ROBUST AND PERSONALIZED FEDERATED LEARNING WITH SPURIOUS FEATURES: AN ADVERSARIAL APPROACH

Anonymous authors
Paper under double-blind review

ABSTRACT

A common approach for personalized federated learning is fine-tuning the global machine learning model to each local client. While this addresses some issues of statistical heterogeneity, we find that such personalization methods are often vulnerable to *spurious features*, leading to bias and diminished generalization performance. However, debiasing the personalized models under spurious features is difficult. To this end, we propose a strategy to mitigate the effect of spurious features based on our observation that the global model in the federated learning step has a low accuracy disparity due to statistical heterogeneity. Then, we estimate and mitigate the *accuracy disparity* of personalized models using the global model and *adversarial transferability* in the personalization step. Empirical results on MNIST, CelebA, and Coil20 datasets show that our method reduces the accuracy disparity of the personalized model on the bias-conflicting data samples from 15.12% to 2.15%, compared to existing personalization approaches, while preserving the benefit of enhanced average accuracy from fine-tuning.

1 INTRODUCTION

Federated learning (FL) is a leading framework for clients to collaboratively train a shared global machine learning (ML) model without releasing their local private datasets (McMahan et al., 2017; Kairouz et al., 2019). The jointly trained global model could be further fine-tuned on each client's local dataset to produce personalized models (Fallah et al., 2020; T. Dinh et al., 2020; Li et al., 2021). While existing theoretical and empirical results highlight how personalized models improve accuracy on local data, few works consider what features the personalized models learn from the local dataset. Our motivating hypothesis is that *not all local features are beneficial*.

For example, consider a gender prediction task using face images, where the ML model learns to predict gender based on hair color because females are more likely to have blond hair (Sagawa et al., 2020). In this case, the hair color is called a *spurious feature* because it only statistically correlates with the gender on biased samples but not necessarily on the overall population. Thus, the accuracy of a model that relies on spurious features such as hair color is likely to drop significantly on the data where the spurious correlation does not hold, e.g., for blond male images (Sagawa et al., 2020). More broadly, the *accuracy disparity* caused by spurious features leads to issues in both fairness (McNamara et al., 2019; Zhao & Gordon, 2019; Agarwal et al., 2019; Chi et al., 2021), i.e., racial bias (Khani & Liang, 2021) and robustness, i.e., accuracy decrease under distribution shift (Zhao et al., 2019; Koh et al., 2021). Compared to the global model, because of the local fine-tuning, the personalized models are more vulnerable to spurious features and have a larger accuracy disparity.

Empirically, we observe that the typical non-i.i.d. data distributions in FL settings reduce the accuracy disparity of the global model. One potential explanation is that the distributional shift under spurious features across users is larger than that of non-spurious features. As a result, the aggregation operation in the central server averages out the diverse shifts under the spurious features (Figure 1) so that the global model becomes more robust against spurious features on the benchmark datasets. On the other hand, the local fine-tuning step will result in a personalized model that entangles spurious features and becomes biased. Hence, it remains a key challenge how to debias the personalized models.

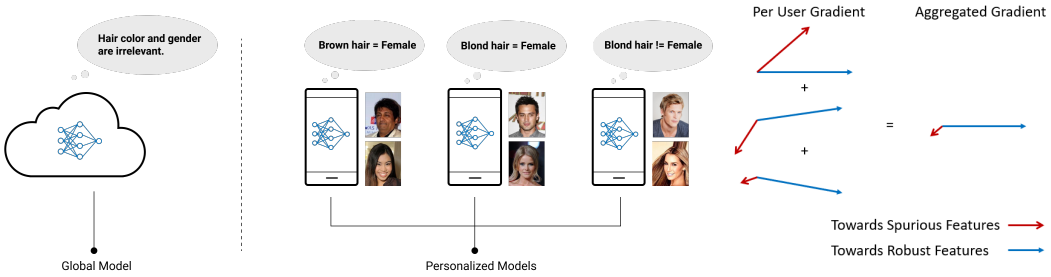

(a) Spurious Correlation Varies across Users

(b) Gradient Divergence between Users

Figure 1: The variation of spurious features leads to diverse gradients. The variation leads to a global ML model with a low accuracy disparity in a federated learning setting because the aggregation step in the central server averages out the gradients resulted from local spurious features.

Various methods (Wang et al., 2019; Sagawa et al., 2020; Liu et al., 2021) have been developed to disentangle spurious features from ML models, but few can be easily applied to personalized models under the federated learning setting. The main reason lies in their reliance on bias-conflicting samples. The bias-conflicting samples can be rare. For example, in the CelebA dataset, only around 1.7k samples are bias-conflicting out of more than170k samples. If we distribute the CelebA dataset across users, nearly half do not have any bias-conflicting samples. Prior work (Li & Wang, 2019) has tried using a global proxy dataset for training FL models. However, a global proxy dataset may not characterize the local samples well. Furthermore, estimating the accuracy disparity of the personalized model becomes difficult without access to bias-conflicting samples.

To approach this problem, we propose a novel method to reduce the accuracy disparity of personalized models. Our proposed method does not rely on the bias-conflicting samples, e.g., blond male images in the aforementioned task of gender prediction, which are often rare and may not be available for every user. Instead, inspired by prior works (Tramèr et al., 2017; Liang et al., 2021) on the transferability of adversarial examples, we use the global model as a reference and the *adversarial transferability* between the global model and the personalized models as a proxy to estimate the accuracy disparity of the personalized models. The intuition is that if two ML models use disjoint subsets of features, the adversarial examples that one ML model generates do not transfer to the other ML model. Based on this intuition, we propose the following hypothesis:

> *If the personalized models entangle spurious features (thus increasing the accuracy disparity), the adversarial examples generated by the global model (that uses non-spurious features) do not transfer to the personalized models.*

Empirically, we show that the adversarial transferability between the global and personalized models strongly correlates with the accuracy disparity between the global and personalized models (Section 4), validating our hypothesis. Based on the empirical observations, we develop a method that enforces adversarial transferability between the global and the personalized models to reduce the accuracy disparity of personalized models (Section 5). Besides entangling spurious features, we find that forgetting bias-conflicting during personalization also results in accuracy disparity. To address the forgetting issue, we further propose a regularization method based on loss function approximation (Section 5). Combining adversarial transferability and our regularization method that mitigates forgetting, we can reduce the accuracy disparity of personalized models. Our contributions are summarized as follows:

- We empirically evaluate the accuracy disparity of the global and personalized models in a federated learning setting with spurious features, highlighting a risk of existing personalization methods.
- We design a method to estimate the accuracy disparity of the personalized models, based on the low accuracy disparity global model and the adversarial transferability between the global and personalized models.

- We develop two methods to reduce the accuracy disparity of personalized models by (1) enforcing the adversarial transferability between the global and personalized models and (2) preventing the personalized model from forgetting bias-conflicting samples, which are rare or missing in the local dataset.

Empirically, we conduct extensive experiments to validate the effectiveness of the proposed methods in mitigating accuracy disparity under the FL setting. Our experiments on MNIST (Deng, 2012), CelebA (Liu et al., 2015), and Coil20 (Nene et al., 1996) datasets show that the proposed approach reduces the accuracy disparity of personalized models from 15.12% to 2.15%, which is closer to that of the global model (-0.63%). Our method also preserves the benefit of the enhanced average accuracy from fine-tuning, resulting in 3.43% accuracy improvement on the biased test set and 0.85% accuracy improvement on the biased-conflicting test set.

2 Related Work

Personalized Federated Learning Fine-tuning is typical for personalization methods. The meta-learning-based method first trains a global model and fine-tunes the global model locally (Fallah et al., 2020). Other methods using multi-task learning (Li et al., 2021) or Moreau envelopes (T. Dinh et al., 2020) have an interpretation as fine-tuning the local model along with training the global model. Fine-tuning is also compatible with clustering-based method (Ghosh et al., 2020).

Debiasing ML Models A few prior works (Li & Vasconcelos, 2019; Sagawa et al., 2020) utilize group labels, which might require human annotation, to debias ML models. For example, the group distributional robust optimization (DRO) method (Sagawa et al., 2020) aims to optimize the worst-case error rate of ML models across different (often manually annotated) groups. Some groups contain bias-conflicting samples while others do not. Residual learning-based methods (He et al., 2019; Nam et al., 2020; Liu et al., 2021) train a biased ML model and up-weight the residual, which mainly contains bias-conflicting samples that the biased ML model mis-predicts. Chi et al. (2021) aims to mitigate the accuracy disparity in regression problems via learning the appropriate representations. However, all these methods rely on the explicit access to the bias-conflicting samples, making them difficult to apply on personalized federated learning, where bias-conflicting samples may not be accessible for every client.

3 Preliminaries

Definitions and Notation A data sample is a vector $\boldsymbol{x} = [\boldsymbol{x}_r, \boldsymbol{x}_s]$, where $\boldsymbol{x}_r$ corresponds to the robust and non-spurious features and $\boldsymbol{x}_s$ are the spurious features. Let y be the label, and define $\ell : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ to be the loss function. $\boldsymbol{w}_g$ and $\boldsymbol{w}_p$ are the weights for the global model $f_g : \mathcal{X} \to \mathcal{Y}$ and personalized model $f_p : \mathcal{X} \to \mathcal{Y}$, respectively. Let $\mathrm{supp}(\mathcal{D})$ be the support of distribution $\mathcal{D}$. We define a global data distribution $\mathcal{D}_g$, a biased local data distribution $\mathcal{D}_b$, and assume a bias-conflicting local data distribution $\mathcal{D}_{bc}$. In the personalization step, we only have access to $\mathcal{D}_b$. We define the "pseudo-gradient" as the difference between the updated local model and the global model from the previous round (sometimes we will use the term "gradient" when it is clear from context).

Training and Personalization Methods We train the global model using the federated averaging algorithm (McMahan et al., 2017), which learns a model $f : \mathcal{X} \to \mathcal{Y}$ that minimizes: $\mathcal{L}(f, \mathcal{D}_g) = \sum_{i=1}^{N} |\mathcal{D}_{b_i}|/|\mathcal{D}_g| \cdot \mathbb{E}_{\boldsymbol{x},y\sim\mathcal{D}_{b_i}} \ell(f(\boldsymbol{x}), y)$, where N is the number of clients, $\mathcal{D}_{b_i}$ is the biased local dataset for client i, and $\mathcal{D}_g = \cup_{i=1}^{N} \mathcal{D}_{b_i}$ is the global dataset. When the global model f_g converges, f_g will be sent to local clients for further fine-tuning by minimizing $\mathbb{E}_{\boldsymbol{x},y\sim\mathcal{D}_{b_i}} \ell(f(\boldsymbol{x}), y)$.

Adversarial Examples and Transferability We can generate an adversarial example $\boldsymbol{x}_{adv}$ given data sample $\boldsymbol{x}$ with label y by solving:

$$\boldsymbol{x}_{adv} = \underset{\|\boldsymbol{x}'-\boldsymbol{x}\|\leq\epsilon}{\arg\max}\ \ell(f(\boldsymbol{x}'), y), \tag{1}$$

(a) MNIST (b) CelebA (c) Coil20

Figure 2: Datasets with spurious features. The object color spuriously correlates with the label in MNIST (a) and Coil20 (c) datasets. The hair color spuriously correlates with gender in the CelebA dataset (b).

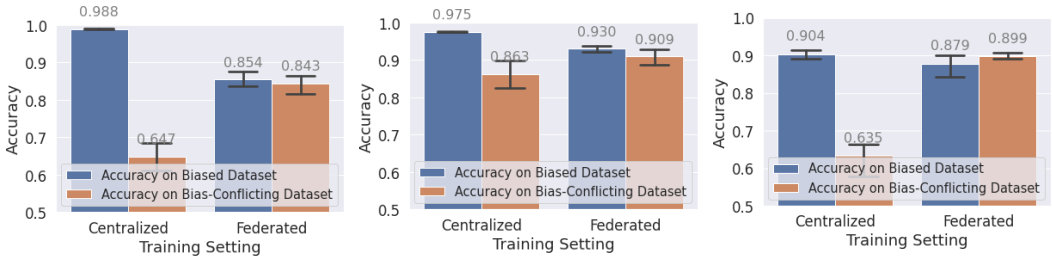

(a) MNIST (b) CelebA (c) Coil20

Figure 3: The accuracy of ML models on Biased Dataset and Bias-Conflicting Dataset under centralized and federated training settings. In the centralized setting, an ML model is trained by a single dataset that contains all the samples with a fixed spurious correlation. The global models in the federated setting achieve smaller accuracy disparities between biased and bias-conflicting datasets.

where f is the victim ML model, and ϵ is the attack budget. In this example, we consider the L_2 attacks, but extensions to general L_p attacks are straightforward. We say an adversarial example to be *transferable* if it also fools another ML model (e.g., a personalized model) other than the original victim model f (e.g., the global model).

4 An Empirical Study with Spurious Features

To gain some insights into the problem, we first perform an empirical study on the accuracy disparity of the global and personalized models in an FL setting. Our results highlight the risk of existing personalization methods and the difficulty of mitigating the risk. We also highlight the correlation between the adversarial transferability and the accuracy disparity between the global and personalized models. The spurious features in the empirical study are as follows.

Spurious Features We consider color as the spurious feature for the MNIST, CelebA, and Coil20 datasets. In the MNIST and Coil20 datasets, we manually color the objects according to their labels to create spurious correlations, as Figure 2 shows. The spurious correlations vary across clients for the MNIST and Coil20 data (e.g., the red color correlates with label zero on the first client and with label one on the second client) to create additional statistical heterogeneity. In the CelebA dataset, the hair color attribute correlates with the gender label. We assign disjoint subsets of celebrities to different users, which naturally increases statistical heterogeneity for the spurious correlation.

Bias-aligned and Bias-conflicting Samples For a given spurious correlation, the bias-aligned samples are the samples on which the spurious correlation holds, and the rest of the samples are bias-conflicting. For example, in CelebA dataset, blond female images are bias-aligned, and blond male images are bias-conflicting. A biased test set contains mainly bias-aligned samples, and a bias-conflicting test set only contains bias-conflicting samples.

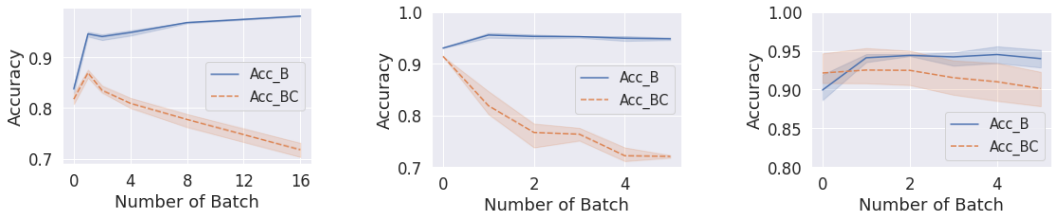

(a) MNIST (b) CelebA (c) Coil20

Figure 4: The accuracy of personalized model on Biased Dataset (Acc_B) and Bias-Conflicting Dataset (Acc_BC) with increasing fine-tuning batches. The personalized models entangle spurious features and increase accuracy disparities between biased and bias-conflicting datasets.

4.1 Statistical Diversity Reduces Accuracy Disparity

Figure 3 shows the accuracy disparity of ML models on bias and bias-conflicting test sets. Compared to the models trained in the centralized setting, where the spurious correlations are fixed, the accuracy disparity of models trained in the federated setting decreased significantly. These empirical results suggest that the global model in FL is more robust to spurious features if the spurious features are non-i.i.d. across clients.

To explain this observation, consider the relationship between the gradient directions and the learned features. For the spurious features, its correlation with the label may change across clients. As an example, in some users' local datasets, the blond hair does not correlate with the gender label because the dataset does not contain any blond female image or the dataset has blond male images, as visualized in Figure 1a. Therefore, the gradient directions for the spurious features are diverse across clients, as shown in Figure 1b. The divergence between the gradient directions, as a consequence, makes learning spurious features difficult. In contrast, the non-spurious features, e.g., shape features, are more consistent across clients, leading to a more consistent gradient direction.

4.2 Personalization May Exacerbate Accuracy Disparity

Although the global model in the federated setting has a lower accuracy disparity than in the centralized setting, the advantage could vanish during the personalization step.

The Personalized Model Entangles Spurious Features As can be observed in Figure 4, the accuracy first increases and decreases on the MNIST bias-conflicting test set and slowly decreases on Coil20. These two observations indicate that the personalized model entangles spurious features and exacerbates the accuracy disparity in a few batches. Although in principle, one may resort to early stopping, this is not feasible when the bias-conflicting dataset is unavailable or scarce.

4.3 Adversarial Transferability Indicates Accuracy Disparity

Because the bias-conflicting test set is often unavailable, it is infeasible to directly measure the accuracy disparity across personalized models. To this end, in this section, we focus on methods that implicitly measure the accuracy disparity and determine whether the personalized models entangle spurious features. Following our hypothesis in Section 1, we consider using the adversarial transferability between the global and personalized models as a proxy for the accuracy disparity measurement. Figure 5 plots the accuracy disparity and the adversarial transferability during fine-tuning. As the accuracy disparity of the personalized models increases and drifts away from that of the global model, the adversarial transferability between the global and personalized models decreases. This result empirically validates our hypothesis.

Running Example The example in Figure 6 intuitively explains why entangling spurious features decreases adversarial transferability. We consider a setting where a data sample $\boldsymbol{x} = [x_r, x_s]$ has a robust, non-spurious feature x_r, a spurious feature x_s, and a global model f_g only uses x_r. Figure 6a shows an adversarial perturbation $\boldsymbol{\delta}_{f_g} = \boldsymbol{x}_{adv} - \boldsymbol{x}$ whose direction is perpendicular to the decision

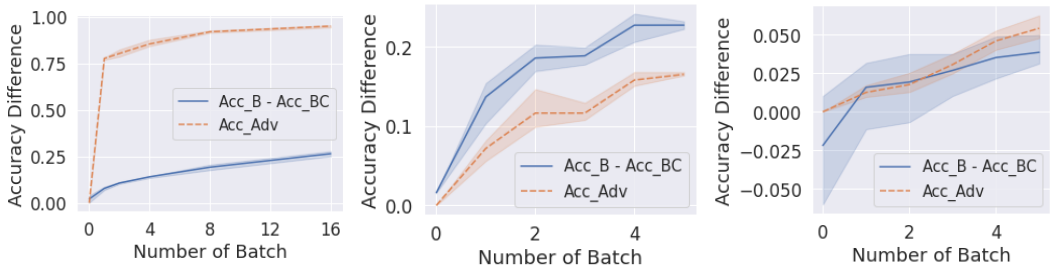

(a) MNIST (b) CelebA (c) Coil20

Figure 5: The accuracy disparity of personalized models on biased dataset and bias-conflicting dataset (Acc_B - Acc_BC) and their accuracy on adversarial examples (Acc_Adv). As the personalized models entangle spurious features and increase the accuracy disparity, the accuracy of the personalized models on adversarial examples increases, which indicates the adversarial transferability between the global and personalized models decreases.

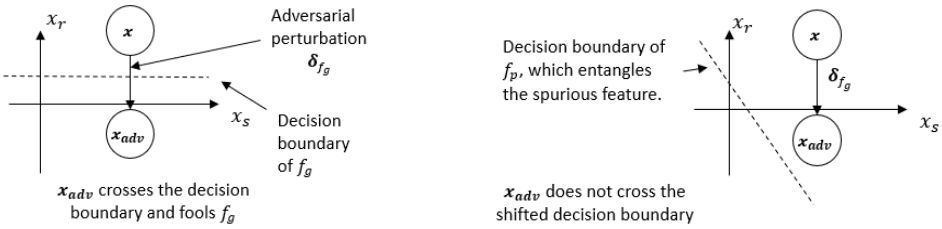

(a) Unbiased Global Model f_g (b) Biased Personalized Model f_p

Figure 6: Entangling spurious features reduces adversarial transferability. Entangling spurious feature x_s shifts the decision boundary. Adversarial example $\boldsymbol{x}_{adv}$ generated using f_g does not transfer to f_p because the shift makes the adversarial perturbation $\boldsymbol{\delta}_{f_g} = \boldsymbol{x}_{adv} - \boldsymbol{x}$ less effective.

boundary of f_g. If the personalized model f_p entangles the spurious feature, then the entanglement could change the decision boundary so that the original adversarial example is no longer adversarial, i.e., it does not transfer to the updated decision boundary, as shown in Figure 6b.

5 METHODS

In Section 4.3, we see that low adversarial transferability between the global and local models indicates high accuracy disparity of the personalized models. Therefore, it is natural to ask if enforcing the adversarial transferability in the personalization step reduces the accuracy disparity. In this section, we first introduce adversarial examples to the personalization step, aiming to *enforce the adversarial transferability*. However, the accuracy disparity still increases, albeit much slower, even if the adversarial transferability remains high. One possible reason is that the personalized model forgets the bias-conflicting samples, which are rare or missing. Specifically, models are known to be susceptible to catastrophic-forgetting – the phenomenon where a neural network forgets the previous task while learning a new task (Kirkpatrick et al., 2017; Adel et al., 2019). To this end, we introduce another method based on loss function approximation (Yin et al., 2020) to *mitigate forgetting*. Combing these two methods addresses the accuracy disparity. We implement the two methods as regularization terms added to the original loss function. Both methods are relatively light-weight from a computational perspective.

5.1 ENFORCING ADVERSARIAL TRANSFERABILITY

We enforce that global and personalized models make consistent predictions on adversarial examples, such that adversarial examples transfer from one to the other.

Generating Adversarial Examples The projected gradient descent (PGD) attack (Madry et al., 2018) is an effective attack method that uses the neural network's first-order gradient, and is easy to compute. The attack solves $\boldsymbol{x}_{adv} = \arg\max_{\|\boldsymbol{x}'-\boldsymbol{x}\|\leq\epsilon} \ell(f(\boldsymbol{x}'), y)$ iteratively. At iteration $t+1$, the adversarial example is:

$$\boldsymbol{x}_{adv}^{t+1} = \mathrm{Proj}_{\|\boldsymbol{x}_{adv}-\boldsymbol{x}\|\leq\epsilon}(\boldsymbol{x} + \alpha \cdot \mathrm{sign}(\nabla_{\boldsymbol{x}_{adv}^t} \ell(f_g(\boldsymbol{x}_{adv}^t), y))), \tag{2}$$

where Proj is a projection operator.

Enforcing Consistent Predictions Both the global model f_g and the personalized model f_p take the adversarial example $\boldsymbol{x}_{adv}$ as input and output $\boldsymbol{z}_g$ and $\boldsymbol{z}_p$ from their last layers, respectively. We enforce the adversarial transferability by adding the following regularization term, which maximizes the cross-entropy between $\boldsymbol{z}_g$ and $\boldsymbol{z}_p$. Since the global model f_g is fixed as a reference in the personalization step and its low accuracy disparity is desirable, we use $\boldsymbol{z}_g$ as the ground-truth:

$$R_{adv}(\boldsymbol{z}_g, \boldsymbol{z}_p) = \sum_{i=1}^{K} [z_{g_i} \cdot \log(z_{p_i}) + (1 - z_{g_i}) \cdot \log(1 - z_{p_i})], \tag{3}$$

where K is the number of classes. The local model has access to the global model, so there is no additional communication overhead for implementing this regularization.

5.2 Mitigating Forgetting

We develop a lightweight method to address the forgetting issue based on recent advances in continual learning (Yin et al., 2020). Many existing methods are computationally and memory expensive because they need to, for example, construct a memory bank (Lopez-Paz & Ranzato, 2017). Our idea is to first approximae the loss function $\mathcal{L}(f_p, \mathcal{D}_g)$ of the personalized model f_p on the global dataset $\mathcal{D}_g$ that contains more bias-conflicting samples. Then, we could penalize the personalized model if the approximated loss increases. To this end, we assume that the loss function is λ-smooth, i.e., $\forall \boldsymbol{w}, \boldsymbol{w}'$, we have $\|\nabla\mathcal{L}(f(\boldsymbol{w}), \mathcal{D}_g) - \nabla\mathcal{L}(f(\boldsymbol{w}'), \mathcal{D}_g)\| \leq \lambda\|\boldsymbol{w} - \boldsymbol{w}'\|$, for some constant $\lambda > 0$.

Loss Function Approximation Direct estimation of the loss $\mathcal{L}(f_p, \mathcal{D}_g)$ is intractable because we do not have access to $\mathcal{D}_g$ on the client. Therefore, we apply a second-order Taylor expansion at the global model f_g with parameter $\boldsymbol{w}_g$ to approximate the loss function:

$$\begin{aligned}\mathcal{L}(f_p, \mathcal{D}_g) &= \mathcal{L}(f_g, \mathcal{D}_g) + (\boldsymbol{w}_p - \boldsymbol{w}_g)^\top \nabla_{\boldsymbol{w}}\mathcal{L}(f_g, \mathcal{D}_g) + \frac{1}{2}(\boldsymbol{w}_p - \boldsymbol{w}_g)^\top \nabla^2_{\tilde{\boldsymbol{w}}}\mathcal{L}(f_g, \mathcal{D}_g)(\boldsymbol{w}_p - \boldsymbol{w}_g) \\ &\leq \mathcal{L}(f_g, \mathcal{D}_g) + (\boldsymbol{w}_p - \boldsymbol{w}_g)^\top \nabla_{\boldsymbol{w}}\mathcal{L}(f_g, \mathcal{D}_g) + \frac{\lambda}{2}\|\boldsymbol{w}_p - \boldsymbol{w}_g\|^2,\end{aligned}$$

where $\boldsymbol{w}_g$ and $\boldsymbol{w}_p$ are the weights for global and personalized model, respectively, and $\tilde{\boldsymbol{w}}$ is a linear interpolation between these two, by the Lagrange's mean-value theorem. Since we know that the global model converges to some minima, by the first-order optimality condition, we have $\nabla_{\boldsymbol{w}}\mathcal{L}(f_g, \mathcal{D}_g) = \mathbf{0}$. Furthermore, by the smoothness assumption on the loss function, we know the spectral norm of the Hessian, $\|\nabla^2_{\tilde{\boldsymbol{w}}}\mathcal{L}(f_g, \mathcal{D}_g)\|$, is upper bounded by λ (Nesterov, 2003, Proof of Theorem 2.1.5). Combining all the above arguments, we have

$$\mathcal{L}(f_p, \mathcal{D}_g) - \mathcal{L}(f_g, \mathcal{D}_g) \leq \frac{\lambda}{2} \cdot \|\boldsymbol{w}_g - \boldsymbol{w}_p\|^2. \tag{4}$$

The upper bound in Eq. (4) implies that, by adding a simple L_2 regularization term $R_{L_2} = \lambda \cdot \|\boldsymbol{w}_g - \boldsymbol{w}_p\|^2$, we can mitigate the forgetting issue. Although prior works (Li et al., 2020; T. Dinh et al., 2020; Li et al., 2021) have explored similar regularization methods, we derive the regularization term from a different perspective.

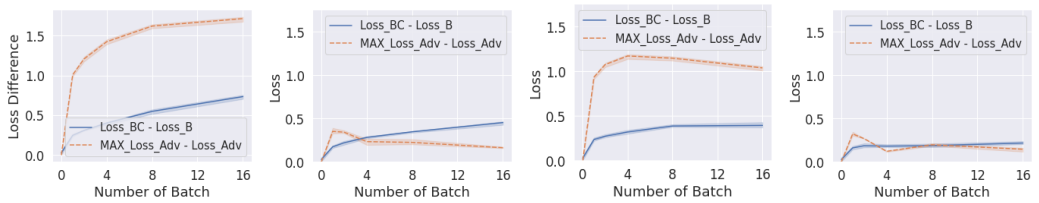

(a) Naive Fine-tuning (b) Enforcing Adversarial Transferability (c) Mitigating Forgetting (d) Applying Both Methods

Figure 7: The loss disparity (Loss_BC - Loss_B) and adversarial transferability (MAX_Loss_Adv - Loss_Adv) with different methods. Combining the two proposed methods addresses the loss disparity. Loss_BC and Loss_B are the losses on the bias-conflicting and biased test sets, respectively. Loss_Adv is the loss on adversarial examples. MAX_Loss_Adv - Loss_Adv measures the decrease of Loss_Adv, which indicates the decrease of adversarial transferability.

6 Experiments

This section presents our experimental results, demonstrating that our method reduces the accuracy disparity. We also show that the benefit of enhanced average accuracy from fine-tuning is preserved.

6.1 Setting

Data Partition We distribute the MNIST and Coil20 dataset across 50 clients. Each client has data from 5 different classes. The local dataset on each client is further partitioned to train/validation/test set with a ratio of 72:8:20, following prior work (Li et al., 2021). The test set here is biased. We alternate the spurious features in biased test sets by recoloring the data to create a bias-conflicting test set. For the CelebA dataset, we consider two partitions. In the first partition, each client represents 20 celebrities. One celebrity only appears on one client. The blond male images in the biased test sets are copied to bias-conflicting test sets. We use all clients (508) for training the global model. In the personalization step, we select the clients who have more than five blond female training samples and more than five blond male test samples. We select these clients because they provide enough samples to create spurious correlations and bias-conflicting test sets. Although the first partition on CelebA is real, the number (161) of blond male images is small. To make the result clearer, we create another synthetic CelebA partition. In the second partition, there are 650 blond male images, which achieve a similar bias-conflicting test set size as prior works Sagawa et al. (2020); Liu et al. (2021). The 650 images are distributed to 3 clients with 2350 other images. The rest of the images are distributed in the same way as the first partition. In the synthetic CelebA partition, there are 650 blond male images, which achieve a similar bias-conflicting test set size as prior works Sagawa et al. (2020); Liu et al. (2021). The 650 images are distributed to 3 clients with 2350 other images. The rest of the images are distributed in the same way as the first partition. Tables 2, 3, and 4 in Appendix A.1 provide more details about the 3 clients.

Hyper-parameters We use Adam optimizer (Kingma & Ba, 2015) through our experiments with learning rate 1e-4. Although stochastic gradient descent (SGD) optimizer is more common in vision-related tasks, we find that the Adam optimizer always leads to lower accuracy disparity. We train the global model for 500 rounds. 5 clients are selected per round, and each performs 5 epochs of local updates. We tune the coefficients of adversarial transferability regularizer and loss approximating regularizer from $\{0.01, 0.1, 1.0, 10.0\}$ and select the largest value that does not decrease the validation accuracy during penalization. We start the attack budget at 0.031 and gradually decrease it such that $30\% - 50\%$ of the attack succeeds. A large attack budget will make the attack too strong and push the adversarial examples far across the decision boundary, making the regularization method less effective. We configure ϵ to 0.031/0.01/0.015 for MNIST/CelebA/Coil20, respectively. We fine-tune the global model for 96/40/30 batches on MNIST/CelebA/Coil20 datasets. Note that we may not select the personalized model with the most fine-tuning batches for reporting. In the just train twice (JTT) method, we up-sample the residuel by a factor of 50. In Ditto, we tune its λ from $\{0.1, 1.0\}$.

Table 1: Accuracy of personalized Models on Biased Test Set (Acc_B) and Bias-Conflicting Test Set (Acc_B). Our proposed method achieves the lowest accuracy disparity (2.15%) compared to other personalization methods (15.12%/15.38%), and 3.43% accuracy improvement on the biased test set and 0.85% improvement on the biased-conflicting test set compared to the global model.

Method	MNIST		CelebA_S		CelebA_R		Coil20	
	Acc_B	Acc_BC	Acc_B	Acc_BC	Acc_B	Acc_BC	Acc_B	Acc_BC
Global	.852 ± 2e-4	.847 ± 6e-4	.930 ± 5e-5	.910 ± 3e-4	.909 ± 6e-5	.929 ± 5e-5	.882 ± .6e-4	.903 ± 7e-4
FT	.989 ± 6e-7	.704 ± 3e-4	.952 ± 6e-5	.786 ± 5e-4	.963 ± 6e-6	.849 ± 1e-3	.931 ± 1e-4	.891 ± 2e-4
Ditto	.982 ± 3e-6	.724 ± 1e-3	.948 ± 5e-5	.715 ± 5e-4	.966 ± 1e-5	.884 ± 2e-4	.939 ± 4e-5	.897 ± 3e-4
UW	.968 ± 2e-5	.823 ± 7e-4	.930 ± 7e-6	.889 ± 5e-4	.936 ± 1e-5	.895 ± 4e-4	N/A	N/A
JTT	.985 ± 2e-7	.707 ± 6e-5	.952 ± 2e-6	.817 ± 3e-4	.956 ± 2e-5	.836 ± 1e-3	N/A	N/A
Ours	.951 ± 2e-5	.870 ± 8e-4	.932 ± 3e-5	.910 ± 2e-4	.925 ± 2e-5	.927 ± 8e-5	.901 ± 3e-4	.916 ± 5e-8

6.2 The Effectiveness of Proposed Methods

We conduct an ablation study on the MNIST dataset. Figure 7b demonstrates the effectiveness of enforcing adversarial transferability. Compared to naive fine-tuning, which is reported in Figure 7a, enforcing adversarial transferability mitigates the loss disparity by $\sim 50\%$. However, the loss disparity still increases during personalization, which is potentially caused by forgetting bias-conflicting samples. Applying the $L2$ regularization term that is derived through loss function approximation and enforcing adversarial transferability addresses the loss disparity as Figure 7d shows. Figure 7c furthers show that applying the $L2$ regularization term alone does not address the accuracy disparity.

6.3 Analysis

We compare our method to no personalization (Global), naive fine-tuning (FT), Ditto (Li et al., 2021), up-weighting (UW) (Sagawa et al., 2020), and just train twice (JTT) (Liu et al., 2021). The up-weighting method is implemented via sampling bias-aligned and bias-conflicting samples with equal probability (Sagawa et al., 2020). Up-weighting and JTT are not applicable to Coil20 due to the lack of bias-conflicting samples. We use the local finetuning version of the Ditto solver because the local finetuning solver performs fewer local updates than that of the joint optimization solver and therefore entangles spurious features less. Each experiment is repeated 9 times with 3 random seeds for the federated learning step and 3 for the personalization step. We select models using the validation accuracy minus the decrease of adversarial transferability and using the validation accuracy for other baseline and competitor methods.

Tables 1 shows the main result. Our method reduces the accuracy disparity of personalized models from 15.12% to 2.15%, compared to other personalization methods. Our method also preserves the enhanced average accuracy from fine-tuning, resulting in 3.43% accuracy improvement on the biased test set and 0.85% improvement on the biased-conflicting test set compared to the global model. In contrast, the naive fine-tuning method sacrifices the accuracy on the bias-conflicting test set by up to 14.3% and increase the accuracy disparity by 15.12%. We also find that our methods outperform the supervised up-weighting method and the unsupervised JTT method, which increase the average accuracy disparity to 7.56% and 17.76%, respectively. One possible reason is that the diversity of the up-weighted bias-conflicting samples are small. Therefore, the neural network could memorize them instead of discarding spurious features.

7 Conclusion

In this work, we explored the risk of federated learning personalization methods with spurious features, which lead to high accuracy disparity between the global and local models. Through empirical explorations, we identity a strong positive correlation between accuracy disparity and the adversarial transferability between the global and local models. Inspired by this observation, we developed a strategy by enforcing the adversarial transferability between the global and personalized models to mitigate this risk and hence reduce the accuracy disparity. Our empirical results show that enforcing transferability reduces the accuracy disparity of personalized models.

8 ETHICS STATEMENT

Our method mitigates the issue of spurious features, which lead to bias towards minority groups, in personalized federated learning. However, completely disentangling spurious features remains challenging and is an issue for many federated learning methods.

9 REPRODUCIBILITY STATEMENT

We link the source code[1] for our methods that are presented in Section 5. The data partition in our experiments is detailed in Appendix A.1. The datasets in our experiments are publicly available.

REFERENCES

Tameem Adel, Han Zhao, and Richard E Turner. Continual learning with adaptive weights (claw). In *International Conference on Learning Representations*, 2019.

Alekh Agarwal, Miroslav Dudík, and Zhiwei Steven Wu. Fair regression: Quantitative definitions and reduction-based algorithms. In *International Conference on Machine Learning*, pp. 120–129. PMLR, 2019.

Jianfeng Chi, Yuan Tian, Geoffrey J Gordon, and Han Zhao. Understanding and mitigating accuracy disparity in regression. In *International Conference on Machine Learning (ICML)*, 2021.

Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.

Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. *Advances in Neural Information Processing Systems*, 33, 2020.

Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for clustered federated learning. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 19586–19597. Curran Associates, Inc., 2020. URL `https://proceedings.neurips.cc/paper/2020/file/e32cc80bf07915058ce90722ee17bb71-Paper.pdf`.

He He, Sheng Zha, and Haohan Wang. Unlearn dataset bias in natural language inference by fitting the residual. *arXiv preprint arXiv:1908.10763*, 2019.

Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.

Fereshte Khani and Percy Liang. Removing spurious features can hurt accuracy and affect groups disproportionately. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 196–205, 2021.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL `http://arxiv.org/abs/1412.6980`.

James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114 (13):3521–3526, 2017.

[1] https://drive.google.com/drive/folders/1TrXnjmc8qzFordsQJ8wkxnZgRcxzl9Iq?usp=sharing

Pang Wei Koh, Shiori Sagawa, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, Tony Lee, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021.

Daliang Li and Junpu Wang. Fedmd: Heterogenous federated learning via model distillation. *arXiv preprint arXiv:1910.03581*, 2019.

Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In I. Dhillon, D. Papailiopoulos, and V. Sze (eds.), *Proceedings of Machine Learning and Systems*, volume 2, pp. 429–450, 2020. URL https://proceedings.mlsys.org/paper/2020/file/38af86134b65d0f10fe33d30dd76442e-Paper.pdf.

Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pp. 6357–6368. PMLR, 2021.

Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9572–9581, 2019.

Kaizhao Liang, Jacky Y Zhang, Boxin Wang, Zhuolin Yang, Sanmi Koyejo, and Bo Li. Uncovering the connections between adversarial transferability and knowledge transferability. In *International Conference on Machine Learning*, pp. 6577–6587. PMLR, 2021.

Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning*, pp. 6781–6792. PMLR, 2021.

Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30:6467–6476, 2017.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJzIBfZAb.

Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pp. 1273–1282. PMLR, 2017.

Daniel McNamara, Cheng Soon Ong, and Robert C Williamson. Costs and benefits of fair representation learning. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 263–270, 2019.

Junhyun Nam, Hyuntak Cha, Sungsoo Ahn, Jaeho Lee, and Jinwoo Shin. Learning from failure: Training debiased classifier from biased classifier. In *Advances in Neural Information Processing Systems*, 2020.

Sameer A. Nene, Shree K. Nayar, and Hiroshi Murase. Columbia object image library (coil-20. Technical report, 1996.

Yurii Nesterov. *Introductory lectures on convex optimization: A basic course*, volume 87. Springer Science & Business Media, 2003.

Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=ryxGuJrFvS.

Canh T. Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with moreau envelopes. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 21394–21405. Curran Associates, Inc., 2020.

Florian Tramèr, Nicolas Papernot, Ian Goodfellow, Dan Boneh, and Patrick McDaniel. The space of transferable adversarial examples. *arXiv preprint arXiv:1704.03453*, 2017.

Haohan Wang, Zexue He, Zachary L. Lipton, and Eric P. Xing. Learning robust representations by projecting superficial statistics out. In *International Conference on Learning Representations*, 2019. URL `https://openreview.net/forum?id=rJEjjoR9K7`.

Dong Yin, Mehrdad Farajtabar, Ang Li, Nir Levine, and Alex Mott. Optimization and generalization of regularization-based continual learning: a loss approximation viewpoint. *arXiv preprint arXiv:2006.10974*, 2020.

Han Zhao and Geoff Gordon. Inherent tradeoffs in learning fair representations. *Advances in neural information processing systems*, 32:15675–15685, 2019.

Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *International Conference on Machine Learning*, pp. 7523–7532. PMLR, 2019.

A MORE EXPERIMENTAL SETTING

A.1 DATA PARTITION

Table 2: Number of Train and Validation Samples in CelebA_S

Client ID	Non-blond Female	Non-Blond Male	Blond Female	Blond Male
0	55	31	12	2
1	30	68	0	2
2	59	28	11	2

Table 3: Number of Biased Test Samples in CelebA_S

Client ID	Non-blond Female	Non-Blond Male	Blond Female	Blond Male
0	115	60	45	2
1	60	75	79	0
2	86	111	14	0

Table 4: Number of Biased-Conflicting Test Samples in CelebA_S

Client ID	Non-blond Female	Non-Blond Male	Blond Female	Blond Male
0	0	0	0	200
1	0	0	0	203
2	0	0	0	204

A.2 NEURAL NETWORK ARCHITECTURE

We use CNN 28x28 for MNIST dataset and CNN 64x64 for CelebA and Coil20 datasets.

Table 5: Neural Network Architecture

CNN 28x28	CNN 64x64
Input: $\mathbb{R}^{3\times28\times28}$	Input: $\mathbb{R}^{3\times64\times64}$
4×4 conv, 64 BN LReLU, stride 2	4×4 conv, 64 BN LReLU, stride 2
4×4 conv, 128 BN LReLU, stride 2	4×4 conv, 64 BN LReLU, stride 2
FC 4096 ReLU	FC 4096 ReLU
FC 10	FC 10