

FedFair³: Unlocking Threefold Fairness in Federated Learning

Simin Javaherian Sanjeev Panta Shelby Williams Md Sirajul Islam Li Chen
School of Computing and Informatics, University of Louisiana at Lafayette
Email: {simin.javaherian1, sanjeev.panta1, shelby, md-sirajul.islam1, li.chen}@louisiana.edu

Abstract—Federated Learning (FL) is an emerging paradigm in machine learning without exposing clients’ raw data. In practical scenarios with numerous clients, encouraging fair and efficient client participation in federated learning is of utmost importance, which is also challenging given the heterogeneity in data distribution and device properties. Existing works have proposed different client-selection methods that consider fairness; however, they fail to select clients with high utilities while simultaneously achieving fair accuracy levels. In this paper, we propose a fair client-selection approach that unlocks threefold fairness in federated learning. In addition to having a fair client-selection strategy, we enforce an equitable number of rounds for client participation and ensure a fair accuracy distribution over the clients. The experimental results demonstrate that FedFair³, in comparison to the state-of-the-art baselines, achieves 18.15% less accuracy variance on the IID data and 54.78% on the non-IID data, without decreasing the global accuracy. Furthermore, it shows 24.36% less wall-clock training time on average.

Index Terms—Accuracy, Convergence Analysis, Fairness, Federated Learning, Importance Sampling, Participant Selection

I. INTRODUCTION

Traditional machine learning relies on centralized servers for data gathering and training, which falls short when dealing with privacy-sensitive data residing in a vast number of edge devices. To address the challenge, Federated Learning (FL) has emerged as a promising alternative, which operates within a loosely federated network of clients, thereby enabling collaborative model updates. In FL, participants retain their data locally on their devices, perform local training, and subsequently share their model updates with a central server for global aggregation.

Despite its potential for distributed learning over decentralized data, FL has encountered open challenges to be addressed for large-scale real-world deployments. One important challenge is to enable a fair system where all clients are encouraged to participate in updating their models. Otherwise, with an unfair system, clients do not receive fair rewards and thus become reluctant to participate. On the other hand, efficiently learning a high-quality model in FL is challenging, considering the diverse set of client devices and heterogeneous data distributions. This requires judiciously giving different priorities to clients according to their capabilities of updating models and communicating with servers.

Recent efforts such as Oort [1] have sought to enhance FL performance through guided participant selection, albeit without considering client features and the distribution of

accuracy among clients. On the other hand, q-FFL [2] takes fair accuracy distribution into account but lacks a comprehensive consideration of selective client participation. Focused on communication efficiency, [3]–[7] did not consider fairness, inversely impacting client willingness in FL participation. All the existing works fail to simultaneously address efficient learning and fair client selection in FL. Moreover, there has not even been any consensus on the definition of fairness.

To this end, this paper aims to answer the following research question: in large-scale FL with heterogeneous clients, how should we elicit and achieve the best possible fairness among all the clients, while ensuring an efficient collaborative model learning? This involves addressing the open challenges of defining appropriate fairness constraints and designing an optimal client selection strategy throughout the learning process toward efficiency and fairness goals. To address these challenges, we present a novel approach, named *FedFair³*, which takes into account client features while maintaining fairness in client participation from three perspectives, without compromising performance and accuracy. In particular, the client features including data size, energy consumption, round duration, power consumption and local loss are comprehensively considered in the probabilistic client selection framework in *FedFair³*. Different from the existing probability-based frameworks ([8], [9]), our approach integrates three simultaneous notions of fairness. Firstly, participants with similar resources and capabilities are selected with equal probability to ensure fairness in participant selection. Secondly, we introduce the concept of accuracy fairness, where we aim to ensure that the performance of each participant, in terms of the model accuracy level, is proportional to its resource and capability. Thirdly, clients are penalized after a specific number of rounds, to allow fair participation throughout the whole FL process.

To shed light on this, we provide a glimpse of our results through a toy example on Fig. 1. Evaluating the variance of accuracy of these clients in each round, it becomes evident that FedFair³ significantly outperforms Oort, which serves as the baseline for our approach. The experimental results show that our approach exhibits an 18.15% reduction in the variance of accuracy over IID data and a remarkable 54.78% reduction over non-IID data, showcasing its superiority over existing methods. Furthermore, regarding the efficiency, FedFair³ also outperforms the baseline algorithms, achieving a 24.36% reduction in wall clock time, thereby advancing us closer to

the goal of achieving target accuracy in less time, a critical objective in the realm of FL [10], [11].




Fig. 1: A toy example of comparison of local accuracy (Oort vs. FedFair³) considering 6 clients in each round.

Our Contributions are summarized as follows:

- This paper proposes a new approach that selects clients in a fine-grained manner considering their demands and resources including computational power, data size, time duration and energy consumption.
- It also considers fairness regarding the accuracy variance and the number of rounds that each client can participate in training. We achieve this by giving larger probabilities to those clients that have more capabilities.
- We show the experimental results, demonstrating the deviation of accuracy, the final accuracy and the wall-clock training time for IID and non-IID datasets.

II. BACKGROUND AND RELATED WORKS

In this section, we provide an overview of FL, including its general concept and key components. We then summarize and discuss related works to motivate our proposed design.

A. Preliminaries of Federated Learning

FL typically involves a vast amount of edge devices, such as smartphones and laptops, and a server that periodically synchronizes the global model, denoted as w , across devices or clients. Considering $f_i(w)$ as the local loss over the n_i data samples held by the client i and n denote the total number of samples, the learning goal in FL is to solve the following optimization problem:

$$\begin{aligned} \min_w \quad & F(w) = \sum_{i=1}^N p_i \alpha_i F_i(w) \\ \text{s.t.} \quad & F_i(w) = \frac{1}{n_i} \sum_{i=1}^{n_i} f_i(w) \\ & \sum_i n_i = n, \quad \sum_i p_i = 1, \quad p_i \geq 0 \end{aligned} \quad (1)$$

p_i is the probability of the client selection that usually is $p_i = \frac{n_i}{n}$, and α_i is the weight of client i that is usually considered 1. The rule of the model update is defined as below:

$$w^{t+1} = w^t - \eta \alpha_i \nabla F(w^t)$$

B. Fairness and Efficiency Related Works

Fairness in Federated Learning (FL) has been explored from various angles in previous research. The core aim of fairness is to keep all clients engaged and motivated in the learning process. Moreover, using the importance sampling [8], the probability can be selected as a proportion of loss. That is, $p_i \propto F(w)$. We use the importance sampling in FL, and the experiments show that importance sampling is advantageous over random sampling as it strategically prioritizes training samples based on their loss values, allowing the optimization algorithm to focus more on informative instances, resulting in improved convergence and generalization performance. However, we also take the client utilities and features into account, making a more efficient and fair system. We briefly summarized the different fairness researches as follows.

Counterfactual fairness [12] is another fairness notion in which all individuals, considering the protected attributes in a casual sense, should receive same distribution of prediction, regardless the group they belong to (including race, gender, etc.). Barocas et al. [13] and Mitchell et al. [14] also mentioned that different sensitive groups should receive same patterns of outcomes; otherwise, it would violate the demographic fairness. There are also some surveys which discuss the FL challenges with a focus on fairness issues [15]

In optimization, Mohri et al. [16] proposed a new framework of agnostic FL that uses a minimax optimization approach, but it optimizes the model just for a small number of clients [2]. On a larger scale, Li et al. [2] introduced a parameter q as the q-FFL method that provides a uniform accuracy for the network. AFL is a special case of q-FFL, when the q is large enough. The q-FFL approach minimizes an aggregate reweighted loss parameterized by q such that the devices with higher loss are given higher relative weight. However, it did not consider participant selection. FedProx [10] incorporates a proximal term into the local training objective with the aim of maintaining the proximity of local models to the global model. The authors in [17] address the trade-off between fairness and global model performance with using an entropy perspective. Chu et al. in [18] propose a formal FL fairness definition, fairness via agent-awareness (FAA), which takes the heterogeneity of different agents into account. In HFFL [19], agents which contribute more to FL are rewarded more in this framework. The agents at the different contribution levels thus receive different model updates. Fan et al. in [20] proposed a new approach that uses a matrix containing clients' contributions and data to address this problem. Another work introduced collaborative fairness in FL (CFFL) [21], which considers fairness regarding the clients' contribution. It evaluates the clients' contributions and updates this information steadily, so it knows the clients' reputation and can fairly distribute updated models. Since working with the non-IID data over a distributed system is unfair, Ray et al. [22] introduced a new fairness concept called core-stable fairness.

C. Motivation

Trade-off between accuracy and fairness. The trade-off between accuracy and fairness is an open challenge. We hope to answer the question: is there any solution to improve the fairness without decreasing the accuracy?

Considering client features to increase the efficiency and fairness. In FL, clients come with varying computational resources, data sizes, and energy limitations, making it crucial to foster fairness. While previous research has introduced various fairness strategies [22]–[24], our work identifies a gap between fairness and client selection probabilities.

Considering more fairness notions into account. Lacking a consensus on fairness, we naturally wonder: Can we have a client selection approach that simultaneously explores more fairness notions?

In response, we propose a novel approach to tackle the multifaceted issue of client selection. Our approach takes a detailed view, considering individual client demands and resources, including computing power, data volume, time constraints, and energy usage. Different from all the existing work, we extend our commitment to fairness beyond just client selection. Our approach also addresses accuracy variance and the number of training rounds allotted to each client, giving preference to those with greater capabilities.

III. METHODOLOGY

In this section, we present our new approach, FedFair³, to address the fairness challenge of client selection in FL. We begin by outlining our assumptions and presenting the objective function, which serves as the basis for our fairness notions regarding accuracy and participant selection. We then present our client selection method, which works in a fair way by using a non-uniform probability distribution instead of a random participant selection. Finally, we show that our approach leads to a uniform and fair system regarding the accuracy variance.

TABLE I: Notations

The Meaning of the Notations	
T	The developer's preference of duration of each round
t_i	The amount of time that client i takes to perform the training
τ	Local rounds
K	The number of global aggregations
I	The total number of local iterations
r_i	Duration of local round that client i takes to perform its training
$l(\cdot)$	$\cdot$ is a Boolean value, if it is true, $l(\cdot)$ is 1, otherwise it is 0
β	A penalty value for those clients that take more time to complete their round
d_i	The size of the data of the client i
c_i	Computational power of the client i
q_i	Energy consumption of each client i
γ_i	$\in \{0, 1\}$, a binary indicator for the selection of client i
κ	Batch size
s_ϑ	The current total accumulated type ϑ resource usage
r_l	Local resource consumption type ϑ
r_g	Global resource consumption type ϑ
Π_ϑ	Resource budget type ϑ
L	Lipschitz constant
E	Explored clients for updating the priority

Assumptions: We assume the following for each client i :

1) $F_i(w)$ is convex [25], i.e.,

$$F(\delta w_1 + (1 - \delta)w_2) \leq \delta F(w_1) + (1 - \delta)F(w_2) \quad (2)$$

2) $F_i(w)$ is L-Lipschitz [26], satisfying

$$\|F_i(w) - F_i(w')\| \leq L\|w - w'\|, \forall w, w' \quad (3)$$

A. FedFair³ Algorithm

Our approach is designed to ensure fairness and efficiency in the selection of clients for federated learning. We accomplish this by strategically choosing clients with higher losses, while taking into account their distinct features. To achieve a uniform distribution of accuracy across clients, we introduce the concept of weighting in our objective function, denoted as α_i . In essence, with this weighting mechanism, each client contributes fairly to the learning process. In more practical terms, our approach not only prioritizes clients with higher loss functions but also factors in their resource capabilities. This means that clients with greater computational resources, larger datasets, lower energy consumption, and quicker training times within our defined limit, denoted as T, have an increased likelihood of being selected. We also ensure that we do not repeatedly select the same clients, thus giving priority to those who have not yet been chosen.

To achieve this, the server plays a crucial role in aggregating clients' characteristics before initiating the model update. This aggregation allows us to rank clients based on their unique features and available resources, ensuring a well-balanced selection process. However, we also consider time as a valuable resource. Clients that exceed the predefined time limit, S , face a penalty denoted as β , which reduces their selection probability. Conversely, clients that require less time than the set limit have their selection priority determined by their available resources, provided they are selected. In each training round, we carefully monitor the consumption of resources. If the cumulative resource consumption surpasses our predefined budget, the server will signal the clients to stop. This dynamic resource management strategy not only ensures fairness but also maintains the efficiency of the federated learning process. We call the resource budget Π_ϑ for the resource type ϑ . Thus, we have the following resource constraint:

$$(I + 1)r_l + (K + 1)r_g \leq \Pi_\vartheta, \forall \vartheta \text{ s.t. } I = K\tau \quad (4)$$

Our algorithm relies on the design of client selection probability p_i and client weight α_i for the FL optimization in Eq. (1) as follows:

$$S, \alpha_i = \frac{p_i^q}{N(q + 1)} \quad (5)$$

where

$$U_i = |\kappa| \sqrt{\frac{1}{|\kappa|} \times \sum_{i \in \kappa} \text{Loss}(i)^2} \cdot \left(\frac{T}{t_i}\right)^{l(T < t_i)\beta} \cdot \lambda^{l(T > t_i)\gamma_i},$$

$$\lambda = \frac{c_i d_i}{q_i r_i}, \gamma_i = \begin{cases} 1, & \text{if client } i \text{ is selected} \\ 0, & \text{otherwise} \end{cases}$$

U_i defines the client *Utility* which refers to a measure that combines the loss values of data samples to assess the clients' significance in improving model performance during various

training tasks with T representing the preferred round duration, t_i denoting the time taken by client i to process training data, λ as the clients' priority considering the clients features, and $l(x)$ serving as an indicator function that evaluates to 1 when x is true and 0 otherwise. Table I demonstrates the meaning of the notations.

We show that with the design of α_i as Eq. (5) and substituting in Eq. (1), we have the equivalent objective function as in q-fairness [2].

$$\begin{aligned} \sum_{i=1}^N p_i \alpha_i F_i(w) &= \sum_{i=1}^N p_i \frac{p_i^q}{N(q+1)} F_i(w) \\ &\propto \sum_{i=1}^N \frac{p_i F_i(w)^q}{N(q+1)} F_i(w) = \sum_{i=1}^N \frac{p_i}{N(q+1)} F_i(w)^{q+1} \end{aligned} \quad (6)$$

Building upon the concepts of q-fairness, our algorithm aims to address fairness considerations in the context of FL, while also striving for equitable and efficient collaboration among clients. With this context in mind, let's proceed to the summary of Algorithm 1 steps as follows:

- **Step1:** The server aggregates clients' features to calculate individual client probabilities for participation in the FL round (Line 3-16).
- **Step2:** Leveraging calculated probabilities, the server prioritizes clients and samples a set based on their priority, penalizing those with lower priority determined by their loss values (Line 17).
- **Step3:** Models are sent to the sampled clients, emphasizing those with higher priority, enabling them to contribute to the federated model update (Line 18-22).
- **Step4:** The server checks the resource budget by considering previous values and current available resources, ensuring that the federated learning process aligns with resource constraints (Line 23-24).
- **Step5:** Clients, whose participation adheres to the resource budget, engage in model training, promoting collaboration while respecting resource limitations (Line 29-32).

B. Fairness Qualification

This section proves that our objective function has a uniform distribution over clients. For the ease of mathematical exposition, we consider the following objective function of FL with limiting q to 0 as a conventional objective function or 1 as our objective function: $\min_w F_q(w) = \sum_{i=1}^N p_i \alpha_i^q F_{i,q}(w)^{q+1}$.

We show this uniform distribution with two definitions as follows:

Definition 1. *Uniformity of variance of the performance distribution:* we say the distribution of N clients $\{F_{1,q}(w), \dots, F_{N,q}(w)\}$ under solution w is more uniform than w' if:

$$\text{Var}_{\mathbf{p}}[F_{1,q}(w), \dots, F_{N,q}(w)] < \text{Var}_{\mathbf{p}}[F_{1,q}(w'), \dots, F_{N,q}(w')] \quad (7)$$

Algorithm 1: FedFair³

```

1 Participant selection for each round
2 At The Aggregator
3 while True do
4   Alleviating preferred time  $T$ 
5   Calculate client probability:
6   for client  $i \in E$  do
7     calculate  $LossF(i)$ 
8      $U_i \leftarrow |\kappa| \sqrt{\frac{\sum_{i \in \kappa} F(i)^2}{|\kappa|}}$ 
9     if  $T < t_i$  then
10       $U_i \leftarrow U_i \times \left(\frac{T}{t_i}\right)^\beta$ 
11    else
12       $\lambda \leftarrow \frac{c_i d_i}{q_i r_i}$ 
13       $U_i \leftarrow U_i \times \lambda$ 
14    end
15     $p_i = \frac{U_i}{\sum U_i}$ 
16  end
17  sample clients by priority
18  for selected clients do
19    Receive  $p_i, \alpha_i, \nabla F_i(w), c_\vartheta$  for all  $i$ 
20     $\nabla F(w) \leftarrow \nabla F(w) + p_i \alpha_i \nabla F_i(w)$ 
21     $w^{t+1} \leftarrow w^t - \eta \alpha_i \nabla F(w^t)$ 
22  end
23  if  $\exists \vartheta | s_\vartheta + \sum_{j \in N} (r_{j,\vartheta} \tau + 2r_g) > \Pi_\vartheta$  then
24    Send a "Stop" message to the clients
25  end
26  Return  $w$ 
27 end
28 At the selected clients
29 while they did not receive a "Stop" message from the
    aggregator do
30   Receive  $w$  and  $\tau$ 
31   Perform updating  $w_i$ 
32   Send  $w_i, \nabla F(w_i), r_{\vartheta,i}$ 
33 end

```

Our algorithm selects clients with higher losses. As a result, we have a system with less variance. According to our first notation of fairness, our system is fair regarding its variance.

Definition 2. *Uniformity of cosine similarity between the performance distribution and $\mathbf{1}$:* we say the distribution of cosine similarity over N clients $\{F_{1,q}(w), \dots, F_{N,q}(w)\}$ under solution w and $\mathbf{1}$ is more uniform than w' and $\mathbf{1}$ if:

$$\frac{\frac{1}{N} \sum_{i=1}^N F_{i,q}(w)}{\sqrt{\frac{1}{N} \sum_{i=1}^N F_{i,q}^2(w)}} \geq \frac{\frac{1}{N} \sum_{i=1}^N F_{i,q}(w')}{\sqrt{\frac{1}{N} \sum_{i=1}^N F_{i,q}^2(w')}} \quad (8)$$

Considering w^* as the optimal solution of $\min F_q(w)$, we have: $\frac{1}{N} \sum_{i=1}^N F_{i,1}(w^*) \geq \frac{1}{N} \sum_{i=1}^N F_{i,0}(w^*)$ and $\frac{1}{N} \sum_{i=1}^N F_{i,1}(w^*)^2 \geq \frac{1}{N} \sum_{i=1}^N F_{i,0}(w^*)^2$. Omitting similar

steps as in [2] due to the space limit, we have:

$$\frac{\frac{1}{N} \sum_{i=1}^N F_{i,1}(w^*)}{\sqrt{\frac{1}{N} \sum_{i=1}^N F_{i,1}^2(w^*)}} \geq \frac{\frac{1}{N} \sum_{i=1}^N F_{i,0}(w^*)}{\sqrt{\frac{1}{N} \sum_{i=1}^N F_{i,0}^2(w^*)}}$$

IV. EXPERIMENTAL RESULT

A. Experimental setup

Dataset and Models. We evaluate the performance of FedFair³ on three different popular benchmark datasets including MNIST [27], FashionMNIST [28], and CIFAR10 [29] on IID and non-IID datasets.

Implementation. Our experiments were conducted using the Plato framework [30]. We ran each experiment three times using the table II parameters with a penalty factor 2. We executed all experiments on a server, which is equipped with NVIDIA GeForce RTX 3080Ti GPU, Intel(R) Core(TM) i9-10900X CPU, and 64G RAM. We run both the server and clients on the same machine, a configuration supported by the fact that the performance metrics we evaluate are independent of the physical separation between the server and clients.




Fig. 2: The variance of accuracy of FedFair³ versus the FedAvg, AFL, FedProx and Oort.

B. Results and Analysis

Performance Comparison. Table IV and Fig. 2, demonstrate the results of our experimental evaluation. We compared the proposed approach, FedFair³, with the baseline algorithms: Oort, FedProx, AFL and FedAvg. As depicted, the observed

TABLE II: The objective used for experimental results

Hyperparameters			
Objectives	MNIST	CIFAR10	FashionMNIST
q	2	2	2
Learning Rate	0.01	0.001	0.01
# Clients/Round	10	20	6
# Total Clients	100	300	100
# Rounds	100	150	100
Batch Size	100	64	100
Optimizer	SGD	Adam	Adam
Model Name	LeNet-5 [31]	VGG16 [32]	LeNet-5 [31]

TABLE III: The wall clock time of FedFair³ vs. baseline algorithms.

Wall Clock Time Comparison(in second)				
Dataset	AFL	FedProx	Oort	FedFair ³
FMNIST IID	16908	16686	16683	16568
FMNIST Non-IID	33598	33073	33307	33070
MNIST IID	30458	2317	2932	2198
MNIST Non-IID	15954	3818	2332	2281
CIFAR10 IID	12199	12414	12701	12040
CIFAR10 Non-IID	12712	16276	12404	12068

results indicate that there is not a substantial difference in terms of accuracy between the compared methods. However, it is worth noting that our algorithm demonstrates a notable advantage in terms of reduced variance of accuracy across different clients. Furthermore, based on the data provided in the table IV, it can be observed that the proposed algorithm, on the average, exhibits 21.31% less accuracy variation in CIFAR10, 48.48% less variation in FashionMNIST, and 45.11% less variation in MNIST non-IID datasets compared to the Oort algorithm. In addition to the improved accuracy variance, our proposed algorithm, FedFair³, also offers enhanced efficiency by considering the features of clients during the client selection process. As illustrated in table III, it can be observed that FedFair³ exhibits lower wall clock time, on the average 24.36 %, compared to other algorithms, particularly in non-IID data scenarios. This indicates that FedFair³ achieves faster execution and demonstrates its advantage in terms of time efficiency.

TABLE IV: Global accuracy and accuracy variance for IID and non-IID datasets across MNIST, CIFAR10, and FashionMNIST

CIFAR10 Dataset				
IID			Non-IID	
Alg.	Accuracy	Variance	Accuracy	Variance
FedAvg	86.82 ±0.7	4.81 ±0.61	76.74 ±0.87	17.71±1.02
AFL	86.91 ±0.98	4.16±0.4	76.1±0.3	11.01±1.1
FedProx	86.22 ±0.34	3.95 ±0.92	76.25±0.88	14.73±1.15
Oort	86.13 ±0.97	3.59 ±0.83	75.16 ±0.07	13.13 ±0.54
FedFair ³	86.93 ±0.88	3.25 ±0.6	76.98±0.38	10.32 ±0.3
FashionMNIST Dataset				
IID			Non-IID	
Alg.	Accuracy	Variance	Accuracy	Variance
FedAvg	78.25±0.35	6.27±1.03	76.31±0.75	27.7±0.74
AFL	79.58±0.75	5.62 ±0.93	74.11±0.93	56.47±1.35
FedProx	79.72 ±0.2	4.23 ±0.44	76.5±0.93	53.77±1.03
Oort	75.95±0.67	5.01 ±0.65	75.73 ±0.98	54.98±1.2
FedFair ³	80.79 ±0.84	4.25±0.74	76.8±0.99	10.13±0.93
MNIST Dataset				
IID			Non-IID	
Alg.	Accuracy	Variance	Accuracy	Variance
FedAvg	96.69 ±0.95	6.79±1.07	89.27±0.92	27.61±1.24
AFL	96.59±0.89	5.69 ±1.39	88.92±0.98	24.46±1.04
FedProx	96.23 ±0.85	8.27±1.03	85.01 ±0.76	10.27±0.95
Oort	96.46 ±1.18	7.56 ±0.93	86.73 ±0.83	27.78 ±1.3
FedFair ³	96.74 ±0.8	5.15 ±1.23	89.79±0.78	9.02 ±1.52

V. LIMITATION AND FUTURE WORK

In this section, we outline the limitations of our approach. One of the assumptions of our paper is that $F(w)$ is a convex function, which can be limiting in real-world scenarios. So, we plan to enhance its practical applicability by handling non-convex functions. Furthermore, because of the simplicity of the FedAvg, we excluded this algorithm from the elapsed time figures. The wall clock time of the FedFair³ is slightly more than FedAvg algorithm in some cases. Another challenge of the approach was finding an optimal value for hyper parameters such as preferred T or q; we used some specific values of q to check the result. It is more efficient to check the q value adaptively. Furthermore, we will try to achieve comparable results with less computation overhead in the future leveraging some approaches such as pruning or dropout during model training that have negligible impact on the output.

VI. CONCLUSION

In conclusion, our study has highlighted the importance of fairness in FL and proposed a novel approach to achieve the goal by incorporating client resources and demands. By introducing a weighted loss in the FedFair³ algorithm, we have ensured a fair distribution of accuracy over the clients while optimizing resource utilization and improving system performance. Our experimental results have demonstrated that the FedFair³ algorithm outperforms existing methods including AFL, FedProx, Oort and FedAvg by achieving a more uniform distribution of accuracy over the clients. By considering client resources and demands, the FedFair³ algorithm has been successful in achieving fairness and efficiency in distributed machine learning.

VII. ACKNOWLEDGEMENT

This work was supported in part by the National Science Foundation under Award Numbers OIA-2019511 and OIA-2327452, in part by the Louisiana Board of Regents under Contract Number LEQSF(2019-22)-RD-A-21.

REFERENCES

- [1] F. Lai, X. Zhu, H. V. Madhyastha, and M. Chowdhury, "Oort: Efficient federated learning via guided participant selection." in *OSDI*, 2021, pp. 19–35.
- [2] T. Li, M. Sanjabi, A. Beirami, and V. Smith, "Fair resource allocation in federated learning," *arXiv preprint arXiv:1905.10497*, 2019.
- [3] H. Yu, S. Yang, and S. Zhu, "Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 5693–5700.
- [4] 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*. PMLR, 2017, pp. 1273–1282.
- [5] Y. Zhang, M. J. Wainwright, and J. C. Duchi, "Communication-efficient algorithms for statistical optimization," *Advances in neural information processing systems*, vol. 25, 2012.
- [6] S. Caldas, J. Konečný, H. B. McMahan, and A. Talwalkar, "Expanding the reach of federated learning by reducing client resource requirements," *arXiv preprint arXiv:1812.07210*, 2018.
- [7] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [8] A. Katharopoulos and F. Fleuret, "Biased importance sampling for deep neural network training," *arXiv preprint arXiv:1706.00043*, 2017.
- [9] J. Goetz, K. Malik, D. Bui, S. Moon, H. Liu, and A. Kumar, "Active federated learning," *arXiv preprint arXiv:1909.12641*, 2019.
- [10] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine learning and systems*, vol. 2, pp. 429–450, 2020.
- [11] S. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan, "Adaptive federated optimization," *arXiv preprint arXiv:2003.00295*, 2020.
- [12] M. J. Kusner, J. Loftus, C. Russell, and R. Silva, "Counterfactual fairness," *Advances in neural information processing systems*, vol. 30, 2017.
- [13] S. Barocas, M. Hardt, and A. Narayanan, "Fairness in machine learning," *Nips tutorial*, vol. 1, p. 2017, 2017.
- [14] S. Mitchell, E. Potash, S. Barocas, A. D'Amour, and K. Lum, "Prediction-based decisions and fairness: A catalogue of choices, assumptions, and definitions," *arXiv preprint arXiv:1811.07867*, 2018.
- [15] Y. Shi, H. Yu, and C. Leung, "A survey of fairness-aware federated learning," *arXiv preprint arXiv:2111.01872*, 2021.
- [16] M. Mohri, G. Sivek, and A. T. Suresh, "Agnostic federated learning," in *International Conference on Machine Learning*. PMLR, 2019, pp. 4615–4625.
- [17] L. Wang, Z. Wang, and X. Tang, "Fedeba+: Towards fair and effective federated learning via entropy-based model," *arXiv preprint arXiv:2301.12407*, 2023.
- [18] W. Chu, C. Xie, B. Wang, L. Li, L. Yin, H. Zhao, and B. Li, "Focus: Fairness via agent-awareness for federated learning on heterogeneous data," *arXiv preprint arXiv:2207.10265*, 2022.
- [19] J. Zhang, C. Li, A. Robles-Kelly, and M. Kankanhalli, "Hierarchically fair federated learning," *arXiv preprint arXiv:2004.10386*, 2020.
- [20] Z. Fan, H. Fang, Z. Zhou, J. Pei, M. P. Friedlander, C. Liu, and Y. Zhang, "Improving fairness for data valuation in horizontal federated learning," in *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. IEEE, 2022, pp. 2440–2453.
- [21] L. Lyu, X. Xu, Q. Wang, and H. Yu, "Collaborative fairness in federated learning," *Federated Learning: Privacy and Incentive*, pp. 189–204, 2020.
- [22] B. Ray Chaudhury, L. Li, M. Kang, B. Li, and R. Mehta, "Fairness in federated learning via core-stability," *Advances in neural information processing systems*, vol. 35, pp. 5738–5750, 2022.
- [23] A. Papadaki, N. Martinez, M. A. Bertran, G. Sapiro, and M. R. Rodrigues, "Federated fairness without access to demographics," in *Workshop on Federated Learning: Recent Advances and New Challenges (in Conjunction with NeurIPS 2022)*, 2022.
- [24] D. Y. Zhang, Z. Kou, and D. Wang, "Fairfl: A fair federated learning approach to reducing demographic bias in privacy-sensitive classification models," in *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, 2020, pp. 1051–1060.
- [25] A. Armacki, D. Bajovic, D. Jakovetic, and S. Kar, "Personalized federated learning via convex clustering," in *2022 IEEE International Smart Cities Conference (ISC2)*. IEEE, 2022, pp. 1–7.
- [26] S. Wang, T. Tuor, T. Salonidis, K. K. Leung, C. Makaya, T. He, and K. Chan, "Adaptive federated learning in resource constrained edge computing systems," *IEEE journal on selected areas in communications*, vol. 37, no. 6, pp. 1205–1221, 2019.
- [27] L. Deng, "The mnist database of handwritten digit images for machine learning research [best of the web]," *IEEE signal processing magazine*, vol. 29, no. 6, pp. 141–142, 2012.
- [28] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *arXiv preprint arXiv:1708.07747*, 2017.
- [29] A. Krizhevsky, G. Hinton *et al.*, "Learning multiple layers of features from tiny images," *cs.utoronto.ca*, 2009.
- [30] N. Su and B. Li, "Asynchronous federated unlearning," *IEEE International Conference on Computer Communications*, 2023.
- [31] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural computation*, vol. 1, no. 4, pp. 541–551, 1989.
- [32] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.