Anti-Matthew FL: Bridging the Performance Gap in Federated Learning to Counteract the Matthew Effect

Jiashi Gao¹, Xin Yao² and Xuetao Wei^{1,*}

¹Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, China

²Department of Computing and Decision Sciences, Lingnan University, Hong Kong SAR, China

Abstract. Federated learning (FL) stands as a paradigmatic approach that facilitates model training across heterogeneous and diverse datasets originating from various data providers. However, conventional FLs fall short of achieving consistent performance, potentially leading to performance degradation for clients who are disadvantaged in data resources. Influenced by the Matthew effect, deploying a performance-imbalanced global model in applications further impedes the generation of high-quality data from disadvantaged clients, exacerbating the disparities in data resources among clients. In this work, we propose anti-Matthew fairness for the global model at the client level, requiring equal accuracy and equal decision bias across clients. To balance the trade-off between achieving anti-Matthew fairness and performance optimality, we formalize the anti-Matthew effect federated learning (anti-Matthew FL) as a multi-constrained multi-objectives optimization (MCMOO) problem and propose a three-stage multi-gradient descent algorithm to obtain the Pareto optimality. We theoretically analyze the convergence and time complexity of our proposed algorithms. Additionally, through extensive experimentation, we demonstrate that our proposed anti-Matthew FL outperforms other state-of-the-art FL algorithms in achieving a high-performance global model while effectively bridging performance gaps among clients. We hope this work provides valuable insights into the manifestation of the Matthew effect in FL and other decentralized learning scenarios and can contribute to designing fairer learning mechanisms, ultimately fostering societal welfare.

1 Introduction

Federated learning (FL) [18] has emerged as a significant learning paradigm in which clients utilize their local data to train a global model collaboratively without sharing data, and has attracted researchers from various fields, especially in domains where data privacy and security are critical, such as healthcare, finance, and social networks [21, 30, 17].

1.1 Matthew Effect in FL

In the real world, clients may have unequal data resources due to historical or unavoidable social factors. They deserve fair treatment based on social welfare and equality principles. However, the existing *contribution fairness* [28, 16], which requires that the model

performance on each client is proportional to their data resource contributed, worsens the resource plight of poorer clients. For instance, when hospitals with non-i.i.d. datasets collaborate to train a disease diagnosis model, the hospitals with lower data resources will receive a model that does not fit well with their data distributions, as high-resource hospitals dominate the collaborative model training more. Therefore, the local model performance in a low-resource hospital may exhibit uncertain accuracy and decision bias. Such a low-trustworthy model may affect the subsequent diagnosis of lowresource clients, leading to persistent resource inequality and the deterioration of social welfare. This phenomenon is referred to as the **Matthew effect** [19], a social psychological phenomenon that describes how the rich get richer and the poor get poorer in terms of resources such as education, economy, and data information, etc.

1.2 Anti-Matthew Fairness

To mitigate the model performance disparities arising from the Matthew effect, we introduce the concept of *anti-Matthew fairness*, which requires the global model to exhibit equal performance across clients. We focus on two critical performances: **accuracy and decision bias**. Accuracy reflects how effectively the global model fits the local data of clients; hence, *equal accuracy performance can improve decision quality for clients with limited resources*, naturally mitigating their disadvantage and enhancing their subsequent ability to leverage model-empowered solutions to bridge the gap with other advantaged clients. Decision bias reflects the fairness of the global model's decisions concerning different protected groups, such as gender, race, or religion, within the client. Achieving *equal decision bias performance can enhance the reputation and decision credibility of the disadvantaged clients*, indirectly strengthening their ability to generate and contribute more valuable data resources.

1.2.1 Challenges and Considerations

Attaining *anti-Matthew fairness* in FL settings confronts the following challenges: **O** First, *anti-Matthew fairness* requires equal performance across clients, introducing a trade-off with achieving the highest performance during model training. This trade-off becomes particularly evident when dealing with heterogeneous clients, where the global model exhibits varied performance. For example, an advantaged client must make a trade-off by sacrificing local performance

^{*} Corresponding Author. Email: weixt@sustech.edu.cn.



Figure 1. Training goals in anti-Matthew FL.

for the sake of *anti-Matthew fairness*, as exceptional high local performance might result from the model's poor performance on other clients, which is undesirable from the standpoints of social welfare and ethics. Our goal is to strike a balance in this trade-off, ensuring that all clients achieve both high and equitable performance; **@** The second challenge arises from considering *anti-Matthew fairness* in both accuracy and decision bias. It is noted that there exists a potential trade-off between accuracy and decision bias [29].

1.2.2 Solutions

To address the challenges, we propose a novel FL framework named anti-Matthew FL to bridge performance gaps across clients. We first quantify anti-Matthew fairness, allowing the training objective to be formulated as a multi-constrained multi-objectives optimization (MCMOO) problem. As shown in Fig. 1, the objectives are to minimize the local empirical risk losses $\{l_1, ..., l_N\}$ on N clients. The local decision biases $\{f_1, ..., f_N\}$ on N clients are constrained to be below an acceptable bias budget. These two goals jointly ensure higher performance of model h across clients. We impose constraints on the deviations of local empirical risk loss and local decision bias from their mean values to achieve anti-Matthew fairness with equal performance among clients. We then propose a three-stage multigradient descent algorithm for global model training that converges to Pareto optimal solutions, in which the global model performance on each objective is maximized within the decision space and cannot be further improved without harming others.

1.3 Contribution

In summary, this paper makes the following contributions: **1** We propose anti-Matthew fairness, aiming to simultaneously bridge the gaps in accuracy and bias of the global model across clients. This is particularly necessary to mitigate the Matthew effect in collaborative training and is valuable to consider for welfare-oriented collaborative training; **2** We formally define anti-Matthew federated learning as a multi-constrained multi-objectives optimization (MCMOO) problem and propose a three-stage multi-gradient descent algorithm that achieves Pareto optimal solutions. We theoretically analyze the convergence and time complexity of the proposed algorithms; 3 We perform comprehensive experiments on both synthetic and real-world datasets to validate the effectiveness, convergence, hyperparameter sensitivity, rationality (via ablation study), and robustness of the proposed anti-Matthew FL. The results show that anti-Matthew FL outperforms state-of-the-art (SOTA) fair FL baselines in terms of overall performance and anti-Matthew fairness across clients.

2 Related Work

Most of the existing fairness research [8, 9, 27, 1, 26] assumes that the training process can access the whole training dataset in a centralized manner. However, this assumption does not hold when privacy protections prevent clients from sharing their data with a central server. In FL, several studies [7, 23, 2, 6, 3, 10] have focused on reducing the global model's decision bias towards various protected groups, such as gender, race, or age. However, due to the heterogeneity of client data, these efforts cannot guarantee an equitable distribution of decision bias across clients. Recently, there has been an increasing interest in promoting fairness among clients in FL. Li et al. [14] have introduced Ditto, allowing clients to fine-tune the global model using their local data. The primary objective of Ditto is to optimize local performance rather than to mitigate performance disparities. To mitigate performance disparities across clients, Mohri et al. [20] have proposed agnostic federated learning (AFL), a min-max multi-objective optimization that improves accuracy for the worstperforming client. Cui et al. [3] have proposed fair and consistent federated learning (FCFL), a multi-objective optimization designed to simultaneously maximize performance and ensure accuracy consistency across different local clients. The consistency in accuracy is also achieved by optimizing for the worst performing client, reducing only the gap between the worst and the best clients, and it does not satisfy the requirement for equitable performance to mitigate the Matthew effect, which demands minimized variance in performance distribution across clients. Furthermore, the disparity in decision bias among clients has not been considered. Li et al. [13] have proposed q-FFL, a heuristic method designed to achieve uniform accuracy across clients by adjusting the weights during aggregation to amplify the impact of clients with resource disadvantages. However, this approach does not guarantee that the global model achieves Pareto optimality. Pan et al. [22] have proposed FedMDFG, which incorporates cosine similarity between the loss vectors of clients and the unit vector as a fairness objective in local loss functions to achieve equitable accuracy among clients. However, this work does not address decision bias. Given the inherent trade-off between accuracy and decision bias for each client, as identified by Wang et al. [29], it is also necessary to consider the impact on decision bias when adjusting the accuracy distribution across clients. As simultaneously achieving both optimality and equity in accuracy and decision bias across clients is still an unexplored area and is essential for alleviating the impact of the Matthew effect in welfare-oriented collaborative training, we propose anti-Matthew effect federated learning (anti-Matthew FL) in this work to attain these objectives.

3 Preliminaries

3.1 Federated Learning

We focus on horizontal FL [31], which involves N clients, each associated with a specific dataset $\mathcal{D}_k = \{X_k, A_k, Y_k\}$, where $k \in \{1, ..., N\}$, X_k denotes the general attributes of the data without protected information, A_k denote a protected attribute, such as gender, race, or religion, and Y_k denoted truth label. The FL procedure involves multiple rounds of communication between the server and the clients. In each round, the server sends the global model h_{θ} with parameter θ to the clients, who then train their local models on their local private datasets $\{\mathcal{D}_1, ..., \mathcal{D}_N\}$, resulting in local models $\{h_{\theta_1}, ..., h_{\theta_N}\}$. The server then aggregates the local parameters and updates the global model for the next communication round [18]. The original FL [18] aims to minimize the average empirical risk loss over all the clients' datasets, and the optimal hypothesis parameter θ^* satisfies:

$$\theta^* = \arg \min_{\theta \in \Theta} \sum_{k=1}^{N} \left(\frac{|\mathcal{D}_k|}{\sum_{k=1}^{N} |\mathcal{D}_k|} l_k \left(\hat{Y}_k, Y_k \right) \right), \tag{1}$$

where $\hat{Y}_k = h_{\theta}(X_k, A_k)$ is the output of the hypothesis h_{θ} when input (X_k, A_k) and $l_k(\cdot)$ is the loss function for k-th client.

We instantiate the FL framework within the context of a binary classification problem, incorporating a binary sensitive attribute. For the optimization process, we employ the *binary cross entropy loss* to instantiate the loss function l_k .

$$l_k(h) = -\frac{1}{|Y|} \sum_{i=1}^{|Y|} \left(Y_k^i \log\left(\hat{Y}_k^i\right) + \left(1 - Y_k^i\right) \log\left(1 - \hat{Y}_k^i\right) \right).$$
(2)

The decision bias refers to the statistical disparity observed in model outcomes across different groups divided by protected attributes, such as gender, race, and region. We employ two specific decision bias metrics, namely the *True Positive Rate Parity Standard Deviation* (TPSD) and the *Accuracy Parity Standard Deviation* (APSD) [24]. The TPSD and APSD for the *k*-th client are defined as:

TPSD:
$$f_k(h) = \sqrt{\frac{\sum_{i=1}^{M} \left(\Pr(\hat{Y}_k = 1 | A_k = i, Y_k = 1) - \mu \right)^2}{M}},$$
 (3)
APSD: $f_k(h) = \sqrt{\frac{\sum_{i=1}^{M} \left(\Pr(\hat{Y}_k = Y_k | A_k = i) - \mu \right)^2}{M}},$

where μ is the average *True Positive Rate* (TPR) or average accuracy under all groups divided by the values of the protected attribute, and M is the number of possible values for the sensitive attribute A_k . A hypothesis h_{θ} satisfies ϵ_b -decision bias on k-th client if $f_k(h) \leq \epsilon_b$, where ϵ_b is the predefined budget for the decision bias.

3.2 Anti-Matthew Fairness

The anti-Matthew fairness refers to the model providing equal performance across clients. Pan et al. [22] propose cosine similarity between local losses and a unit vector p = 1 to assess the equity of model losses across all clients. This metric, however, lacks granularity in distinguishing each client's performance and fails to impose precise performance constraints. To address this limitation, we propose an alternative approach for evaluating model performance equality across clients. We measure this equality through the absolute deviation of each client's performance from the mean performance of all clients. Specifically, a hypothesis h satisfies ϵ_{vl} -anti-Matthew fairness concerning accuracy performance and ϵ_{vb} -anti-Matthew fairness concerning decision bias performance if:

$$\left|l_k(h) - \bar{l}(h)\right| \le \epsilon_{vl}, \left|f_k(h) - \bar{f}(h)\right| \le \epsilon_{vb}, k \in \{1, \dots, N\},$$
(4)

where $\bar{l}(h) = \frac{1}{N} \sum_{k=1}^{N} l_k(h)$ and $\bar{f}(h) = \frac{1}{N} \sum_{k=1}^{N} f_k(h)$ are the average empirical risk loss and average decision bias, respectively, and ϵ_{vl} and ϵ_{vb} are the predefined budgets for the *anti-Matthew fairness* on accuracy and decision bias, respectively.

3.3 Anti-Matthew Federated Learning

To achieve a global model that provides both high and *anti-Matthew fairness* across clients, we propose a novel framework called *anti-Matthew effect federated learning (anti-Matthew FL)*, in which the training goals can be formulated as a multi-constrained multi-objectives optimization (MCMOO) problem.

Definition 1. (*Anti-Matthew FL*) *We formalize the anti-Matthew FL training objective as follows:*

$$\min_{h \in \mathcal{H}^{*}} \{l_{1}(h), ..., l_{N}(h)\}, s.t. \{f_{k}(h)\}_{k=1}^{N} \leq \epsilon_{b},
\{|l_{k}(h) - \bar{l}(h)|\}_{k=1}^{N} \leq \epsilon_{vl}, \{|f_{k}(h) - \bar{f}(h)|\}_{k=1}^{N} \leq \epsilon_{vb},$$
(5)

where h is a hypothesis from a hypothesis set \mathcal{H}^* .

The MCMOO problem seeks to minimize the empirical risk losses for all clients while ensuring each client has a ϵ_b -decision bias. It also satisfies ϵ_{vl} -anti-Matthew fairness for accuracy and ϵ_{vb} -anti-Matthew fairness for decision bias. Finding the optimal solution to the MCMOO problem is nontrivial, as the objectives may conflict. Therefore, we aim to identify the Pareto-optimal hypothesis h, which is not dominated by any other $h' \in \mathcal{H}$. The definitions of Pareto optimal and Pareto front [15] are as follows:

Definition 2. (Pareto Optimal and Pareto Front) In a multiobjective optimization problem with loss function $l(h) = \{l_1(h), ..., l_N(h)\}$, we say that for $h_1, h_2 \in \mathcal{H}$, h_1 is dominated by h_2 if $\forall i \in [N], l_i(h_2) \leq l_i(h_1)$ and $\exists i \in [N], l_i(h_2) < l_i(h_1)$. A solution h is Pareto optimal if it is not dominated by any other $h' \in \mathcal{H}$. The collection of Pareto optimal solutions is called the Pareto set. The image of the Pareto set in the loss function space is called the Pareto front.

4 Detailed Mechanism of Anti-Matthew FL

4.1 Stages to Obtain Pareto Optimal

Fig. 2 illustrates the feasible decision space of Eq. (5), which is bounded by the intersection of two hypothesis sets: \mathcal{H}_B , and \mathcal{H}_E . The \mathcal{H}_B contains hypotheses satisfying ϵ_b -decision bias in each client,

$$\{f_k(h)\}_{k=1}^N \le \epsilon_b, \forall h \in \mathcal{H}_B.$$
(6)

The \mathcal{H}_E contains hypotheses that satisfy ϵ_{vl} -anti-Matthew fairness on accuracy and ϵ_{vb} -anti-Matthew fairness on decision bias across all clients,

$$\left\{ \left| l_k \left(h \right) - \bar{l} \left(h \right) \right| \right\}_{k=1}^N \le \epsilon_{vl}, \\ \left\{ \left| f_k \left(h \right) - \bar{f} \left(h \right) \right| \right\}_{k=1}^N \le \epsilon_{vb}, \forall h \in \mathcal{H}_E.$$
 (7)

The set $\mathcal{H}^* \subset \mathcal{H}_B \cap \mathcal{H}_E$ represents the Pareto set corresponding to Eq. (5). Formally, for any hypotheses h in \mathcal{H}^* and h' in $\mathcal{H}_B \cap \mathcal{H}_E$, it holds that $h' \not\prec h$, where \prec denotes the Pareto dominance relation.

Proposition 1. (*Existence of Feasible Hypothesis*) The feasible decision space, $\mathcal{H}_B \cap \mathcal{H}_E \neq \emptyset$, is guaranteed to be non-empty.

Proof. For $h \in \mathcal{H}_E$, if $\bar{f}(h) \leq \epsilon_b - \epsilon_{vb}$, then $\{|f_k(h) - \bar{f}(h)|\}_{k=1}^N \leq \epsilon_{vb}$, this implies that h also satisfies $\{f_k(h)\}_{k=1}^N \leq \epsilon_b$ and belongs to \mathcal{H}_B . Therefore, $\mathcal{H}_B \cap \mathcal{H}_E \neq \emptyset$ and includes at least h that fulfills the requirement $\bar{f}(h) \leq \epsilon_b - \epsilon_{vb}$. \Box

Finding the Pareto set for *anti-Matthew FL* poses significant challenges due to the highly constrained nature of the feasible decision space. Additionally, when dealing with a large number of objectives, the optimization of one objective may have adverse effects on others. To address this issue, we employ a linear scalarization technique to construct an approximate Pareto front. Average weights are assigned to each objective, transforming the multi-objectives into a single surrogate objective. This surrogate objective forms the convex segment



Figure 2. Optimization paths to achieve a Pareto solution $h \in \mathcal{H}^*$ in *anti-Matthew FL*.

of the Pareto front, as illustrated in Fig. 2, which is denoted as $\mathcal{H}_{\bar{L}}$. The hypothesis in the $\mathcal{H}_{\bar{L}}$ satisfies:

$$\bar{l}(h) \leq \bar{l}(h'), \forall h \in \mathcal{H}_{\bar{L}}, h' \notin \mathcal{H}_{\bar{L}}.$$
(8)

Compared to \mathcal{H}^* , $\mathcal{H}_{\bar{L}}$ is easier to obtain and can serve as an intermediate set, from which we propose a three stages optimization path: $h^0 \to \mathcal{H}_B \cap \mathcal{H}_{\bar{L}} \to \mathcal{H}_B \cap \mathcal{H}_E \cap \mathcal{H}_{\bar{L}} \to \mathcal{H}^*$ (purple arrows in Fig. 2), and decompose the *anti-Matthew FL* into three sub-problems as follows:

Stage 1: Constrained Minimization Problem. We define a constrained minimization problem on the hypothesis set \mathcal{H} to obtain a hypothesis $h' \in \mathcal{H}_B \cap \mathcal{H}_{\bar{L}}$,

$$\min_{h \in \mathcal{H}} \bar{l}(h), \text{ s.t. } \{f_k(h)\}_{k=1}^N \le \epsilon_b.$$
(9)

By solving Eq. (9), we obtain h' that 1) satisfies ϵ_b -decision bias for each client and 2) minimizes the average empirical risk loss among all clients.

Stage 2: Multi-Constrained Optimization Problem. We formulate a multi-constrained optimization problem to obtain a hypothesis $h'' \in \mathcal{H}_B \cap \mathcal{H}_E \cap \mathcal{H}_{\bar{L}}$,

$$\min_{h \in \mathcal{H}} \bar{l}(h), \text{ s.t. } \{ |l_k(h) - \bar{l}(h)| \}_{k=1}^N \leq \epsilon_{vl}, \\ \{ |f_k(h) - \bar{f}(h)| \}_{k=1}^N \leq \epsilon_{vb}, \{ f_k(h) \}_{k=1}^N \leq \epsilon_b.$$
(10)

By solving Eq. (10), we obtain h'' that, compared to h', exhibits the following properties: 1) it provides ϵ_{vl} -anti-Matthew fairness on accuracy; and 2) it provides ϵ_{vb} -anti-Matthew fairness on decision bias.

Stage 3: Multi-Constrained Pareto Optimization Problem. Solely prioritizing the minimization of the weighted sum $\bar{l}(h)$ during optimization may harm the performance of individual clients. To mitigate this concern, we formulate a multi-constrained Pareto optimization problem aimed at refining h'' to $h^* \in \mathcal{H}^*$. In this process, the empirical risk loss of each client is further reduced until Pareto optimality is attained. After this stage, the loss of each client cannot be further minimized without causing a negative impact on the loss of other clients.

$$\min_{\substack{h \in \mathcal{H} \\ b_{k} \in \mathcal{H}}} \{l_{1}(h), ..., l_{N}(h)\},$$
s.t. $\{f_{k}(h)\}_{k=1}^{N} \leq \epsilon_{b}, \{|l_{k}(h) - \bar{l}(h)|\}_{k=1}^{N} \leq \epsilon_{vl},$ (11)
 $\{|f_{k}(h) - \bar{f}(h)|\}_{k=1}^{N} \leq \epsilon_{vb}, \bar{l}(h) \leq \bar{l}(h'').$

4.2 Three-Stage Gradient Descent to Obtain \mathcal{H}^*

To obtain the convergent solution for the sub-problems defined in Eq. (9) to Eq. (11), we propose a three-stage multi-gradient descent

algorithm for acquiring $h^* \in \mathcal{H}^*$ that is well-suited for implementation within *federated stochastic gradient descent* (FedSGD). Given a hypothesis h_{θ^t} parameterized by θ^t , at iteration t + 1, the update rule of the parameters is $\theta^{t+1} = \theta^t + \eta d$, where d represents gradient descent direction, and η is the step size. In the context of an optimization problem with N objectives, i.e., $\min \{l_1(h_\theta), ..., l_N(h_\theta)\}$, the gradient descent direction d is considered effective in guiding the optimization towards minimization if $\{d^{*T} \nabla_{\theta} l_i(h_\theta)\}_{i=1}^N \leq 0$. As the gradient direction d resides within the convex hull of the gradients of all objectives, denoted as $G = [\nabla_{\theta} l_1(h_{\theta}), ..., \nabla_{\theta} l_N(h_{\theta})]$ [4], we can obtain the optimal gradient descent direction d^* by performing a linear transformation on G using an N-dimensional vector α^* ,

$$d^* = \alpha^{*T} G, \text{ where } \alpha^* = \arg\min_{\alpha} \sum_{i=1}^N \alpha_i \nabla_{\theta} l_i (h_{\theta}),$$

s.t. $\sum_{i=1}^N \alpha_i = 1, \alpha_i \ge 0, \forall i \in [N].$ (12)

Optimal Gradient Descent Direction in Stage 1. We initially convert Eq. (9) into an equivalent single-constraint optimization problem by imposing a constraint solely on the maximum value, as follows:

$$\min_{h \in \mathcal{H}} \bar{l}(h), \text{ s.t. } \max\left\{f_k(h)\right\}_{k=1}^N \le \epsilon_b.$$
(13)

Denoting the max $\{f_k(h)\}_{k=1}^N$ as $f_{max}(h)$, the descent gradient of Eq. (13) lies in the convex hull of $G' = [\nabla_{\theta} \bar{l}(h), \nabla_{\theta} f_{max}(h)]$. Drawing inspiration from Cui et al.'s work [3], we adopt an alternating optimization strategy to alleviate computational burden: if the ϵ_b -decision bias is satisfied within the worst-case client, only $\bar{l}(h)$ is further optimized,

$$d^* = \arg \min_{d \in G'} d^T \nabla_{\theta} \bar{l}(h), \text{ if } f_{max}(h) \le \epsilon_b.$$
(14)

Otherwise, we optimize towards a descent direction d, which minimizes $f_{max}(h)$ while ensuring that $\bar{l}(h)$ does not increase, as follows:

$$d^* = \arg \min_{d \in G'} d^T \nabla_{\theta} f_{max}(h),$$

s.t. $d^T \nabla_{\theta} \bar{l}(h) \le 0$, if $f_{max}(h) > \epsilon_b.$ (15)

The gradient direction in Eq. (14) and Eq. (15) is optimized to minimize the loss while better satisfying the ϵ_b -decision bias constraint. This optimization process results in a hypothesis h' that achieves a balanced trade-off between accuracy and decision bias.

Optimal Gradient Descent Direction in Stage 2. To reduce the computational complexity associated with handling O(N) constraints in Eq. (10), we focus on optimizing *anti-Matthew fairness* specifically for the worst-case client. Furthermore, to better achieve ϵ_{vl} -anti-Matthew fairness on accuracy and ϵ_{vb} -anti-Matthew fairness on decision bias, we modify Eq. (10) by treating anti-Matthew fairness as objectives and introducing a constraint $\overline{l}(h) \leq \overline{l}(h')$ to prevent the degradation of model accuracy performance.

We optimize *anti-Matthew fairness* of accuracy and decision bias alternately, i.e., if $\{|l_k(h) - \bar{l}(h)|\}_{k=1}^N \leq \epsilon_{vl}$,

$$\min_{h \in \mathcal{H}} \max\left\{ \left| f_k\left(h\right) - \bar{f}(h) \right| - \epsilon_{vb} \right\}_{k=1}^N, \\ \text{s.t.} \max\left\{ f_k(h) \right\}_{k=1}^N \le \epsilon_b, \bar{l}\left(h\right) \le \bar{l}\left(h'\right),$$
(16)

else,

$$\min_{h \in \mathcal{H}} \max \left\{ \left| l_{k} (h) - \bar{l}(h) \right| - \epsilon_{vl} \right\}_{k=1}^{N}, \\
\text{s.t.} \max \left\{ \left| f_{k} (h) - \bar{f}(h) \right| \right\}_{k=1}^{N} \le \epsilon_{vb}, \\
\max \left\{ f_{k}(h) \right\}_{k=1}^{N} \le \epsilon_{b}, \bar{l}(h) \le \bar{l}(h').$$
(17)

Denoting the $\max\left\{\left|l_{k}\left(h\right)-\bar{l}(h)\right|-\epsilon_{vl}\right\}_{k=1}^{N}$ and $\max\left\{\left|f_{k}\left(h\right)-\bar{f}(h)\right|-\epsilon_{vb}\right\}_{k=1}^{N}$ as $\hat{l}_{max}(h)$ and $\hat{f}_{max}(h)$,

respectively, the gradient descent direction of Eq. (16) lies in the convex hull of $G'' = \left[\nabla_{\theta} \hat{f}_{max}, \nabla_{\theta} f_{max}, \nabla_{\theta} \bar{l} \right]$. We obtain the optimal d^* as follows:

$$d^* = \arg\min_{d \in G''} d^T \nabla_{\theta} \hat{f}_{max}(h),$$

s.t. $d^T \nabla_{\theta} \bar{l}(h) \le 0, d^T \nabla_{\theta} f_{max}(h) \le 0$ if $f_{max}(h) > \epsilon_b.$ (18)

The gradient descent direction of Eq. (17) lies in the convex hull of $G'' = \left[\nabla_{\theta} \hat{l}_{max}, \nabla_{\theta} \hat{f}_{max}, \nabla_{\theta} f_{max}, \nabla_{\theta} \bar{l} \right]$. We obtain the optimal d^* as follows:

$$d^{*} = \arg\min_{d \in G''} d^{T} \nabla_{\theta} \hat{l}_{max}(h), \text{ s.t. } d^{T} \nabla_{\theta} \bar{l}(h) \leq 0,$$

$$d^{T} \nabla_{\theta} \hat{f}_{max}(h) \leq 0 \text{ if } \hat{f}_{max}(h) > \epsilon_{vb},$$

$$d^{T} \nabla_{\theta} f_{max}(h) \leq 0 \text{ if } f_{max}(h) > \epsilon_{b}.$$
(19)

The constraints are dynamically imposed based on whether the current hypothesis h satisfies ϵ_b -decision bias and ϵ_{vb} -anti-Matthew fairness on the decision bias. The optimal gradient direction in Eq. (18) and Eq. (19) is optimized to enhance the equality of performance among clients without compromising the overall model performance. This optimization process results in a hypothesis h'' that strikes a balance between anti-Matthew fairness and maximizing model performance.

Optimal Gradient Descent Direction in Stage 3. To reduce the computational complexity of minimizing N objectives and handling O(N)-constraints in Eq. (11), we optimize the empirical risk loss for the worst-case client and impose constraints l_k (h) $\leq l_k$ (h'') to prevent the degradation of performance for other clients.

$$\min_{h \in \mathcal{H}} \max \{ l_{1}(h), ..., l_{N}(h) \}, \\
\text{s.t. } l_{k}(h) \leq l_{k}(h''), \forall k \in [N], \\
\bar{l}(h) \leq \bar{l}(h''), \max \{ f_{k}(h) \}_{k=1}^{N} \leq \epsilon_{b}, \\
\max \{ |l_{k}(h) - \bar{l}(h)| \}_{k=1}^{N} \leq \epsilon_{vl}, \\
\max \{ |f_{k}(h) - \bar{f}(h)| \}_{k=1}^{N} \leq \epsilon_{vb}.$$
(20)

Denoting the max $\{l_k(h)\}_{k=1}^N$ as $l_{max}(h)$, the gradient descent direction lies in the convex hull of:

$$G^* = [\nabla_{\theta} l_{max}(h), \nabla_{\theta} l_1(h), ..., \nabla_{\theta} l_N(h), \nabla_{\theta} \bar{l}(h), \\ \nabla_{\theta} f_{max}(h), \nabla_{\theta} \hat{l}_{max}(h), \nabla_{\theta} \hat{f}_{max}(h)].$$

We obtain the optimal d^* as follows,

$$d^{*} = \arg\min_{\substack{d \in G^{*} \\ d^{T} \nabla_{\theta} f_{max}(h) \leq 0 \text{ if } f_{max}(h) > \epsilon_{b}, \\ d^{T} \nabla_{\theta} f_{max}(h) \leq 0 \text{ if } f_{max}(h) > \epsilon_{vl}, \\ d^{T} \nabla_{\theta} \hat{f}_{max}(h) \leq 0 \text{ if } \hat{f}_{max}(h) > \epsilon_{vb}, \\ d^{T} \nabla_{\theta} \hat{f}_{max}(h) \leq 0 \text{ if } \hat{f}_{max}(h) > \epsilon_{vb}, \\ d^{T} \nabla_{\theta} l_{i}(h) \leq 0, \forall i \in [N] \text{ and } i \neq \arg\max\left\{l_{k}(h)\right\}_{k=1}^{N}. \end{cases}$$
(21)

The optimal gradient direction in Eq. (21) is tailored to minimize the empirical risk loss on the worst client without adversely affecting the performance of other clients. This leads to further refinement from the approximation Pareto front, formed by linear scalarization, to a Pareto optimal hypothesis h^* of Eq. (5).

4.3 Convergence and Time Complexity Analysis

Proposition 2. (Convergence-Preserving Gradient Descent) For N optimization objectives $\{l_1(\theta^t), ..., l_N(\theta^t)\}$ and the model parameter updating rule under gradient descent direction $d: \theta^{t+1} = \theta^t + \eta d$; if $d^T \nabla_{\theta} l_i \leq 0$, there exists η_0 such that for $\forall \eta \in [0, \eta_0]$, the objectives $\{l_1(\theta^t), ..., l_N(\theta^t)\}$ will not increase.

Based on Prop. 2, and with each new stage constraining the gradient direction without compromising the performance achieved in the previous stage, the proposed three-stage gradient descent enables a convergent FL training process towards Pareto optimality.

Proposition 3. (Polynomial Time-Complexity) In our proposed three-stage gradient descent, solving direction d^* has a polynomial time complexity and does not exceed $\tilde{O}^*(2^{2.38}) + \tilde{O}^*(4^{2.38}) + \tilde{O}^*((N+4)^{2.38})$, where N represents the number of clients.

5 Experiments

5.1 Settings

Datasets. O Synthetic dataset: we generate a synthetic dataset with a protected attribute: $A \sim Ber(0.5)$, two general attributes: $X_1 \sim$ $\mathcal{N}(0,1), X_2 \sim \mathcal{N}(\mathbb{1}(a > 0), 2)$. The label is setting by $Y \sim$ $Ber(u^{l}\mathbb{1}(x_{1}+x_{2} \leq 0)+u^{h}\mathbb{1}(x_{1}+x_{2} > 0)), \text{ where } \{u^{l}, u_{h}\} =$ $\{0.3, 0.6\}$ if A = 0, otherwise, $\{u^l, u^h\} = \{0.1, 0.9\}$. We split the dataset into two clients based on whether $x_1 \leq -0.5$ to make the clients heterogeneous in distribution and size; 2 Adult dataset [12]: Adult is a binary classification dataset with more than 40000 adult records for predicting whether the annual income is greater than 50K. We split the dataset into two clients based on whether the individual's education level is a Ph.D and select race as a protected attribute; I EICU dataset [11]: the eICU dataset includes data records about clinical information and hospital details of patients who are admitted to ICUs. We filter out the hospitals with data points less than 1500, leaving 11 hospitals in our experiments. Naturally, we treat each hospital as a client and select race as a protected attribute; $\mathbf{0}$ ACSPublicCoverage dataset [5]: ACSPublicCoverage dataset is used to predict whether an individual is covered by public health insurance in the United States. The dataset is collected in the year 2022. Naturally, we treat each state as a client, generating 51 clients in our experiments and select gender as a protected attribute.

Baselines. • *FedAvg* [18]: the original FL algorithm for distributed training of private data. It does not consider fairness for different demographic groups and different clients; • *FedAvg* + *FairBatch* [25]: each client adopts the state-of-the-art FairBatch in-processing debiasing strategy on its local training data and then aggregation uses FedAvg; • *FedAvg*+*FairReg*: a local processing method by optimizing the linear scalarized objective with the fairness regularizations of all clients; • *Ditto* [14]: an FL framework to achieve training fairness by allowing clients to fine-tune the received global model on the local data; • *q*-*FFL* [13]: an FL framework to achieve accuracy consistency among clients by weighing different local clients differently; • *FCFL* [3]: an FL framework to achieve accuracy consistency across different local clients meanwhile to minimize decision bias; • *FedMDFG* [22]: an FL framework to achieve accuracy consistency across different local clients.

Hyperparameters. We divide the communication rounds into three stages, each with 750, 750, and 500 rounds, respectively, to ensure that the global model is fully updated and converges in each stage. In the constraint budgets setting, we set the decision bias budget ϵ_b , the *anti-Matthew fairness* budget on accuracy ϵ_{vl} , and the *anti-Matthew fairness* budget on decision bias ϵ_{vb} to half of the relatedperformance achieved by the original *FedAvg*. For example, as shown in Tab. 1, on the synthetic dataset experiments, the Avg. local TPSD of *FedAvg* is 0.2480, so we set $\epsilon_b = 0.1$. The Std. local accuracy of *FedAvg* is 0.0283, so we set $\epsilon_{vl} = 0.01$. The Std. local TPSD of *FedAvg* is 0.0819, so we set $\epsilon_{vb} = 0.04$. We use the same parameter-setting strategy for other datasets. Since the constraints may conflict

		Model Performance			
Dataset	Method	Local Acc.		Local TPSD	
		Avg.	$Std.(\epsilon_{vl})$	Avg.(ϵ_b)	$\operatorname{Std.}(\epsilon_{vb})$
	FedAvg	.7735	.0283(×)	.2480(×)	.0819(×)
	q-FFL	.7735	.0283(×)	.2480(×)	.0819(×)
Synthetic	Ditto	.7229	.0132(×)	.2703(×)	$.0566(\times)$
$\epsilon_b = 0.1$	FedMDFG	.7717	.0068()	.2473(×)	$.0662(\times)$
$\epsilon_{vl} = 0.01$	FedAvg+FairBatch	.6360	$.0643(\times)$.1040(≈)	.0798(×)
$\epsilon_{ub} = 0.04$	FedAvg+FairReg	.6227	$.0394(\times)$.0952()	.0463(×)
00	FCFL	.6330	.0177(×)	.0812()	.0435(≈)
	Anti-Matthew FL (Ours)	.6327	.0087(√)	.0801()	.0359(√)
	FedAvg	.7767	.0592(×)	$.0328(\times)$.0093 (×)
	q-FFL	.7662	.0400(×)	.0472(×)	.0046()
Adult	Ditto	.7210	.0039(1)	$.0169(\times)$.0111(×)
$\epsilon_h = 0.01$	FedMDFG	.7656	.0397(×)	.0455(×)	.0069(×)
$\epsilon_{v,l} = 0.03$	FedAvg+FairBatch	.7756	.0556(×)	.0136(×)	.0128(×)
$\epsilon_{vh} = 0.005$	FedAvg+FairReg	.7663	$.0686(\times)$.0089()	.0066(×)
00	FCFL	.7638	.0487(×)	.0143(×)	.0159(×)
	Anti-Matthew FL (Ours)	.7685	.0281(√)	.0036()	.0009 (√)
	FedAvg	.6560	.0427(×)	.0371(×)	$.0409(\times)$
	q-FFL	.6565	$.0425(\times)$	$.0371(\times)$.0405(×)
eICU	Ditto	.6311	.0216(≈)	.0472(×)	.0447(×)
$\epsilon_h = 0.02$	FedMDFG	.6479	$.0227(\times)$	$.0311(\times)$	$.0266(\times)$
$\epsilon_{al} = 0.02$	FedAvg+FairBatch	.6441	.0413(×)	.0304(×)	.0298(×)
$\epsilon_{ub}^{01} = 0.02$	FedAvg+FairReg	.6455	$.0408(\times)$	$.0322(\times)$.0266(×)
00	FCFL	.6550	$.0272(\times)$.0344(×)	.0246(×)
	Anti-Matthew FL (Ours)	.6530	.0195()	.0209(≈)	.0201(≈)
	FedAvg	.7015	.0815(×)	.0307(×)	$.0242(\times)$
	q-FFL	.5888	.0406(≈)	.0236(×)	.0202(×)
ACSPublicCoverage	Ditto	.6578	.0584(×)	.0260(×)	.0328(×)
$\epsilon_{h} = 0.015$	FedMDFG	.5978	.0453(×)	.0215(×)	.0187(×)
$\epsilon_{vl} = 0.04$	FedAvg+FairBatch	.5972	.0454(×)	.0214(×)	.0188(×)
$\epsilon_{vb} = 0.015$	FedAvg+FairReg	.5964	.0453(×)	$.0202(\times)$.0194(×)
	FCFL	.6285	.0391()	.0215(×)	.0187(×)
	Anti-Matthew FL (Ours)	.6237	.0384(1/)	.0147	.0128(1/)

 Table 1.
 Local performance under TPSD metric.

Red : Methods that consider reducing decision bias

Bold : Best performance compared to all algorithms (×) : Violation of constraint exceeds 10%.

(×): Violation of constraint exceeds 10%.
 (≈): Close to constraint, with violation of constraint not exceeding 10%.

 $(\sqrt{)}$: Satisfy constraint.



Figure 3. Model performance distribution across different clients.

with each other, this setting allows us to better evaluate the superior performance of our proposed *anti-Matthew FL* and avoid making a constraint too tight, which may result in a solution that is only optimal on that constraint.

Evaluation Metrics. To evaluate the effectiveness and equality of the global model's performance across all clients, we introduce two evaluation metrics: **1** Avg.: the average performance of the global model across all clients; **2** Std.: the standard deviation of the performance of the global model across clients. We utilize the TPSD and APSD in Eq. (3) as metrics for decision bias.

Normalization. We suggest normalizing the gradients **g** before computing the optimal gradient descent direction d^* when there is a large disparity among the gradients, as in the eICU and ACSPublicCoverage experiments where the numbers of clients are large,

$$\mathbf{g} = \mathbf{g} / \sqrt{\sum_{i=1}^{|g|} g_i^2}.$$
(22)

5.2 Performance Comparison

We compare the global model performance of *anti-Matthew FL* with other SOTA baselines on four datasets. Tab. 1 and 2 respectively illustrate the model performance under TPSD and APSD as decision

 Table 2.
 Local performance under APSD metric.

	Method	Model Performance			
Dataset		Local Acc.		Local APSD	
		Avg.	$Std.(\epsilon_{vl})$	Avg.(ϵ_b)	$Std.(\epsilon_{vb})$
	FedAvg	.7735	$.0283(\times)$.1663(×)	.0791(×)
	q-FFL	.7735	$.0283(\times)$.1663(×)	.0791(×)
Synthetic	Ditto	.7229	.0132(×)	.1977(×)	.0615(×)
$\epsilon_b = 0.08$	FedMDFG	.7717	.0068()	.1667(×)	.0718(×)
$\epsilon_{vl} = 0.01$	FedAvg+FairBatch	.6695	.0397(×)	$.0999(\times)$.0401()
$\epsilon_{vb} = 0.04$	FedAvg+FairReg	.6191	$.0250(\times)$.0976(×)	.0720(×)
	FCFL	.6302	$.0165(\times)$.0687(√)	$0453(\times)$
	Anti-Matthew FL (Ours)	.6269	.0029(√)	.0621(\sqrt{)}	.0430(≈)
	FedAvg	.7767	.0592(×)	.0494(×)	.0257(×)
	q-FFL	.7662	.0400(×)	.0386(×)	.0089(\/)
Adult	Ditto	.7210	.0039(√)	.0450 (×)	.0481 (×)
$\epsilon_{h} = 0.02$	FedMDFG	.7656	.0397(×)	$.0436(\times)$.0068()
$\epsilon_{vl} = 0.03$	FedAvg+FairBatch	.7726	.0556 (×)	.0218(≈)	.0076 ()
$\epsilon_{vb} = 0.01$	FedAvg+FairReg	.7446	$.0484(\times)$.0109()	$.0186(\times)$
	FCFL	.7583	.0487(×)	.0109()	.0195(×)
	Anti-Matthew FL (Ours)	.7549	.0284(√)	.0101(\sqrt{)}	.0067 $()$
	FedAvg	.6560	.0427(×)	$.0386(\times)$	$.0310(\times)$
	q-FFL	.6565	.0425(×)	$.0384(\times)$.0307(×)
eICU	Ditto	.6311	.0216(≈)	.0399(×)	.0405(×)
$\epsilon_h = 0.02$	FedMDFG	.6479	.0227(×)	.0395(×)	.0301(×)
$\epsilon_{vl} = 0.02$	FedAvg+FairBatch	.6441	.0213(≈)	.0301(×)	.0227(×)
$\epsilon_{vb}^{vi} = 0.02$	FedAvg+FairReg	.6346	$.0765(\times)$	$.0335(\times)$	$.0267(\times)$
	FCFL	.6569	$.0362(\times)$	$.0331(\times)$.0200()
	Anti-Matthew FL (Ours)	.6544	.0192()	.0207(pprox)	.0182 $()$
ACSPublicCoverage	FedAvg	.7015	$.0815(\times)$	$.0198(\times)$	$.0139(\times)$
	q-FFL	.5888	.0406(≈)	.0160(×)	.0140(×)
	Ditto	.6578	.0584(×)	.0162(×)	$.0118(\times)$
$\epsilon_{b} = 0.01$	FedMDFG	.5978	.0453(×)	.0158(×)	.0123(×)
$\bar{\epsilon_{vl}} = 0.04$	FedAvg+FairBatch	.5946	.0441(×)	.0152(×)	.0128(×)
$\epsilon_{vb} = 0.01$	FedAvg+FairReg	.5977	.0453(×)	.0158(×)	.0122(×)
	FCFL	.6239	.0403(≈)	$.0114(\times)$.0095()
	Anti-Matthew FL (Ours)	.6221	.0402(pprox)	.0106(\sqrt{)}	.0080(\sqrt{)}

bias metrics. It is evident that anti-Matthew FL achieves the best satisfaction of the fairness constraints. In terms of accuracy, the methods that do not consider decision bias (FedAvg, q-FFL, Ditto, and Fed-MDFG) have higher accuracy. However, as there is a trade-off between accuracy and decision bias, we compare our method with the baselines that also aim to reduce decision bias (*FedAvg+FairBatch*, FedAvg+FairReg, FCFL). Under TPSD, anti-Matthew FL achieves the best constraint satisfaction with only 0.3% decrease in accuracy on the synthetic dataset, 0.7% decrease in accuracy on the adult dataset, 0.2% decrease in accuracy on the eICU dataset, and 0.7%on the ACSPublicCoverage dataset. Under APSD, anti-Matthew FL achieves the best constraint satisfaction with only 0.3% decrease in accuracy on the synthetic and adult dataset, 0.2% decrease in accuracy on the eICU and ACSPublicCoverage dataset. The results further validate the scalability of the anti-Matthew FL, with clients ranging from 2 to 51.

Fig. 3 illustrates the distribution of model performance across various clients. We use TPSD as decision bias metric. For a clearer visualization, we prioritize the display of baselines considering decision bias (*FedAvg+Fairbatch, FedAvg+FairReg, FCFL*) on the eICU dataset, which encompasses 11 clients. The results demonstrate that *anti-Matthew FL* ensures a more equitable performance distribution among clients, thereby indicating enhanced *anti-Matthew fairness*.

5.3 Convergence

Fig. 4 illustrates the convergence of *anti-Matthew FL* within 2000 communication rounds on Adult and eICU datasets. We use TPSD as decision bias metric. The experiments are repeated 20 times. We observe that the accuracy of the global model converges as the communication rounds increase, while the decision bias and the *anti-Matthew fairness* of accuracy and decision bias converge to the predefined fairness budgets (indicated by the colored dashed lines in Fig. 4).



Figure 4. Testing results during 2000 communication rounds.

5.4 Hyperparameters Sensitivity

In *anti-Matthew FL*, we introduce three constraint budgets, namely, ϵ_b , ϵ_{vl} , and ϵ_{vb} , to regulate the decision bias and the *anti-Matthew fairness*. To investigate the impact of varied constraint configurations on the *anti-Matthew FL* algorithm, we manipulated the values of ϵ_b , ϵ_{vl} , and ϵ_{vb} within the range of [0.01, 0.02, 0.05, 0.1], respectively.

As shown in Fig. 5, the decision bias of the global model is inversely proportional to ϵ_b , which regulates the bias level. The ϵ_{vl} and ϵ_{vb} regulate the equality of the model performance across the clients, specifically, the smaller the ϵ_{vl} , the smaller the standard deviation of accuracy among the clients; the smaller the ϵ_{vb} , the smaller the standard deviation of decision bias among the clients.



Figure 5. Effects of budgets ϵ_b , ϵ_{vl} and ϵ_{vb} on the Avg. of local TPSDs, Std. of local accuracies and Std. of local TPSDs, respectively.

5.5 Ablation Study

Compared to solving a MCMOO problem in a single-stage, a multistages method in *anti-Matthew FL* offers the following advantages:

- Balanced Trade-offs: The staged method promotes a more nuanced and balanced consideration of trade-offs;
- Partial Constraint Focus: Dividing the problem into three distinct stages allows us to concentrate on satisfying partial constraints at each stage. This focused approach proves beneficial for navigating more effectively through a progressively stringent decision space in subsequent stages;
- Enhanced Convergence Speed: The staged method contributes to a heightened convergence speed. In each stage, the emphasis is on solving a specific subproblem, thereby mitigating overall complexity and significantly improving the speed of arriving at solutions.

In this experiment, we conducted an ablation study to systematically assess the indispensability of each stage within the three-stage algorithm. The results presented in Tab. 3 demonstrate the realization of the desired hypothesis within the decision space constrained by fairness, achieved solely by integrating all three stages.

5.6 Robustness

We conduct robustness validation experiments on the eICU dataset and randomly select 4 clients to be malicious. The malicious clients

Table 3. Ablation study of each stage in anti-Matthew FL.

Dataset	Stage	Avg. Acc. \pm Std. Acc.	Avg. TPSD \pm Std. TPSD
Synthetic	Stage 1	$0.6519 \pm 0.0169 (imes)$	$0.1033(\approx)\pm0.0477(\times)$
$\epsilon_b = 0.1$	Stage 2	$0.5740 \pm 0.0661 (imes)$	$0.0772 (\sqrt{)} \pm 0.0740 (\times)$
$\epsilon_{vl} = 0.01$	Stage 3	$0.6840 \pm 0.0983(\times)$	$0.2483(\times)\pm 0.0670(\times)$
$\epsilon_{vb} = 0.04$	Stage 1+Stage 2+Stage 3	$0.6327 \pm 0.0087(\checkmark)$	0.0801 (√)±0.0359(√)
Adult	Stage 1	$0.7778 \pm 0.0530 (imes)$	$0.0167(\times)\pm0.0212(\times)$
$\epsilon_b = 0.01$	Stage 2	$0.7611 \pm 0.0450 (imes)$	$0.0069(\sqrt{)\pm 0.0094(\times)}$
$\epsilon_{vl} = 0.03$	Stage 3	$0.7383 \pm 0.0206(\sqrt{)}$	$0.0109(\approx)\pm0.0102(\times)$
$\epsilon_{vb} = 0.005$	Stage 1+Stage 2+Stage 3	$0.7685 \pm 0.0281(\sqrt{)}$	$0.0036(\sqrt{)\pm 0.0009(\sqrt{)}}$
eICU	Stage 1	$0.6488 \pm 0.0223(\times)$	$0.0189(\sqrt{)\pm0.0263(\times)}$
$\epsilon_b = 0.02$	Stage 2	$0.6503 \pm 0.0207 (\approx)$	$0.0356(\times)\pm0.0230(\times)$
$\epsilon_{vl} = 0.02$	Stage 3	$0.6446 \pm 0.0226(\times)$	$0.0309(\times)\pm0.0209(\approx)$
$\epsilon_{vb} = 0.02$	Stage 1+Stage 2+Stage 3	$0.6530 \pm 0.0195(\sqrt{)}$	$0.0209(pprox)\pm0.0201(pprox)$
ACSPublicCoverage	Stage 1	$0.6215 \pm 0.0526(\times)$	$0.0194 (\times) \pm 0.0160 (\approx)$
$\epsilon_{b} = 0.015$	Stage 2	$0.6441 \pm 0.0567(\times)$	$0.0211(\times)\pm0.0158(\approx)$
$\epsilon_{vl} = 0.04$	Stage 3	$0.6913 \pm 0.0754 (imes)$	$0.0071(\checkmark) \pm 0.0065(\checkmark)$
$\epsilon_{vb} = 0.015$	Stage 1+Stage 2+Stage 3	0.6237 ± 0.0384 (\checkmark)	0.0147 (√)±0.0128 (√)

adopt the following attacks to dominate or disrupt the FL process: (1) Enlarge: the malicious clients enlarge the local gradients or local model parameters sent to the server to enhance their influence in the training process. In our experiments, we set the enlarging factor to 10; (2) Random: the malicious clients send randomly generated local gradients or model parameters to the server to disrupt the training process; (3) Zero: The malicious clients send zero local gradients or zero model parameters to the server to disrupt the training process.

Tab. 4 shows the testing performance of the SOTA baselines and our proposed *anti-Matthew FL* under attacks. The results demonstrate that our *anti-Matthew FL* is robust and the performance of the global model does not degrade due to the presence of malicious clients, while other baselines suffer from different degrees of decline in testing performance.

Table 4. The test performance under attacks.

Attack	Avg. Acc. \pm Std. Acc./ Avg. TPSD \pm Std. TPSD				
	Enlarge	Random	Zero		
FedAvg	.5307±.0414/.0205±.0669	.5129±.0743/.0242±.0291	.6225±.0503/.0245±.0411		
q-FFL	.6278±.0261/.0242±.0250	$.5824 \pm .0461 / .0225 \pm .0371$.6492±.0386/.0216±.0379		
Ditto	$.6162 {\pm} .0283 {\rm /} .0203 {\pm} .0247$	$.5130 \pm .0771 / .0273 \pm .0331$.6236±.0498/.0231±.0399		
FedMDFG	$.6512 \pm .0345 / .0217 \pm .0340$	$.6502 \pm .0332 / .0229 \pm .0318$.6466±.0271/.0213±.0211		
FedAvg+FairBatch	$.6079 \pm .0280 / .0199 \pm .0200$	$.5095 \pm .0724 / .0297 \pm .0335$.6225±.0503/.0245±.0411		
FedAvg+FairReg	$.5643 {\pm} .0582 {\rm /} .0272 {\pm} .0372$	$.6510 \pm .0299 / .0225 \pm .0275$.6511±.0292/.0228±.0283		
FCFL	$.6280 \pm .0296 / .0232 \pm .0237$.6535±.0263/.0225±.0233	.6254±.0317/.0215±.0295		
Anti-Matthew FL	.6536 ±.0278/ .0178 ± .0185	.6510±.0244/.0203±.0158	.6522±.0259/.0206±.0183		

6 Conclusion

To mitigate **Matthew effect** in federated learning (FL), which exacerbates resource disparities among clients and jeopardizes the sustainability of the FL systems, we defined a novel fairness concept, *anti-Matthew fairness*. We analyzed the challenges of achieving *anti-Matthew fairness* in the FL setting. Specifically, we formally defined *anti-Matthew FL* as a multi-constrained, multi-objectives optimization problem. Subsequently, we designed an effective threestage multi-gradient descent algorithm to achieve the Pareto optimal global model. We also provided theoretical analyses of the algorithm's convergence properties. Finally, we conducted comprehensive empirical evaluations to demonstrate that our proposed *anti-Matthew FL* outperforms other state-of-the-art baselines in achieving a high-performance global model with enhanced *anti-Matthew fairness* among all clients.

Acknowledgements

This work was supported by Key Programs of Guangdong Province under Grant 2021QN02X166. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding parties.

References

- Y. Choi, M. Dang, and G. Van den Broeck. Group fairness by probabilistic modeling with latent fair decisions. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12051–12059, Virtual Event, May 2021. AAAI Press.
- [2] L. Chu, L. Wang, Y. Dong, J. Pei, Z. Zhou, and Y. Zhang. Fedfair: Training fair models in cross-silo federated learning. arXiv:2109.05662, September 2021.
- [3] S. Cui, W. Pan, J. Liang, C. Zhang, and F. Wang. Addressing algorithmic disparity and performance inconsistency in federated learning. *Advances in Neural Information Processing Systems*, 34:26091–26102, December 2021.
- J.-A. Désidéri. Multiple-gradient descent algorithm (mgda) for multiobjective optimization. *Comptes Rendus Mathematique*, 350(5-6):313– 318, March 2012.
- [5] F. Ding, M. Hardt, J. Miller, and L. Schmidt. Retiring adult: New datasets for fair machine learning. *Advances in Neural Information Processing Systems*, 34:6478–6490, 2021.
- [6] W. Du, D. Xu, X. Wu, and H. Tong. Fairness-aware agnostic federated learning. In *Proceedings of the 2021 SIAM International Conference* on Data Mining (SDM), pages 181–189, Philadelphia, PA, April 2021. SIAM.
- [7] Y. H. Ezzeldin, S. Yan, C. He, E. Ferrara, and A. S. Avestimehr. Fairfed: Enabling group fairness in federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7494–7502, Washington, USA, February 2023. AAAI Press.
- [8] R. Ghani, K. T. Rodolfa, P. Saleiro, and S. Jesus. Addressing bias and fairness in machine learning: A practical guide and hands-on tutorial. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 5779–5780, Long Beach, USA, August 2023. ACM.
- [9] U. Gohar, S. Biswas, and H. Rajan. Towards understanding fairness and its composition in ensemble machine learning. In 2023 IEEE/ACM 45th International Conference on Software Engineering, pages 1533–1545, Melbourne, Australia, May 2023. IEEE.
- [10] S. Hu, Z. S. Wu, and V. Smith. Fair federated learning via bounded group loss. arXiv:2203.10190, March 2022.
- [11] A. E. Johnson, T. J. Pollard, and T. Naumann. Generalizability of predictive models for intensive care unit patients. arXiv:1812.02275, December 2018.
- [12] R. Kohavi and B. Becker. Uci adult data set. UCI Meachine Learning Repository, 5, 1996.
- [13] T. Li, M. Sanjabi, A. Beirami, and V. Smith. Fair resource allocation in federated learning. In *International Conference on Learning Representations*, New Orleans, USA, May 2019. OpenReview.
- [14] T. Li, S. Hu, A. Beirami, and V. Smith. Ditto: Fair and robust federated learning through personalization. In *International Conference on Machine Learning*, pages 6357–6368, Virtual Event, July 2021. PMLR.
- [15] X. Lin, H.-L. Zhen, Z. Li, Q.-F. Zhang, and S. Kwong. Pareto multitask learning. Advances in neural information processing systems, 32: 12060–12070, December 2019.
- [16] Z. Liu, Y. Chen, Y. Zhao, H. Yu, Y. Liu, R. Bao, J. Jiang, Z. Nie, Q. Xu, and Q. Yang. Contribution-aware federated learning for smart healthcare. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 12396–12404, Virtual Event, February 2022. AAAI Press.
- [17] G. Long, Y. Tan, J. Jiang, and C. Zhang. Federated Learning for Open Banking. Springer International Publishing, 2020.
- [18] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282, Ft. Lauderdale, USA, April 2017. PMLR.
- [19] R. K. Merton. The matthew effect in science: The reward and communication systems of science are considered. *Science*, 159(3810):56–63, January 1968.
- [20] M. Mohri, G. Sivek, and A. T. Suresh. Agnostic federated learning. In *International Conference on Machine Learning*, pages 4615–4625, Long Beach, USA., June 2019. PMLR.

- [21] D. C. Nguyen, Q.-V. Pham, P. N. Pathirana, M. Ding, A. Seneviratne, Z. Lin, O. Dobre, and W.-J. Hwang. Federated learning for smart healthcare: A survey. ACM Comput. Surv., 55, 2022.
- [22] Z. Pan, S. Wang, C. Li, H. Wang, X. Tang, and J. Zhao. Fedmdfg: Federated learning with multi-gradient descent and fair guidance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9364–9371, Washington, USA, February 2023. AAAI Press.
- [23] A. Papadaki, N. Martinez, M. Bertran, G. Sapiro, and M. Rodrigues. Minimax demographic group fairness in federated learning. In Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, pages 142–159, New York, USA, June 2022. ACM.
- [24] R. Poulain, M. F. Bin Tarek, and R. Beheshti. Improving fairness in ai models on electronic health records: The case for federated learning methods. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, pages 1599–1608, New York, USA, June 2023. ACM.
- [25] Y. Roh, K. Lee, S. E. Whang, and C. Suh. Fairbatch: Batch selection for model fairness. In *International Conference on Learning Representations*, Virtual Event, May 2021. OpenReview.
- [26] G. Sankar, A. Louis, M. Nasre, and P. Nimbhorkar. Matchings with group fairness constraints: Online and offline algorithms. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 377–383, Virtual Event, August 2021. IJCAI.
- [27] X. Shen, Y. Wong, and M. Kankanhalli. Fair representation: guaranteeing approximate multiple group fairness for unknown tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):525– 538, February 2022.
- [28] S. S. Tay, X. Xu, C. S. Foo, and B. K. H. Low. Incentivizing collaboration in machine learning via synthetic data rewards. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 9448–9456, Virtual Event, September 2022. AAAI Press.
- [29] Y. Wang, X. Wang, A. Beutel, F. Prost, J. Chen, and E. H. Chi. Understanding and improving fairness-accuracy trade-offs in multi-task learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 1748–1757, Virtual Event, August 2021. ACM.
- [30] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. ACM Trans. Intell. Syst. Technol., 10, 2019.
- [31] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. ACM Trans. Intell. Syst. Technol., 10(2):1– 19, January 2019.