

FEDNS: IMPROVING FEDERATED LEARNING FOR COLLABORATIVE IMAGE CLASSIFICATION ON MOBILE CLIENTS

Yaoxin Zhuo, Baoxin Li

Arizona State University
{yzhuo6, baoxin.li}@asu.edu

ABSTRACT

Federated Learning (FL) is a paradigm that aims to support loosely connected clients in learning a global model collaboratively with the help of a centralized server. The most popular FL algorithm is Federated Averaging (FedAvg), which is based on taking weighted average of the client models, with the weights determined largely based on dataset sizes at the clients. In this paper, we propose a new approach, termed Federated Node Selection (FedNS), for the server's global model aggregation in the FL setting. FedNS filters and re-weights the clients' models at the node/kernel level, hence leading to a potentially better global model by fusing the best components of the clients. Using collaborative image classification as an example, we show with experiments from multiple datasets and networks that FedNS can consistently achieve improved performance over FedAvg.

Index Terms— Federated learning, Model aggregation, Image classification

1. INTRODUCTION

Mobile devices with cameras have become ubiquitous and they potentially produce the amount of data necessary for training deep learning models that have seen increasing adoption in applications such as computer vision and robotics. In most real applications, it is not realistic to expect that a centralized server may train a global model by collecting data from all involved mobile clients, due to bandwidth constraint (e.g., too much traffic between the server and millions of clients) and/or privacy concerns (e.g., users may not be willing to share images from their devices). The *Federated Learning* [1] was proposed and advocated as a good strategy to address these issues, which allows loosely connected clients to contribute to the learning of a global model collaboratively.

In a typical FL setup, training data are distributed over clients, each of which learns to update a global model using its local data. A server maintains and updates a global

model by aggregating the local models from all the clients. A commonly used algorithm, named *Federated Averaging* (FedAvg) [1], updates the global model by a weighted average of the client models, where the weights are proportional to the amount of data used by each involved client (in real implementations, not all clients are involved in every update).

While FedAvg delivered state-of-the-art results in some reports, it still has some limitations. First, its assumption about the data distributions across clients is often violated in real-world scenarios. Second, simple weighted average of client models may not be optimal, as analyzed in some recent work [2]. There are also other variants of FL, such as CMFL [3] and FedMA [4], where typically additional constraints or processing steps are introduced (e.g., more hyper-parameters and more computations at the clients), making the method less generally applicable compared with the simpler FedAvg.

In this paper, we propose a new approach, termed Federated Node Selection (FedNS), for the server's global model aggregation in the FL setting. The approach filters and re-weights the clients' models at the node/kernel level on the server, hence leading to a potentially better global model by fusing the best components of the clients. We focus on a better way of module aggregation on the server while keeping the simplicity of the FL paradigm. Using collaborative image classification as a case study, we demonstrate the benefits of the proposed method from the following perspectives. First, the learning paradigm follows the FL framework without extra computation requirement for the clients or extra communication burden between the clients and the server. Second, our method can handle iid or non-iid distributions without pre-defined hyper-parameters. Third, our method allows the clients to have varying amounts of data, making it closer to real-world scenarios (where mobile clients could see and process dynamically changing data every day). Experiments with three different datasets and networks are used to assess the performance of the proposed method.

The remainder of this paper is organized as follows. In Section 2, we review related work and discuss the disadvantages wherever applicable. In Section 3, we first introduce common baselines and then present our method FedNS. We design and perform experiments in Section 4 for evaluation. And finally, we conclude with Section 5.

978-1-6654-3864-3/21/\$31.00 ©2021 IEEE

The work was supported in part by a grant from ONR (N00014-19-1-2119). Any opinions expressed in this material are those of the authors and do not necessarily reflect the views of ONR.

2. RELATED WORK

Since the initial paper by McMahan *et al.* [1], FL has received considerable attention and there have been numerous research papers in this regard, of which we can only briefly review some that are most relevant to our work. A couple of recent surveys [5, 6, 7] may provide an ample picture for an interested reader. One line of research has been on saving the communication between the clients and the server. In [8], a compressed structure was designed to reduce the communication cost. Horvath *et al.* [9] introduced another compression technique for lowering the training time. Alistarh *et al.* [10] proposed a family of compression schemes (QSGD), allowing a user to trade off between bandwidth and communication rounds (and hence modifying the basic FL framework). Federated Dropout [11] was proposed to allow the clients to train on only a subset of the global model, reducing parameters that need to be communicated to the server. In general, as the focus is on saving communication, such methods typically do not improve the learned global model, if they do not cause degradation.

Another direction is on improving the effectiveness of FL. Wang *et al.* [3] proposed CMFL that picks a fraction of all communicated clients to update the global model. The main limitation is the need for searching optimal hyper-parameters (thresholds). Nishio *et al.* [12] tried to do client-level selective aggregation, based on the client's computing capability. Zhao *et al.* [13] designed a strategy for improving accuracy when every client has only a single class of data. In [14], a method was proposed to filter the local data at client. But the filtering is based on pre-trained models. Yao *et al.* [15] proposed FedFusion, an approach that could improve the performance by fusing the two model features, although extra computation is needed for choosing the best fusion approach. FedMA [4] constructs the global model in a layer-wise manner, which appears to cause significant burden on the clients.

There are also some efforts on improving the FL for specific tasks/applications. For example, Huang *et al.* [16] designed a community-based FL algorithm that clusters distributed patient data into meaningful communities. Jenny *et al.* [17] designed an approach for reducing the cost of communication, focusing on enabling a large model to be learned on relatively small memory of the clients.

While some of the aforementioned methods improve upon the more foundational FedAvg approach under peculiar settings, they often do not readily apply to more general FL settings, owing to the assumptions they introduced to constrain the problem. Our objective of this research is to design a method that focuses on better model aggregation with minimal modification to the basic FL framework. For example, we relax the assumption about the knowledge of the data distribution. Hence our method has a better potential for wider applicability, especially for real-world applications.

3. PROPOSED METHOD

In this section, we begin by introducing the baseline method of FL: Federated Averaging (FedAvg) in Section 3.1. Then we present a straightforward extension in Section 3.2, termed FedAvg+lastFC, which attempts to do a better weighted averaging for the last FC layer under the non-iid condition. We propose our FedNS method in Section 3.3, which is an improvement over FedAvg+lastFC in that better weighted averaging is done for all layers.

3.1. The FedAvg Baseline

FedAvg was first proposed in [1], where the problem was formulated as one of optimizing the following non-convex neural network objective function:

$$f(w) = \frac{1}{K} \sum_{k=1}^K \ell(x_k, y_k; w_k) \quad (1)$$

where k is the client index ranging from 1 to K .

The loss function $\ell(x_k, y_k; w_k)$ is the k -th client's loss, i.e., the loss of model w_k under data x_k with corresponding label y_k . In the first step of FedAvg, the server initializes a global model w_0 and sends it to a fraction of communicated clients. Each communicated client k will train the model with its local data, i.e., calculating the local gradient g_k^i and updating the local model after training certain E epochs (assuming with the learning rate η):

$$w_{k,t}^{i+1} = w_{k,t}^i - \eta g_k^i \quad (2)$$

where t indexes the communication rounds, ranging from 1 to N , and i is the client's local epoch range from 1 to E .

The initial model of the client is the weights received from the server. The clients will upload their updated local models back to the server. The server will then aggregate models from all clients using a weighted averaging scheme:

$$w_{s,t+1} = \sum_{k=1}^K \frac{n_k}{n} w_{k,t} \quad (3)$$

where w_s is the weights of the global model on the server, n_k is the total number of training samples on client k and n is the sum of all samples from all communicated clients. It is worth noting that Eq.n. 3 will become simple averaging when all clients have the same amount of data. In the next communication round, the server will start from sending aggregated server model $w_{s,t}$ to all the clients and repeat the previous steps until reach the target communication rounds. The complete pseudo-code for this baseline is given in Algorithm1.

3.2. FedAvg+lastFC: A Simple Extension to FedAvg

One may build a naive and straightforward extension to FedAvg for potential improvement, by using weights based on the number of samples for each class (instead of the number

Algorithm 1 Federated Averaging (FedAvg)

Server executes:initialize w_0 as $w_{s,1}$ **for** each communication round t from 1 to N **do** $S_t \leftarrow$ (the fraction of communicated clients)**for** each client $k \in S_t$ **in parallel do** $w_{k,t} \leftarrow \text{ClientUpdate}(k, w_{s,t})$ **end for****Model Aggregation:** $w_{s,t+1} \leftarrow \sum_{k=1}^K \frac{n_k}{n} w_{k,t}$ **end for****ClientUpdate**($k, w_{k,t}$): $B \leftarrow$ (split local data into batches of local batch size B)**for** each local epoch i from 1 to E **do****for** batch $b \in B$ **do** $g_k^i = \nabla(w_{k,t}^i; b)$ $w_{k,t}^{i+1} \leftarrow w_{k,t}^i - \eta g_k^i$ **end for****end for**return $w_{k,t}$ to server

of samples) at each client when doing the weighted averaging for the last fully connected (FC) layer. That is, we define

$$w_{s,t+1}^{l=L,c} = \sum_{k=1}^K \frac{n_k^c}{n^c} w_{k,t}^{l=L,c} \quad (4)$$

where l is the index of layers in a model that has L layers, c is the index for the last FC layer's hidden nodes, ranging from 1 to C . And C is also equals to the total number of classes in image classification models. So n_k^c is client k 's number of samples in the class that corresponding to node c and n^c is the total number of data in class c within all communicated clients. This leads to Algorithm 2, where only the model aggregation pseudo-code is given to save space.

3.3. FedNS: the Proposed Method

While FedAvg+lastFC employs an intuitive way of weighted averaging for the last FC layer, the method cannot be easily extended to previous layers since they lack explicit correspondence to the number of classes. In the meantime, using only the number of samples at each client for weighted averaging is not necessarily optimal. Hence we propose to use the variance of the weight update to evaluate the contribution of a node/kernel to the client's local model (for layers other than the last FC layer). During the aggregation stage, the server could easily calculate the differential weights for each node and calculate its variance (and derive from these weights for the clients) since it has the global model from the previous communication round. We term this approach Federated Node Selection (FedNS). It first calculates the fol-

Algorithm 2 Federated Averaging with weighted averaging last FC layer (FedAvg+lastFC)

Server executes:

.....

Model Aggregation:**for** each layer l from 1 to L **do****if** $l < N$ **then** $w_{s,t+1} \leftarrow \sum_{k=1}^K \frac{n_k}{n} w_{k,t}$ **else****for** each node c in the layer L from 1 to C **do** $w_{s,t+1}^{l=L,c} = \sum_{k=1}^K \frac{n_k^c}{n^c} w_{k,t}^{l=L,c}$ **end for****end if****end for**.....

lowing:

$$v_{k,t+1}^{l \neq L,c} = \text{variance}(w_{k,t+1}^{l \neq L,c} - w_{k,t}^{l \neq L,c}) \quad (5)$$

where c is the node index ranges from 1 to C . C is the total node/kernel number for the layer l in a model that has L layers. We calculate their mean ($\mu_{t+1}^{l \neq L,c}$) and sigma ($\sigma_{t+1}^{l \neq L,c}$) after getting all the variances. If the node's variance is out of range $[\mu_{t+1}^{l \neq L,c} - 2\sigma_{t+1}^{l \neq L,c}, \mu_{t+1}^{l \neq L,c} + 2\sigma_{t+1}^{l \neq L,c}]$ (95% of confidence interval if the underlying distribution is normal), we will remove that node in the global model aggregation step. Then we re-normalize the remaining nodes' variances and get the new weights. The server do a new weighted averaging for all layers other than the last FC layer at node/kernel level by the new weights in the following equation:

$$w_{s,t+1}^{l \neq L,c} = \sum_{k=1}^K \frac{v_{k,t}^{l \neq L,c}}{v_t^{l \neq L,c}} w_{k,t}^{l \neq L,c} \quad (6)$$

where c is the remaining node's index, $v_{k,t}^{l \neq L,c}$ is node c 's variance of differential weights. And $v_t^{l \neq L,c}$ is the sum of variances from all remaining nodes. This new weighted averaging method leads to Algorithm 3, where only the model aggregation pseudo-code is given to save space.

4. EXPERIMENTS

In this section, we present the empirical study of FedNS and comprehensively evaluate it. We compare FedNS with FedAvg and FedAvg+lastFC in two aspects: one is the final global model's performance and the other is every communication round's performance. The first one shows the model's performance after FL reaches the target communication rounds. And the second one shows which method is in a leading position during the FL process, which is also an important factor for real applications (e.g., whether a mobile client can still have good performance in a long process of incremental learning).

Algorithm 3 Federated Node Selection (FedNS)

Server executes:

.....

Model Aggregation:

for each layer l from 1 to L **do**

if $l < L$ **then**

for each node c in the layer l from 1 to C **do**

$v_{k,t}^{l \neq L,c} = \text{variance}(w_{k,t}^{l \neq L,c} - w_{k,t-1}^{l \neq L,c})$

 filter nodes and re-normalize remaining variances

end for

$w_{s,t+1}^{l \neq L,c} = \sum_{k=1}^K \frac{v_{k,t}^{l \neq L,c}}{v^c} w_{k,t}^{l \neq L,c}$

else

for each node c in the layer L from 1 to C **do**

$w_{s,t+1}^{l=L,c} = \sum_{k=1}^K \frac{n_k^c}{n^c} w_{k,t}^{l=L,c}$

end for

end if

end for

.....

4.1. Datasets

To do the case study of collaborative image classification on mobile clients, we select three image classification datasets. Different from aforementioned methods that used simple MNIST [18] dataset, we replace it with a more challenging dataset: FashionMNIST [19]. The other two datasets we select are CIFAR10 [20] and tinyImageNet [21]. Both FashionMNIST and CIFAR10 have 10 classes. FashionMNIST has 6,000 training and 1,000 testing images for each class. CIFAR10 has 5,000 training and 1,000 testing images for each class. TinyImageNet is a selective subset of ImageNet [22] and consist of 200 classes' images. It has 500 training and 50 testing images for each class.

4.2. Networks

We use CNN models of different complex levels for the datasets. For the FashionMNIST dataset, we use a simple CNN model that has 2 convolutional layers and 3 fully connected layers. The first 5×5 convolutional layer has 32 channels. The second 5×5 convolutional layer has 64 channels. Both of them are followed with 2×2 ReLU activation and 2×2 max pooling. The 3 FC layers have 1024, 256, and 10 nodes respectively. All of them are followed by ReLU activation. For the CIFAR10 dataset, we use the AlexNet [23]. For the tinyImageNet dataset, we use the ResNet18 [24].

4.3. Simulation Settings

Unlike some existing works that employed simplistic settings of each client having a fixed set of images throughout the entire training process, we allow each client to randomly choose a subset of images as local data in every communication round according to the iid or non-iid conditions. So the experiment

setting is closer to real-world scenarios where mobile clients receive dynamically changing data. Specifically, for FashionMNIST, each client randomly chooses 5 images per class in the iid condition and 1-10 images per class in the non-iid condition; for CIFAR10, each client randomly chooses 50 images per class in the iid condition and 1-100 images per class in the non-iid condition; for tinyImageNet, each client randomly chooses 5 images per class in the iid condition and 1-10 images per class in the non-iid condition.

For other algorithmic parameters, we follow the same settings in [1]. We fix the learning rate η as $1e-2$. And we set communication round (day) $N = 50$, local epoch $E = 5$, local batch size $B = 10$ and the communicated client number $K = 10$, which means the server could randomly communicate with 10 clients in every communication round.

4.4. Evaluation

For FashionMNIST and CIFAR10, we use accuracy, macro precision, macro recall, and macro F-Score to evaluate the models. For tinyImageNet, we use top-1 accuracy and top-5 accuracy to evaluate the models. To visualize the performance difference in every communication round, we plot the differential performance numbers (against the FedAvg baseline). Because of the similar trend of accuracy and macro recall, we omit the macro recall numbers due to space limitation.

4.5. Results

For each dataset and approach, we repeat the experiment ten times and report the averaged results. It is worth noting that FedAvg+lastFC and FedAvg are the same in the iid condition, and thus we only compare FedNS with FedAvg in this case. Figure 1 and Table 1 show the results of FashionMNIST with 5-layer CNN. Figure 2 and Table 2 show the results of CIFAR10 with AlexNet. Figure 3 and Table 3 show the results of tinyImageNet with ResNet18. The tables show the final global model's performance and figures show the progressive performance differences for the methods in each communication round. From the figures and tables, we observe that overall FedNS outperforms the other two methods, especially in the non-iid conditions. Please refer to the supplemental materials for higher-resolution versions of the plots.

5. CONCLUSION AND FUTURE WORK

We proposed FedNS as a new model aggregation method for FL. Without much modification to the FL framework, our method can be easily applied to applications like mobile image classification. Different from the traditional FL setting where each client has a fixed fraction of the data, our approach allows flexible and varying data availability, making it closer to the real-world scenario. Even under such challenging settings, experiments showed that FedNS improves

Table 1. The final performance comparison in FashionMNIST dataset with 5-layer CNN in iid and no-iid conditions

Fashion MNIST	iid		non-iid		
	FedAvg	FedNS	FedAvg	FedAvg +lastFC	FedNS
accuracy	83.53	83.79	82.96	83.50	83.56
macro precision	0.834	0.838	0.830	0.837	0.838
macro F-Score	0.931	0.936	0.926	0.933	0.929

Table 2. The final performance comparison in CIFAR10 dataset with AlexNet in iid and non-iid condition

CIFAR10	iid		non-iid		
	FedAvg	FedNS	FedAvg	FedAvg +lastFC	FedNS
accuracy	66.09	67.07	65.86	65.57	67.55
macro precision	0.658	0.668	0.660	0.658	0.671
macro F-Score	0.740	0.745	0.737	0.737	0.745

upon the baselines, both during the process of learning and at the end of learning. One future direction may be to explore how to extend FedNS to handle extreme data distributions, such as when each client has distinct classes. In such cases, a more subtle question of comparing models learned from distinct classes needs to be addressed.

6. REFERENCES

- [1] Brendan McMahan, Eider Moore, Daniel Ramage, et al., “Communication-efficient learning of deep networks from decentralized data,” in *Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273–1282.
- [2] Peng Xiao et al., “Averaging is probably not the optimum way of aggregating parameters in federated learning,” *Entropy*, vol. 22, no. 3, pp. 314, 2020.
- [3] Luping Wang et al., “Cmfl: Mitigating communication overhead for federated learning,” in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2019, pp. 954–964.
- [4] Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papailiopoulos, and Yasaman Khazaeni, “Federated learning with matched averaging,” in *International Conference on Learning Representations*, 2020.
- [5] Peter Kairouz, Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Bhagoji, Keith Bonawitz, et al., “Advances and open problems in federated learning,” *arXiv preprint arXiv:1912.04977*, 2019.

Table 3. The final performance comparison in tinyImageNet dataset with ResNet18 in iid and no-iid conditions

tiny ImageNet	iid		non-iid		
	FedAvg	FedNS	FedAvg	FedAvg +lastFC	FedNS
top-1 accuracy	18.99	19.11	18.64	18.89	18.91
top-5 accuracy	41.14	41.39	40.80	40.86	41.05

- [6] Li Li et al., “A survey on federated learning,” in *2020 IEEE 16th International Conference on Control & Automation (ICCA)*. IEEE, 2020, pp. 791–796.
- [7] Tian Li et al., “Federated learning: Challenges, methods, and future directions,” *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 50–60, 2020.
- [8] Jakub Konečný et al., “Federated learning: Strategies for improving communication efficiency,” *arXiv preprint arXiv:1610.05492*, 2016.
- [9] Samuel Horvath, Chen-Yu Ho, Ludovit Horvath, Atal Narayan Sahu, Marco Canini, and Peter Richtarik, “Natural compression for distributed deep learning,” *arXiv preprint arXiv:1905.10988*, 2019.
- [10] Dan Alistarh, Demjan Grubic, Jerry Li, et al., “Qsgd: Communication-efficient sgd via gradient quantization and encoding,” in *Advances in Neural Information Processing Systems*, 2017, pp. 1709–1720.
- [11] Nader Bouacida, Jiahui Hou, Hui Zang, and Xin Liu, “Adaptive federated dropout: Improving communication efficiency and generalization for federated learning,” *arXiv preprint arXiv:2011.04050*, 2020.
- [12] Takayuki Nishio and Ryo Yonetani, “Client selection for federated learning with heterogeneous resources in mobile edge,” in *ICC 2019 IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–7.
- [13] Yue Zhao, Meng Li, et al., “Federated learning with non-iid data,” *arXiv preprint arXiv:1806.00582*, 2018.
- [14] Tiffany Tuor, Shiqiang Wang, Bong Ko, et al., “Overcoming noisy and irrelevant data in federated learning,” *arXiv preprint arXiv:2001.08300*, 2020.
- [15] Xin Yao, Tianchi Huang, Chenglei Wu, et al., “Towards faster and better federated learning: A feature fusion approach,” in *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2019, pp. 175–179.
- [16] Li Huang, Andrew Shea, Huining Qian, et al., “Patient clustering improves efficiency of federated machine learning to predict mortality and hospital stay time using distributed electronic medical records,” *Journal of Biomedical Informatics*, vol. 99, pp. 103291, 2019.
- [17] Jenny Hamer, Mehryar Mohri, and Ananda Suresh, “Fedboost: A communication-efficient algorithm for

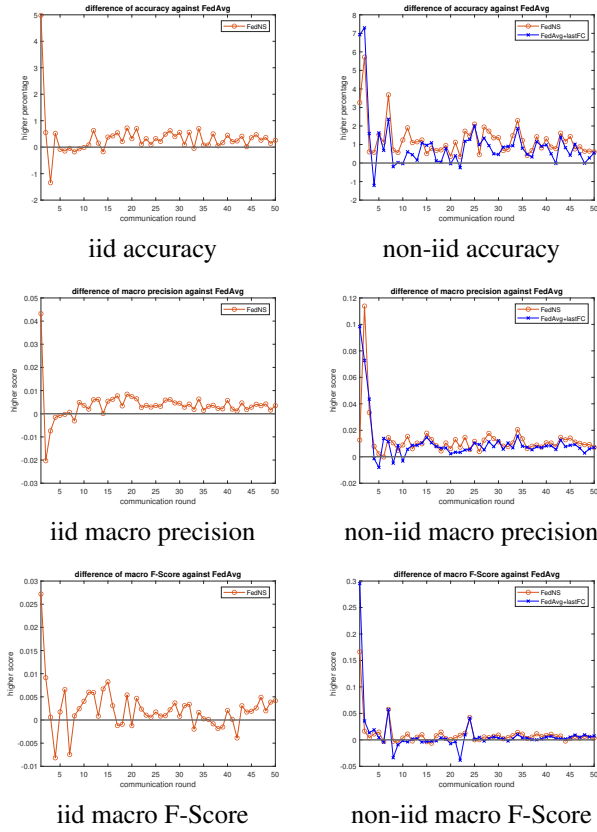


Fig. 1. FashionMNIST iid and non-iid daily comparison

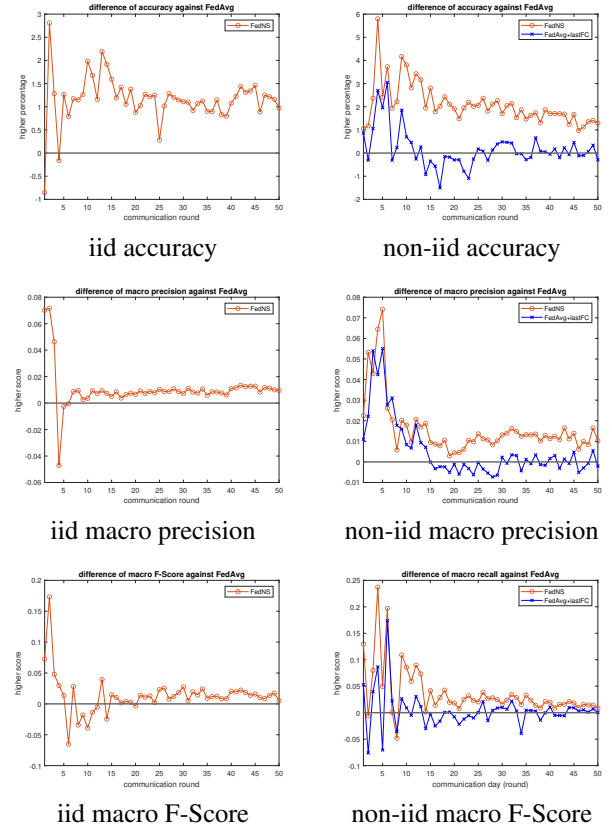


Fig. 2. CIFAR10 iid and non-iid daily comparison

federated learning,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 3973–3983.

- [18] Yann LeCun, Léon Bottou, et al., “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [19] Han Xiao, Kashif Rasul, et al., “Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms,” *arXiv preprint arXiv:1708.07747*, 2017.
- [20] Alex Krizhevsky et al., “Learning multiple layers of features from tiny images,” Tech. Rep., Citeseer, 2009.
- [21] “Tiny imagenet visual recognition challenge,” <https://tiny-imagenet.herokuapp.com/>.
- [22] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 2009, pp. 248–255.
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

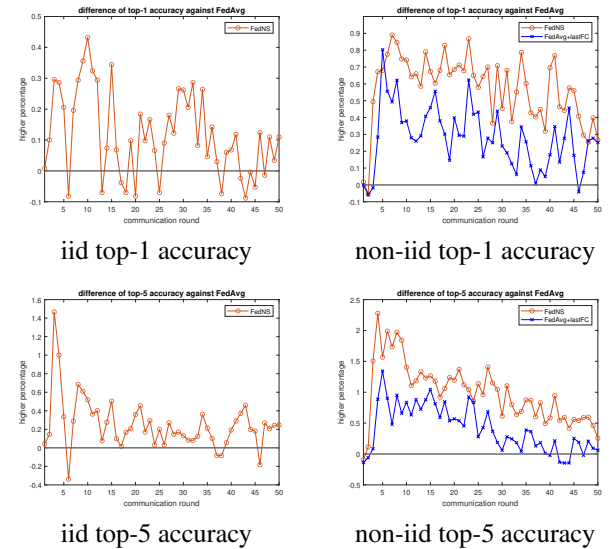


Fig. 3. tinyImageNet iid and non-iid daily comparison