

FEDERATED LEARNING WITH CLASS IMBALANCE REDUCTION

Miao Yang^{†*} Ximin Wang[†] Hongbin Zhu[†] Haifeng Wang[‡] Hua Qian^{*†}

[†]School of Information Science and Technology, ShanghaiTech University

^{*}Shanghai Advanced Research Institute, Chinese Academy of Sciences (CAS)

[‡]Key Laboratory of Wireless Sensor Network Communication, CAS

Emails: {yangmiao, wangxm1, zhuhb1, qianhua}@shanghaitech.edu.cn,
haifeng.wang@wico.sh

ABSTRACT

Federated learning (FL) is a promising technique that enables a large amount of edge computing devices to collaboratively train a global learning model. Due to the communication limitation, only a subset of devices can be engaged to train and transmit the trained model to centralized server for aggregation. Since the local data distribution varies among all devices, class imbalance problem arises along with the unfavorable client selection, resulting in a slow converge rate of the global model. In this paper, we design an estimation scheme to reveal the class distribution without the awareness of raw data. According to the estimation scheme, we propose a multi-arm bandit based algorithm that can select the client set with minimal class imbalance. The proposed algorithm can significantly improve the convergence performance of the global model. Simulation results demonstrate the effectiveness of the proposed algorithm.

Index Terms— federated learning, deep neural networks, privacy concerns, class imbalance, client scheduling, multi-armed bandit.

1. INTRODUCTION

With the growing amount of applications, end devices, e.g., smart phones, tablets, or vehicles, generate massive private data in daily life [1]. The valuable personal data can be harnessed to train the machine learning model and significantly improve the quality of end-users' experience. However, due to privacy concerns, transmitting private data from local client devices to the cloud server is not feasible and appropriate.

To efficiently utilize the end-users' data, federated learning (FL) has emerged as a new paradigm of distributed machine learning that executes model training with the private data in local devices [2]. Engaging FL, locally training model can be transmitted to global server for model aggregation without any data information. Unlike the server-based applications, FL on end devices poses several fundamental challenges, such as the limited connectivity of wireless networks, unstable availability of end devices, and the non identically and independently distributed (non-IID) distributions of client dataset. The concerns mentioned above prohibits the model training on all participating devices from beginning to end. To avoid such a dilemma, only a subset of devices is selected to participate in each round of model training which is a common practice in FL [3].

In recent years, a variety of client selection schemes in FL have been advocated. In [4], the authors proposed a client scheduling approach to achieve a proper trade-off between the learning efficiency and latency per round. The staleness of the received models and instantaneous channel qualities were jointly considered in [5]. Besides, the work in [6] investigated the problem of minimizing energy consumption of edge devices in FL without compromising learning performance. The client data's distribution of the above FL methods is IID, which may not be applicable in practice.

The device usage pattern of different users varies, so the data samples and labels on any individual device may follow a different distribution. A individual data distribution cannot represent the global data distribution. Recently, it has been pointed out that the performance of FL, especially federated averaging (FedAvg) algorithm [7], may significantly degrade in the presence of non-IID data, in terms of the model accuracy and the communication rounds required for convergence [8–10]. Recognizing such criticality, the authors in [11] designed a scheme to mitigate the impact of the class imbalance and introduce a loss function to evaluate the class imbalance. Moreover, the work in [12] exploited learning algorithm for client selection to decrease the communication rounds with target accuracy. However, the algorithm in [12] needs lots of offline training and the data on each device must remain unchanged during the training procedure.

In this paper, the client selection problem concentrated on class imbalance in FL is investigated. With the concern of the users' privacy, we propose a scheme that can reveal the severity of class imbalance without any raw data of client devices. Besides, utilizing reinforcement learning, we propose a client selection scheme to minimize the effect of class imbalance. The proposed algorithm endeavors to learn the class distribution and selects the most balanced clients combination.

The remainder of this paper is organized as follows. The system model is presented in section 2. In section 3, we introduce a scheme to evaluate the class imbalance of clients and propose a learning algorithm to find the best client set. The performance of the proposed algorithm is numerically evaluated in section 4. In section 5, we conclude this paper.

2. SYSTEM MODEL

Consider training a deep neural network (DNN) under FL settings with a set of clients $\mathcal{K} = \{1, 2, \dots, K\}$, each with its own local dataset and a global server for model aggregation, as illustrated in Fig. 1. We model the channel from the device to global server as multiple access channel. With the communication limitation, only a

This work was supported in part by the National Natural Science Foundation of China (Grant No. 61671436) and the Science and Technology Commission Foundation of Shanghai (Grant No. 19DZ1204300).

fixed amount of spectrum is available [13]. Due to the scarce of the spectrum resource, the number of available channels is much smaller than that of the client devices.

We then formally introduce the DNN training for multiclass classification problem in FL. Consider a class classification problem defined over a compact feature space \mathcal{X} and a label space $\mathcal{Y} = \mathcal{C}$ with C classes, where $\mathcal{C} = \{1, \dots, C\}$. Let (\mathbf{x}, y) denote a particular labeled sample generated from distribution p . A function $f : \mathcal{X} \rightarrow \mathcal{D}$ maps \mathbf{x} to the probability for the i th class, where $\mathcal{D} = \{z | \sum_{i=1}^C z_i = 1, z_i \geq 0, \forall i \in \mathcal{C}\}$. Let \mathbf{W} denote the weight of DNN. The cross entropy loss can be harnessed to evaluate the training performance in classification, which is defined as

$$L(\mathbf{w}) = E_{\mathbf{x}, y \sim p} \left[\sum_{i=1}^C \mathbb{1}_{y=i} \log f_i(\mathbf{x}, \mathbf{W}) \right] \\ = \sum_{i=1}^C p(y=i) E_{\mathbf{x} | y=i} [\log f_i(\mathbf{x}, \mathbf{W})]. \quad (1)$$

where $E[\cdot]$ denotes the expectation operation and $\mathbb{1}$ denotes the indicator matrix.

In FL, the training procedure is an iterative process consisting of a number of communication rounds. Let \mathbf{W}_t^k and \mathbf{W}_t^g denote the weights of k th device's DNN model and global DNN model in round t , respectively. In round t , the server chooses a subset $\mathcal{S}_t \subseteq \mathcal{K}$ of the clients and then distributes the weights \mathbf{W}_{t-1}^g of the global model to the selected clients. Then the selected clients synchronize their local models such that $\mathbf{W}_{t-1}^k = \mathbf{W}_{t-1}^g$, and perform the following stochastic gradient descent (SGD) training as

$$\mathbf{W}_t^k = \mathbf{W}_{t-1}^k - \eta_t^k \nabla_{\mathbf{W}} L(\mathbf{W}_{t-1}^k; \xi^k), \quad (2)$$

where η_t^k is the learning rate setting for device k in round t , ξ^k is an example sampled from local dataset in device k .

Once \mathbf{W}_t^k is obtained, each participating device k updates its own model weight difference Δ_t^k to the global server, which is defined as

$$\Delta_t^k = \mathbf{W}_t^k - \mathbf{W}_{t-1}^k. \quad (3)$$

When the global server collects all the updates from client devices, it performs the FedAvg algorithm to update the global model as follows [7]

$$\Delta_t^g = \sum_{k \in \mathcal{S}_t} \frac{n_k}{\sum_{k'=1}^K n_{k'}} \Delta_t^k, \quad (4)$$

$$\mathbf{W}_t^g \leftarrow \mathbf{W}_{t-1}^g + \Delta_t^g. \quad (5)$$

where n_k denotes the number of data samples of device k .

When the data and label distribution on different devices are IID, FedAvg has been shown to perform well approximating the model trained on centrally data [7]. However, in practice, data owned by each device are typically non-IID, i.e., the data distributions of clients vary from different devices due to different user preferences and usage patterns. When data distributions are non-IID, FedAvg algorithm is unstable and may even diverge [8].

The problem arises by the inconsistency between the locally performed SGD algorithm, which aims to minimize the loss value on local samples on each device and the global objective of minimizing the overall loss on server data samples. Since the training model is fitted on different devices to heterogeneous local data, the divergence among the weights \mathbf{W}^k of these local models will be accumulated and eventually degrades the performance of the learning process [14].




Fig. 1. An illustration of FL procedure.

3. ONLINE LEARNING ALGORITHM FOR CLIENT SELECTION

In this section, we first discuss the approach to estimate the class imbalance of each training client. Then we propose an online learning algorithm to find the most balanced clients set.

3.1. Class Estimation Scheme

In FL settings, the raw data of clients could not be obtained due to privacy concerns. However, with the help of the following scheme, we can reveal the class distribution of client services according to their updated gradients.

Consider a DNN contains an input layer, a hidden layer and an output layer. Note that in DNN for the classification problem, the scale of the output layer is equal to the size of class label. That is, each neuron of the output layer corresponding to a class label. Let $\mathbf{W} = [\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_C]$ denotes the weights between hidden layer and output layer. Every element of \mathbf{W} denote the weights connect hidden layer to the specific neuron of output layer.

Without loss of generality, we assume there exists a balanced auxiliary dataset with C classes in global server. Such auxiliary dataset only consists of a few data samples and can be extracted from the test dataset. When the auxiliary data examples are fed to the updated model, we can obtain the gradients vector brought by auxiliary data with respect to the corresponding classes. The gradients vector brought by auxiliary data can be expressed as

$$\left\{ \nabla L^{\text{aux}}(\mathbf{w}_1), \nabla L^{\text{aux}}(\mathbf{w}_2), \dots, \nabla L^{\text{aux}}(\mathbf{w}_C) \right\},$$

where each $\nabla L^{\text{aux}}(\mathbf{w}_i)$ is related to the class \mathcal{C}_i . The following theorem can help to obtain the class distribution from the above gradients vector.

Fact 1. When training DNN in classification problem, the expectations of gradient square for different classes have the following approximate relation [15]:

$$\frac{E \|\nabla L(\mathbf{w}_i)\|^2}{E \|\nabla L(\mathbf{w}_j)\|^2} \approx \frac{n_i^2}{n_j^2}, \quad (6)$$

where L denotes the cost function of the neural network, n_i and n_j are the number of samples for class i and class j , respectively, where $i \neq j$ and $i, j \in \mathcal{C}$.

Fact 1 reveals the correlation between the gradients and class distribution. Then for class \mathcal{C}_i , the estimation of class ratio $\frac{n_i^2}{\sum_j n_j^2}$

can be defined

$$R_i = \frac{e^{\frac{\beta}{\|\nabla L^{aux}(\mathbf{w}_i)\|^2}}}{\sum_j e^{\frac{\beta}{\|\nabla L^{aux}(\mathbf{w}_j)\|^2}}}, \quad (7)$$

where β is a hyperparameter that can be tuned to control the normalization between classes. Then we can obtain the composition vector $\mathbf{R} = [R_1, \dots, R_C]$ that indicates the distribution of raw data. Moreover, the Kullback-Leibler (KL) divergence can be harnessed to evaluate the class imbalance of each client, which is defined as

$$D_{KL}(\mathbf{R}||\mathbf{U}) = \sum_{i \in C} R_i \log \frac{R_i}{U_i}, \quad (8)$$

where \mathbf{U} is a vector of ones with magnitude C .

3.2. Online Learning Framework for Client Selection

Multi-arm bandit problems are motivated by a variety of real-world problems, such as online advertising, dynamic pricing and stock investment. To find the optimal balance client subset, the difficulty lies in how to learn the statistics of class distribution via iterative updated gradient. Combinatorial multi-Armed bandit (CMAB) can be harnessed to solve this problem [16]. In CMAB model, an agent gambles on a bandit machine with a finite set of arms, where each arm has unknown distribution. At round t , several arms defined as a super arm \mathcal{S}_t ($\mathcal{S}_t \subseteq \mathcal{K}$) can be pulled. The reward of the super arm depends on the outcomes of all pulled arms.

In this work, we consider the client selection as a CMAB problem, where each client represents the arm and the client set represents the super arm. In FL training, once the model is updated, the server could obtain the local model of each client device. Utilizing the class estimation scheme, we can reveal the composition vector \mathbf{R}^k of selected client k . Define the reward of client k as

$$r^k = \frac{1}{D_{KL}(\mathbf{R}^k||\mathbf{U})}. \quad (9)$$

Then the global server executes FedAvg algorithm for model aggregation and the reward r of the whole client set can be obtained the same as (9). Note that the super arm's reward is a nonlinear combination of the selected single arms' reward. An algorithm that has good theoretical results with nonlinear reward is the combinatorial upper confidence bounds (CUCB) algorithm [16].

Let T^k denote the number of times that the client k has been selected. Once client k has been selected in a time slot, $T^k \rightarrow T^k + 1$, otherwise, $T^k \rightarrow T^k$. The proposed client selection algorithm based CUCB is shown in **Algorithm 1**.

Algorithm 1 CUCB for Client Selection

- 1: For each client k , choose an arbitrary set $\mathcal{S} \in \mathcal{K}$ such that $k \in \mathcal{S}$ and update variables T^k and \hat{r}^k .
 - 2: $t \leftarrow N$.
 - 3: **while true do**
 - 4: $t \leftarrow t + 1$.
 - 5: For each client i , set $\hat{r}^k = \bar{r}^k + \alpha \sqrt{\frac{3 \ln t}{2T^k}}$.
 - 6: Obtain \mathcal{S}_t using **Algorithm 2** with \hat{r} .
 - 7: Play \mathcal{S}_t and update T^k and \hat{r}^k .
 - 8: **end while**
-

Here α is the exploration factor to balance the trade-off between exploitation and exploration. Step 1 of **Algorithm 1** guarantees that

each client has been selected once at least in the first N round. Then the time slot moves to N . In step 5, notation \bar{r}^k denotes the individual reward sample mean of client k , and \hat{r}^k denotes the perturbed version of \bar{r}^k . The perturbation \hat{r}^k in step 5 promotes the selection of clients that are not selected frequently, by artificially increasing their expected reward estimates. In step 6, the proposed algorithm utilizes the perturbed \hat{r}^k to solve the client selection problem.

In our client selection problem, the class distribution of each client is uncertain. Fortunately, we can reveal the class distribution according to the composition vector. Let $\mathbf{R}^k(t)$ represent the composition vector of client k at time slot t . Thus, the class ratio can be estimated by the sample mean of composition vector, which can be expressed as

$$\bar{\mathbf{R}}^k = \frac{\sum_{t=1}^{T^k} \mathbf{R}^k(t)}{T^k}. \quad (10)$$

With the estimated composition vector $\bar{\mathbf{R}}$ and reward r of each client, we can design the client selection scheme with minimal class imbalance according to **Algorithm 2**.

Algorithm 2 Class Balancing Algorithm

- 1: **Initialize:**
Set $\mathcal{S}_t = \emptyset$ and $\mathbf{R}_{total} = \emptyset$.
 - 2: $k_0 = \arg \max_k \hat{r}^k$.
 - 3: $\mathcal{S}_t \leftarrow \mathcal{S}_t \cup \{k_0\}$.
 - 4: **while** $|\mathcal{S}_t| < K$ **do**
 - 5: Select $k_{min} = \arg \min_k D_{KL}(\mathbf{R}_{total} + \bar{\mathbf{R}}^k || \mathbf{U})$ for $k \in \mathcal{K} \setminus \mathcal{S}_t$.
 - 6: Set $\mathcal{S}_t \leftarrow \mathcal{S}_t \cup \{k_{min}\}$, $\mathbf{R}_{total} \leftarrow \mathbf{R}_{total} + \bar{\mathbf{R}}_{min}^k$.
 - 7: **end while**
 - 8: **Outputs:** \mathcal{S}_t
-

Algorithm 2 determines the participants of FL updating. Combine **Algorithm 1** and **Algorithm 2**, we can find the most suitable client set with class balance.

4. NUMERICAL RESULTS

In this section, we present numerical results to validate the effectiveness of the proposed algorithms¹.

We test our scheme on one of the main benchmarks: CIFAR10 [17]. CIFAR10 dataset consists of 50000 training examples and 10000 testing examples of 32×32 RGB images, categorized by total 10 classes. The number of clients in FL training is set as 100. To model the imbalanced class distribution, we split the whole CIFAR10 dataset to each client with random amount of classes and random amount of data samples. That is, the data distribution of each client is different. Besides, the sizes of auxiliary dataset is 320 with 10 classes equally distributed.

The architecture of our deep model is a standard convolutional neural network (CNN), which comprises 3 convolutional layers followed by rectified linear units (ReLU) nonlinear activations and max-pooling layer and 2 fully connected layers, with totally 122570 parameters. Such standard model can meet our needs to validate the effectiveness of our scheme. We use standard SGD as our optimizer. The learning rate and the learning rate decay of SGD are set as 0.1 and 0.996, respectively. In each training round, the selected clients

¹The source code of our work can be found in <https://github.com/ym1231/fl-cir>.

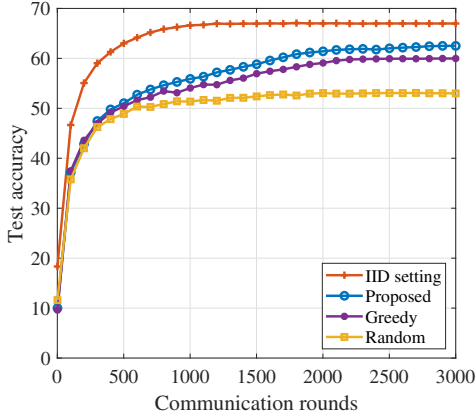


Fig. 2. The global test accuracy with different selection schemes.

train their local models for 5 epochs. The client selects 10 batches with batch size 10 at each training epoch.

In our simulation, we set exploration factor α as 0.2. We fix the normalization hyperparameter β as 1. To evaluate the performance of the proposed algorithm, we compare the performance of the proposed algorithm to the following schemes. (i) Greedy scheme: global server selects the client set with the sample mean information according to the class estimation method; (ii) Random scheme: global server randomly selects the client set. Note that the original FedAvg algorithm engages the random scheme to select devices. In addition, we compare the proposed algorithm in IID setting to show the performance gap. Note that in IID setting, the class distribution and the number of data samples in each client are set as the same. Thus the above selection schemes are the same in practice.

In the first experiment, we examine the convergence performance of global model with the proposed algorithm. At each round, the global server selects 20 clients for model aggregation. Fig. 2 depicts the test accuracy of global model with different schemes. The test accuracy performance represents the learning process of global model. From Fig. 2, we observe that the proposed algorithm achieves faster convergence speed and higher test accuracy compared with the greedy and uniform schemes. The discrepancy of performance between the proposed and random scheme comes from the effect of class imbalance. The proposed algorithm can reduce the class imbalance by carefully selecting the proper client set. Compared with the greedy scheme, the proposed algorithm can balance the trade-off between exploration and exploitation, thus resulting in the exploration of more suitable clients set with balanced class combination.

In the second experiment, we would like to verify the performance of the proposed algorithm with different amounts of selected clients. Fig. 3 plot the FL training performance with respect to the amounts of clients. With the increasing amounts of selected clients, the FL training process can achieve better performance. However, the performance improvement is slighter when the amount of clients become larger. This result indicates that too large amount of clients is not essential to find the best client set when suffering the burden of communication consumption.

We investigate the accuracy performance of the proposed algorithm with different parameters α in Fig. 4. The parameters α can decide the trade-off between exploitation and exploration. When the exploration parameter α is small, the global server would like to

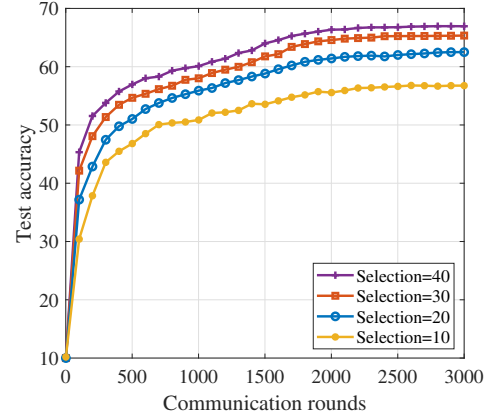


Fig. 3. The global test accuracy of the proposed algorithm with different amounts of clients selection.




Fig. 4. The global test accuracy with different exploration factors.

exploit the history sampling information and thus could not explore sufficiently to find the best client set. As the exploration parameter α increases, the global server prefers to explore the clients with fewer selections. Spending many rounds for exploration could result in the performance deterioration since the clients selected for exploration may be not always proper. A suitable α is essential to improve the convergence performance.

5. CONCLUSIONS

In this paper, we studied the client selection problem with class imbalance in FL. Without the requirement of clients' data information, we designed a scheme to explicitly reveal the class distribution according to the updated gradients. Besides, a client selection algorithm based on the CMAB framework was proposed to reduce the class imbalance effect. Numerical results confirmed that the proposed algorithm could pick the properly balanced client set and improve the convergence performance of the global model.

6. REFERENCES

- [1] M. Chiang and T. Zhang, "Fog and IoT: An overview of research opportunities," *IEEE Internet of Things Journal*, vol. 3, no. 6, pp. 854–864, Dec. 2016.
- [2] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon, "Federated learning: Strategies for improving communication efficiency," *Availabel as ArXiv:1610.05492*, 2016.
- [3] Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally, "Deep gradient compression: Reducing the communication bandwidth for distributed training," *Availabel as ArXiv:1712.01887*, 2017.
- [4] Wenqi Shi, Sheng Zhou, and Zhisheng Niu, "Device scheduling with fast convergence for wireless federated learning," in *IEEE International Conference on Communications (ICC)*, May 2020.
- [5] Howard H Yang, Ahmed Arafa, Tony QS Quek, and H Vincent Poor, "Age-based scheduling policy for federated learning in mobile edge networks," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2020.
- [6] Qunsong Zeng, Yuqing Du, Kaibin Huang, and Kin K Leung, "Energy-efficient radio resource allocation for federated edge learning," in *IEEE International Conference on Communications Workshops (ICC Workshops)*, May 2020.
- [7] 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 (AISTATS)*, Aug. 2017.
- [8] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra, "Federated learning with non-iid data," *Availabel as ArXiv:1806.00582*, 2018.
- [9] F. Sattler, S. Wiedemann, K. R. Müller, and W. Samek, "Robust and communication-efficient federated learning from non-i.i.d. data," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 9, pp. 3400–3413, 2020.
- [10] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang, "On the convergence of fedavg on non-iid data," *Availabel as ArXiv:1907.02189*, 2019.
- [11] Lixu Wang, Shichao Xu, Xiao Wang, and Qi Zhu, "Towards class imbalance in federated learning," *Availabel as ArXiv:2008.06217*, 2020.
- [12] Hao Wang, Zakhary Kaplan, Di Niu, and Baochun Li, "Optimizing federated learning on non-iid data with reinforcement learning," in *IEEE Conference on Computer Communications (INFOCOM)*, Jul. 2020.
- [13] Wenchao Xia, Tony QS Quek, Kun Guo, Wanli Wen, Howard H Yang, and Hongbo Zhu, "Multi-armed bandit based client scheduling for federated learning," *IEEE Transactions on Wireless Communications*, 2020.
- [14] Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh, "Agnostic federated learning," *Availabel as ArXiv:1902.00146*, 2019.
- [15] Rangachari Anand, Kishan G Mehrotra, Chilukuri K Mohan, and Sanjay Ranka, "An improved algorithm for neural network classification of imbalanced training sets," *IEEE Transactions on Neural Networks*, vol. 4, no. 6, pp. 962–969, 1993.
- [16] Wei Chen, Yajun Wang, and Yang Yuan, "Combinatorial multi-armed bandit: General framework and applications," in *International Conference on Machine Learning (ICML)*, Jun. 2013.
- [17] A. Krizhevsky and G.E. Hinton, "Learning multiple layers of features from tiny images," in *Technical report*, 2009.