt, M.; Pitassi, T.; Reingold, O.; and Zemel, R. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, 214–226.
- Ezzeldin, Y. H.; Yan, S.; He, C.; Ferrara, E.; and Avestimehr, S. 2021. Fairfed: Enabling group fairness in federated learning. *arXiv preprint arXiv:2110.00857*.
- Feldman, M.; Friedler, S. A.; Moeller, J.; Scheidegger, C.; and Venkatasubramanian, S. 2015. Certifying and removing disparate impact. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 259–268. ACM.
- Gálvez, B. R.; Granqvist, F.; van Dalen, R.; and Seigel, M. 2021. Enforcing fairness in private federated learning via the modified method of differential multipliers. In *NeurIPS 2021 Workshop Privacy in Machine Learning*.
- Grgić-Hlača, N.; Zafar, M. B.; Gummadi, K. P.; and Weller, A. 2018. Beyond distributive fairness in algorithmic decision making: Feature selection for procedurally fair learning. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Hardt, M.; Price, E.; Srebro, N.; et al. 2016. Equality of opportunity in supervised learning. In *Advances in neural information processing systems*, 3315–3323.
- He, C.; Li, S.; So, J.; Zeng, X.; Zhang, M.; Wang, H.; Wang, X.; Vepakomma, P.; Singh, A.; Qiu, H.; et al. 2020. Fedml: A research library and benchmark for federated machine learning. *arXiv preprint arXiv:2007.13518*.
- He, Y.; Burghardt, K.; and Lerman, K. 2020. A geometric solution to fair representations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 279–285.
- Hsu, H.; Qi, H.; and Brown, M. 2019. Measuring the Effects of Non-Identical Data Distribution for Federated Visual Classification. *arXiv preprint arXiv:1909.06335*.
- Kairouz, P.; McMahan, H. B.; Avent, B.; Bellet, A.; Bennis, M.; Bhagoji, A. N.; Bonawitz, K.; Charles, Z. B.; Cormode, G.; Cummings, R.; D’Oliveira, R. G. L.; Rouayheb, S. Y. E.; Evans, D.; Gardner, J.; Garrett, Z.; Gascón, A.; Ghazi, B.; Gibbons, P. B.; Gruteser, M.; Harchaoui, Z.; He, C.; He, L.; Huo, Z.; Hutchinson, B.; Hsu, J.; Jaggi, M.; Javidi, T.; Joshi, G.; Khodak, M.; Konečný, J.; Korolova, A.; Koushanfar, F.; Koyejo, O.; Lepoint, T.; Liu, Y.; Mittal, P.; Mohri, M.; Nock, R.; Özgür, A.; Pagh, R.; Raykova, M.; Qi, H.; Ramage, D.; Raskar, R.; Song, D. X.; Song, W.; Stich, S. U.; Sun, Z.; Suresh, A. T.; Tramèr, F.; Vepakomma, P.; Wang, J.; Xiong, L.; Xu, Z.; Yang, Q.; Yu, F. X.; Yu, H.; and Zhao, S. 2021. Advances and Open Problems in Federated Learning. *Found. Trends Mach. Learn.*, 14: 1–210.
- Kamiran, F.; and Calders, T. 2012. Data preprocessing techniques for classification without discrimination. *Knowledge and Information Systems*, 33(1): 1–33.
- Kamishima, T.; Akaho, S.; Asoh, H.; and Sakuma, J. 2012. Fairness-aware classifier with prejudice remover regularizer. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 35–50. Springer.
- Kim, M. P.; Ghorbani, A.; and Zou, J. 2019. Multiaccuracy: Black-box post-processing for fairness in classification. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 247–254. ACM.
- Larson, J.; Mattu, S.; Kirchner, L.; and Angwin, J. 2016. How we analyzed the COMPAS recidivism algorithm. *ProPublica (5 2016)*, 9.
- Li, T.; Sanjabi, M.; Beirami, A.; and Smith, V. 2019. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*.
- Lohia, P. K.; Ramamurthy, K. N.; Bhide, M.; Saha, D.; Varshney, K. R.; and Puri, R. 2019. Bias mitigation post-processing for individual and group fairness. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2847–2851. IEEE.
- Lyu, L.; Xu, X.; Wang, Q.; and Yu, H. 2020. Collaborative fairness in federated learning. In *Federated Learning*, 189–204. Springer.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, 1273–1282. PMLR.
- Mohri, M.; Sivek, G.; and Suresh, A. T. 2019. Agnostic federated learning. In *International Conference on Machine Learning*, 4615–4625. PMLR.
- Mundnich, K.; Booth, B. M.; l’Hommedieu, M.; Feng, T.; Girault, B.; L’hommedieu, J.; Wildman, M.; Skaaden, S.; Nadarajan, A.; Villatte, J. L.; et al. 2020. TILES-2018, a longitudinal physiologic and behavioral data set of hospital workers. *Scientific Data*, 7(1): 1–26.

- Papadaki, A.; Martinez, N.; Bertran, M.; Sapiro, G.; and Rodrigues, M. 2021. Federating for Learning Group Fair Models. *arXiv preprint arXiv:2110.01999*.
- Reddi, S.; Charles, Z.; Zaheer, M.; Garrett, Z.; Rush, K.; Konečný, J.; Kumar, S.; and McMahan, H. B. 2020. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*.
- Roh, Y.; Lee, K.; Whang, S. E.; and Suh, C. 2021. Fair-Batch: Batch Selection for Model Fairness. In *International Conference on Learning Representations*.
- Wang, J.; Charles, Z.; Xu, Z.; Joshi, G.; McMahan, H. B.; Al-Shedivat, M.; Andrew, G.; Avestimehr, S.; Daly, K.; Data, D.; et al. 2021a. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*.
- Wang, J.; Liu, Q.; Liang, H.; Joshi, G.; and Poor, H. V. 2020. Tackling the objective inconsistency problem in heterogeneous federated optimization. *arXiv preprint arXiv:2007.07481*.
- Wang, Z.; Fan, X.; Qi, J.; Wen, C.; Wang, C.; and Yu, R. 2021b. Federated Learning with Fair Averaging. In *IJCAI*.
- Zeng, Y.; Chen, H.; and Lee, K. 2021. Improving fairness via federated learning. *arXiv preprint arXiv:2110.15545*.
- Zhang, B. H.; Lemoine, B.; and Mitchell, M. 2018. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, 335–340. ACM.
- Zhang, D. Y.; Kou, Z.; and Wang, D. 2020. Fairfl: A fair federated learning approach to reducing demographic bias in privacy-sensitive classification models. In *2020 IEEE International Conference on Big Data (Big Data)*, 1051–1060. IEEE.