

FLOP: Federated Learning on Medical Datasets using Partial Networks

Qian Yang^{*#}, Jianyi Zhang^{*}, Weituo Hao, Gregory P. Spell, Lawrence Carin
Duke University, USA

ABSTRACT

The outbreak of COVID-19 Disease due to the novel coronavirus has caused a shortage of medical resources. To aid and accelerate the diagnosis process, automatic diagnosis of COVID-19 via deep learning models has recently been explored by researchers across the world. While different data-driven deep learning models have been developed to mitigate the diagnosis of COVID-19, the data itself is still scarce due to patient privacy concerns. Federated Learning (FL) is a natural solution because it allows different organizations to cooperatively learn an effective deep learning model without sharing raw data. However, recent studies show that FL still lacks privacy protection and may cause data leakage. We investigate this challenging problem by proposing a simple yet effective algorithm, named Federated Learning on Medical Datasets using Partial Networks (FLOP), that shares only a partial model between the server and clients. Extensive experiments on benchmark data and real-world healthcare tasks show that our approach achieves comparable or better performance while reducing the privacy and security risks. Of particular interest, we conduct experiments on the COVID-19 dataset and find that our FLOP algorithm can allow different hospitals to collaboratively and effectively train a partially shared model without sharing local patients' data.

CCS CONCEPTS

• Computing methodologies → Supervised learning by classification.

KEYWORDS

Federated Learning; Disease Diagnosis

ACM Reference Format:

Qian Yang^{*#}, Jianyi Zhang^{*}, Weituo Hao, Gregory P. Spell, Lawrence Carin. 2021. FLOP: Federated Learning on Medical Datasets using Partial Networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21)*, August 14–18, 2021, Virtual Event, Singapore. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3447548.3467185>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '21, August 14–18, 2021, Virtual Event, Singapore

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8332-5/21/08...\$15.00

<https://doi.org/10.1145/3447548.3467185>

1 INTRODUCTION

Automatic disease diagnosis using machine learning methods holds immense promise, and innovations in this field may refine health care systems and improve medical practice worldwide. For example, human digestive system cancers — including esophageal, stomach and colorectal cancers — account for about 2.8 million new cases and 1.8 million deaths per year. Automatic detection, recognition, and assessment of pathological findings based on images from inside the gastrointestinal (GI) tract will assist doctors in identifying areas of concern and optimize use of scarce medical resources. Of great concern in 2020 and into 2021, the global COVID-19 (“the coronavirus”) pandemic has caused over 1.32 million deaths, with infections and deaths still increasing [1]. As communities and organizations across the world continue making efforts to control the pandemic, researchers seek to quicken COVID-19 early detection by automatically classifying computed tomography (CT) scan slices (images) of patients' chests [7, 20, 29, 30, 40].

However, there are two major challenges towards utilizing these medical images. One challenge is that, collectively, this data is distributed across a large number of devices or clients located in different hospitals. When relying on data-driven deep learning models to diagnose disease [11, 32], using only the local data isolated on a single device will not be sufficient to train an effective model. A second challenge is the necessity of using the data without compromising patients' privacy and security. The leaking of private data is not only a concern in public media, but also for the hospitals which must protect patients' privacy. To train deep learning models on such data while not compromising patients' privacy, federated learning [22] has become a promising solution by sharing a model between clients and a server, instead of sharing the data itself.

Recent improvements in federated learning include overcoming the statistical challenge in training machine learning models over distributed networks of devices [28, 39], improving security [4, 10], and personalization [6, 28]. The conventional federated learning framework is proved to prevent data leakage against a semi-honest server, if gradients aggregation is operated with SMC [4] or Homomorphic Encryption [3]. However, recent empirical results in [41] show that sharing a model may not fully protect privacy, and gradients exchange will cause *Deep Leakage* [9, 38, 41]. In [41], the authors showed that it is possible to obtain private training data from the publicly shared gradients, including pixel-wise images and token-wise sentences. One strategy to avoid deep leakage is by compressing the gradients. Furthermore, the authors in [9] empirically show that federated averaging is also susceptible to attacks, by successfully reconstructing training images from a convolutional neural network. To overcome these vulnerabilities of federated

^{*}Equal contribution

[#]Corresponding author. Email: laraqianyang@gmail.com.

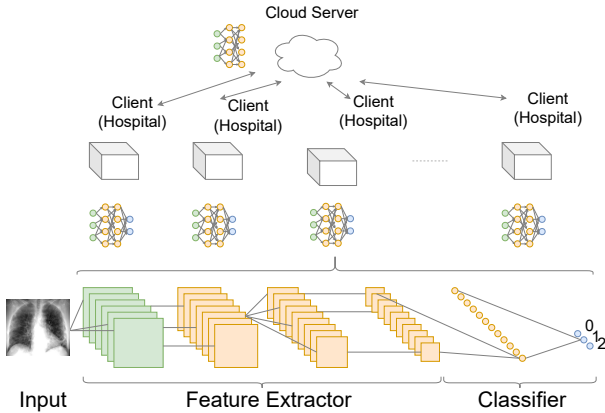


Figure 1: Overview of FLOP, allowing for collaboration among hospitals with small, local datasets to train better machine learning models without loss of privacy.

learning, this paper exploits a new model structure, and we present the attempt at sharing a partial model for federated learning on medical datasets, which is still an unexplored field.

In this paper, we propose a simple yet effective algorithm called Federated Learning on Medical Datasets using Partial Networks (FLOP). Specifically, instead of sharing an entire model between a server and clients in each round of training, clients share only a part of the model for federated averaging and keep the last several layers private. An overview of FLOP is shown in Figure 1, and our contributions are as follows:

- studying the effects of sharing a partial model in a federated learning framework on medical datasets;
- applying the FLOP algorithm to different model architectures (3-layer CNN, VGG11, CovidNet, ResNet50, MobileNet-v2, ResNetXt), and the presentation of extensive experiments on both benchmark data (Fashion-MNIST, CIFAR-10) and real-world medical data (COVIDx, Kvasir);
- showing empirically that FLOP allows for collaboration among clients (such as hospitals) with small, local datasets to train better machine learning models than the baseline algorithm FedAvg [22] without loss of privacy.

2 RELATED WORK AND PRELIMINARIES

With the emergence of tighter privacy regulations in Europe and around the world, researchers have started seeking solutions to train machine learning models from user data without compromising user privacy. Federated learning decentralizes conventional machine learning by removing the need to aggregate data into a single location or server, and has become the most popular solution to meet new data protection regulations [14, 22, 36]. Intuitively, the mechanism of federated learning is as follows: clients download the current model, train it on local data from the client, and then send model updates to the server. The server aggregates and averages the model updates from a set of clients to improve the shared model. All the training data remains on the client devices throughout the learning process.

More formally, we define a set of S data owners $\{C_1, \dots, C_S\}$, with the i -th owner holding a data matrix \mathcal{D}_i . Each row of the matrix \mathcal{D}_i denotes a data sample, and each column represents a particular feature. The data are partitioned by sample identifiers, such as user or device IDs. We denote the feature, label, and sample ID spaces as \mathcal{X} , \mathcal{Y} , and \mathcal{Z} , respectively. These constitute the complete training dataset $(\mathcal{Z}, \mathcal{X}, \mathcal{Y})$. Federated learning is a process whereby clients collaboratively train a model \mathcal{M} , while each \mathcal{D}_i is held locally by each data owner C_i .

Federated Learning can be classified into *horizontal* federated learning, *vertical* federated learning, and *federated transfer learning*. If the clients in federated learning share overlapping data features but differ in data samples, we denote it as horizontal federated learning [36]. The scenario in which clients share overlapping data samples but differ in data features is known as vertical federated learning. Federated transfer learning is the case in which there is no overlap in both data samples or features. For example, when two hospitals serve two different regions, the data samples associated with a specific disease are likely different but with similar feature spaces, as the disease is the same. Therefore, the two hospitals can collaborate in designing better machine learning models through horizontal federated learning, without loss of privacy.

The Federated Learning framework has been applied to many healthcare tasks, such as predicting heart-related hospitalizations [5] and understanding the genetic underpinnings of brain diseases [27]. Recent works [19, 21] focusing on federated learning for COVID-19 typically rely on sharing the full model between clients. Moreover, they do not differentiate between IID and Non-IID data distributions. However, [41] indicates that sharing a full model will cause *Deep Leakage* [9, 38, 41]. To address these shortcomings, this paper investigates a new model framework and presents the attempt to share a partial model for federated learning on medical datasets. We also analyze both IID and Non-IID data distribution cases in our experiments.

The model architectures in this paper are based on Convolutional Neural Networks (CNNs), which have achieved great empirical success in computer vision [12, 18], natural language processing [15, 37] and speech recognition [2, 23]. Although there are many variations of the CNN architecture, a CNN for image classification tasks is typically composed of two basic components: a feature extractor and classifier. The feature extractor includes several convolutional layers followed by max-pooling and an activation function, while the classifier usually consists of fully connected layers. Motivated by this observation, we note a natural way to incorporate a split CNN model into a federated learning architecture: a *shared* feature extractor with general feature domain information and *private* classifier with private label and task information.

3 METHOD

Federated learning addresses data collection/aggregation concerns by communicating model updates only. In this paper, we strengthen the data protection by splitting a model into two parts, choosing a natural split for CNN architectures: a shared part with general feature domain information, and an unshared part with user-specific task information. We summarize the proposed FLOP method in Algorithm 1. Let \mathcal{M} denote the full model, which is partitioned into

Algorithm 1 FLOP Algorithm

Model Training(\mathcal{M}): // Run on user u

procedure

Receive \mathcal{M}_s from the server and let $\mathcal{M}_s^u = \mathcal{M}_s$
Train the model $\mathcal{M}^u = [\mathcal{M}_s^u, \mathcal{M}_p^u]$ on local training set
Return $\Delta\mathcal{M}_s^u$ to the server

end procedure

Model Update

procedure

Initialize \mathcal{M}_s
for each episode $t = 1, 2, \dots$ **do**
 Sample m users U_t
 for each user $u \in U_t$ **in parallel do**
 Send \mathcal{M}_s to user u
 end for
 Receive $\Delta\mathcal{M}_s^u$ from user u ;
 Update $\mathcal{M}_s = \mathcal{M}_s - \beta \frac{1}{|U_t|} \sum_{u \in U_t} \Delta\mathcal{M}_s^u$

end for**end procedure**

a shared part \mathcal{M}_s and private part \mathcal{M}_p . Altogether $\mathcal{M} = [\mathcal{M}_s, \mathcal{M}_p]$. For a particular local client u , we denote $\Delta\mathcal{M}_s^u$ as the update of \mathcal{M}_s^u . Again considering the horizontal federated learning paradigm for illustration: U clients with the same data structure collaboratively learn a machine learning model. The training process of our algorithm is as follows.

- Step 1: Clients receive \mathcal{M}_s from the server, train their own model $\mathcal{M}^u = [\mathcal{M}_s^u, \mathcal{M}_p^u]$ locally, and send the gradients of \mathcal{M}_s^u back to server;
- Step 2: The server performs secure aggregation on the received updates from the participating clients;
- Step 3: The server sends the aggregated results of \mathcal{M}_s back to clients;
- Step 4: Clients update their \mathcal{M}_s^u model with the results from the server.

The above steps are iterated until the loss function converges, concluding the training. The process and the algorithm are agnostic to any specific models, and all the clients will obtain the final shared model parameters.

Privacy is one of the key properties that federated learning aims to ensure. There are different types of privacy attacks in federated learning. Recent empirical results in [41] show that sharing a model may not fully protect the privacy and gradients exchange will cause *Deep Leakage* [9, 38, 41]. However, our FLOP framework addresses this vulnerability because it only shares a partial model. Furthermore, we can achieve guaranteed privacy by masking a selection of gradients with encryption [3], differential privacy [26], or secret sharing [4] techniques in step 1, which is out of the scope of this paper.

4 EXPERIMENTS

In this section, we report the results of our FLOP algorithm for different models. We build upon an open source federated learning framework¹ to implement our FLOP² in the PyTorch deep learning API. Our experiments are conducted on both real-world medical datasets (COVIDx and Kvasir) and benchmark datasets (Fashion-MNIST [33] and CIFAR-10 [17]). Specifically, each client has a sub-dataset derived from the original full dataset. We describe the details on generation of these sub-datasets for each client in Subsection 4.1, using the CIFAR-10 dataset as an example. On the two real-world medical datasets in Subsection 4.2, we use the CovidNet, ResNet50, MobileNet-v2, and ResNetXt model architectures. For the two benchmark datasets in Subsection 4.3, we verify the effectiveness of FLOP using the VGG-11 model architecture and a 3-layer CNN. The task in this paper is image classification, and the datasets across clients follow a non-IID distribution in this section. We discuss the non-IID results in Subsection 4.4 and also analyze the results when datasets are IID in Subsection 4.5.

4.1 Construction of the Non-IID Datasets

In the federated learning setting, the training data on a given client depend on the manner of device use by a particular user. For example, different hospitals may experience different COVID-19 caseloads, resulting in a varying proportion of COVID-19 cases. Any particular client’s **local** dataset will not be representative of the population distribution. Hence, the data distribution on each client is likely to be **Non-IID**. We describe the construction of non-IID datasets across client devices, using the CIFAR-10 dataset as an illustrative example. CIFAR-10 consists of 60,000 color images in 10 classes, with 6,000 images per class. The dataset is split into 50,000 training images and 10,000 test images. Supposing the number of clients for Federated Learning is 50, we distribute to each client 1,000 training images as its local dataset. The assignment of the 50,000 CIFAR-10 training images to each client under the non-IID setting is as follows:

- Step 1: The images are sorted such that all examples with the same category label are together. We note that CIFAR-10 has a uniform class distribution of 5,000 images for each of its 10 classes. Thus, after this step, the sorting yields 5,000 examples of the first class, 5,000 of the second class and so on for the remaining classes.
- Step 2: Set a number of “chunks”, and use these chunks to subdivide each class. For example, if we set the number of chunks for each class to be 25, then each chunk will have $5000/25 = 200$ images. Then the entire training dataset will have 250 chunks, and each client receives 5 chunks.
- Step 3: Randomly distribute the chunks uniformly to each client. In our running example, each client selects 5 chunks from the 250 chunks. The distribution scheme is as follows:
 - (i) The first client chooses the chunks from the first class with the probability of λ and from the rest classes with the probability of $1 - \lambda$. If we set $\lambda = 0.6$, the first client will select chunks randomly from the remaining classes with probability 0.4.

¹<https://github.com/AshwinRJ/Federated-Learning-PyTorch>

²<https://github.com/jianyizhang123/FLOP>

- (ii) Similarly, the second client chooses the chunks from the second class with 0.6 probability, and the tenth client chooses the chunks from the last class with 0.6 probability.
- (iii) After that, the eleventh client chooses chunks from the first class again with 0.6 probability and from the other classes with 0.4 probability.
- (iv) The distribution follows this scheme for clients until each has 5 chunks.
- (v) Once a particular class runs out of images, the current client will choose chunks from the remaining classes with a normalized probability (normalized from the original probability). For example, if there are no images in the first class, and the original probability is 0.6 from the first class and $(1 - 0.6)/9$ from the other nine classes, the current client will choose the chunks from each of the remaining nine classes with $1/9$.

Following the steps above, the clients will receive images with an uneven distribution of classes. Hence, across the clients, the data distribution in each client becomes Non-IID. We use this non-IID dataset distribution scheme for all datasets mentioned in the paper.

4.2 Experiments on Medical Datasets

4.2.1 Dataset.

COVIDx. The task of Covid-19 diagnosis is image classification with three classes: (i) *Normal* (No infection), (ii) *Pneumonia* (Non-COVID-19 infection, e.g., viral, bacterial, etc.), and (iii) *COVID-19* (COVID-19 viral infection). COVIDx [31] is the open-access benchmark dataset with the largest number of COVID-19 positive patient cases, and is the combination of five publicly available COVID-19 data repositories: (1) COVID-19 Image Data Collection [8], (2) COVID-19 Chest X-ray Dataset Initiative³, (3) Actualmed COVID-19 Chest X-ray Dataset⁴, (4) COVID-19 radiography dataset⁵, and (5) RSNA Pneumonia Detection Challenges dataset⁶. We use COVIDx as our training and test dataset. As these datasets are ever-updated during the ongoing pandemic, we specify that for our experiments, the dataset consists of 13,954 images for training and 1,579 for testing. The training dataset contains 7,966 *Normal*, 5,471 *Pneumonia* and 517 *COVID-19* images. The test dataset contains 885 *Normal*, 594 *Pneumonia*, and 100 *COVID-19* images.

Kvasir. The Kvasir dataset [24] concerns image classification for Gastrointestinal disease with eight classes. It includes images showing anatomical landmarks, pathological findings, or endoscopic procedures in the GI tract, which are collected using endoscopic equipment at Vestre Viken Health Trust (VV) in Norway. It consists of 8,000 images in 8 classes and 1,000 images for each class (6,000 for training and 2,000 for testing). The 8 classes show Anatomical Landmarks (Z-line, pylorus, cecum), Pathological Findings (esophagitis, polyps, ulcerative colitis), and Polyp Removal (“dyed and lifted polyp” and “dyed resection margins”) in the GI tract.

³<https://github.com/agchung/Figure1-COVID-chestxray-dataset>

⁴<https://github.com/agchung/Actualmed-COVID-chestxray-dataset>

⁵<https://www.kaggle.com/tawsfurrahman/covid19-radiography-database>

⁶<https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/data>

4.2.2 Model. For both medical datasets above, we apply our FLOP framework on models below to verify the framework’s efficacy.

COVID-Net [31] is a recently proposed deep convolutional neural network designed for the detection of COVID-19 cases from chest X-ray (CXR) images. To compress the network structure, it utilizes projection-expansion- projection-extension (PEPX) while preserving the performance to a large extent.

MobileNet-v2 [25] is a new mobile model which improves the state-of-the-art performance on several tasks, of which the architecture is based on an inverted residual structure. MobileNet-v2 uses lightweight convolutions to process features in the intermediate expansion layer.

ResNet50 [12] is a variant of the ResNet model, which utilizes the the Residual Block to improve the performance of very deep neural networks. It has been widely adopted in many computer vision tasks.

ResNeXt’s [34] topology is as the same as ResNet50. The difference is it uses a “split-transform-merge” strategy (branched paths within a single module) to improve the performance.

4.2.3 Implementations.

Local Testing. The classical federated learning framework uses FedAvg [22] to update the model \mathcal{M}_s shared between server and clients. The recent work in [21] further utilizes FedAvg to detect COVID-19. By contrast, FLOP avoids *Deep Leakage* [9, 38, 41] and protects patients’ privacy by only allowing the server access to a partial model \mathcal{M}_s^u from each client u . The server in [21] derives a full model after each training round and tests it on the test dataset, denoted as **global** testing. However, since the server in FLOP does not maintain a globally shared full model, we instead test FLOP with different $\mathcal{M}^u = [\mathcal{M}_s^u, \mathcal{M}_p^u]$ on each local client u , which is denoted as **local** testing.

With a sub-dataset distributed to each client, we further randomly split the client-local dataset into local training and test. In our experiments, we consider 5 clients in total. At each round, we randomly select 2 clients ($u_i, i = 1, 2$) from the 5 clients. The selected client u_i trains its model $\mathcal{M}^{u_i} = [\mathcal{M}_s^{u_i}, \mathcal{M}_p^{u_i}]$ on its own training dataset for 3 epochs, and then sends $\mathcal{M}_s^{u_i}$ to the sever. After the server aggregates and updates $\hat{\mathcal{M}}_s$, all the clients derive a new model $\hat{\mathcal{M}}^{u_i} = [\hat{\mathcal{M}}_s, \hat{\mathcal{M}}_p^{u_i}]$, where $\hat{\mathcal{M}}_s = \hat{\mathcal{M}}_s$. Then they test $\hat{\mathcal{M}}^{u_i}$ on their own local test datasets. We average the accuracy and loss over all the clients for the comparison.

We follow the previous work [21] to conduct our experiments in a pseudo-distributed setting on $1 \times$ Nvidia RTX 2080 Ti GPU. Clients use the Adam [16] optimizer with a learning rate $r = 2e - 5$ and weight decay $w = 1e - 7$. All other hyperparameters are as the same as found in [21].

4.2.4 Results on COVIDx.

Analysis. While we initially expect that there may be a tradeoff between model performance and privacy protection, we actually find that our FLOP framework outperforms classical FedAvg on four models (COVID-Net, MobileNet-v2, ResNet50, ResNeXt) by 0.5% ~ 2%. We partition each client local dataset into a **local** training

Table 1: Local testing accuracy on COVIDx. We expect a possible tradeoff between protecting privacy and model performance; however, we find FLOP improves local testing accuracy over FedAvg for four models by 0.5% ~ 2%.

Framework	COVID-Net	MobileNet-v2	ResNet50	ResNeXt
FedAvg	90.08±0.40	90.07±1.99	93.64±0.23	93.26±0.08
FLOP	92.10±0.36	91.16±1.76	94.54±0.21	93.72±0.32




Figure 2: Local testing loss on the COVIDx dataset, with testing as described in Section 4.2.3. Note that local testing loss for FLOP decreases more rapidly than that of the FedAvg framework. X-axis: epoch; y-axis: averaged local testing loss.




Figure 3: (a) Local testing loss for the four models altogether under FedAvg; (b) Local testing loss for the four models altogether under FLOP. X-axis: epoch; y-axis: averaged local testing loss.

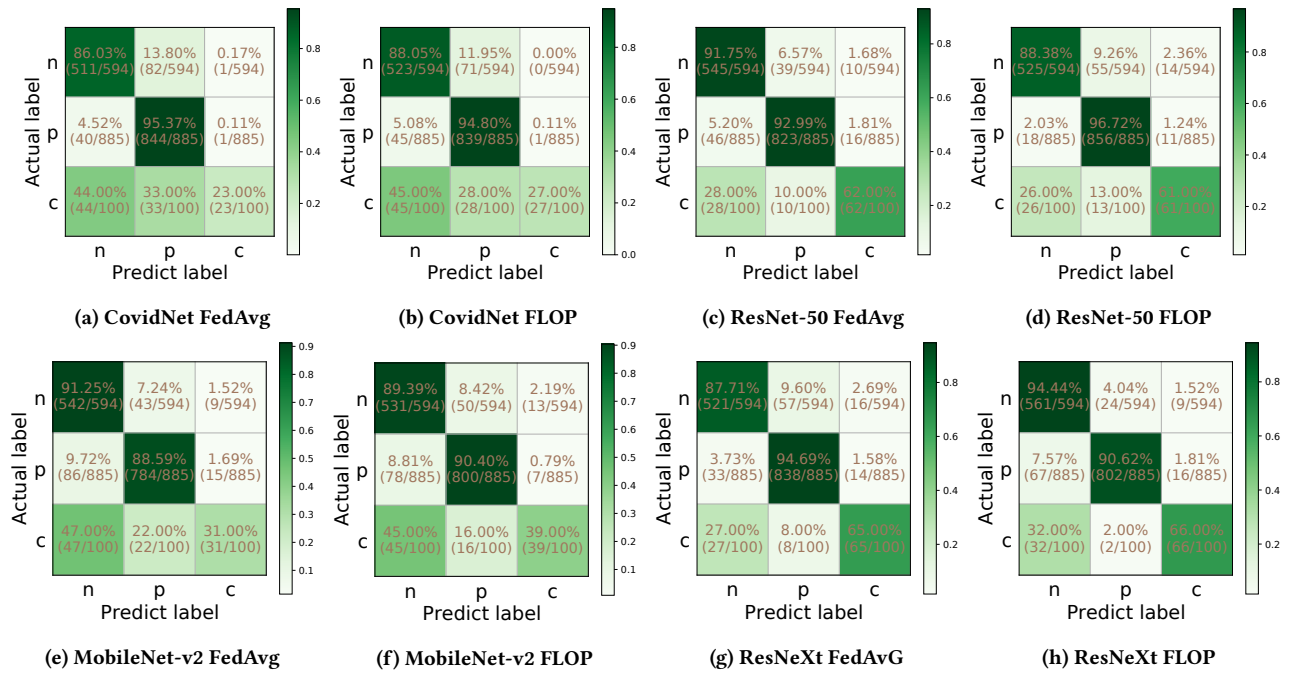


Figure 4: Confusion matrices on the COVIDx dataset. “n” denotes “normal”; “p” – “Pneumonia”; and “c” – “COVID-19”.

Table 2: Global testing accuracy on the COVIDx dataset

Framework	COVID-Net	MobileNet-v2	ResNet50	ResNeXt
FedAvg	87.61±0.38	86.98±3.15	90.86±0.23	90.26±0.12
FLOP	88.14±0.17	87.01±2.93	90.84±0.39	90.31±0.13

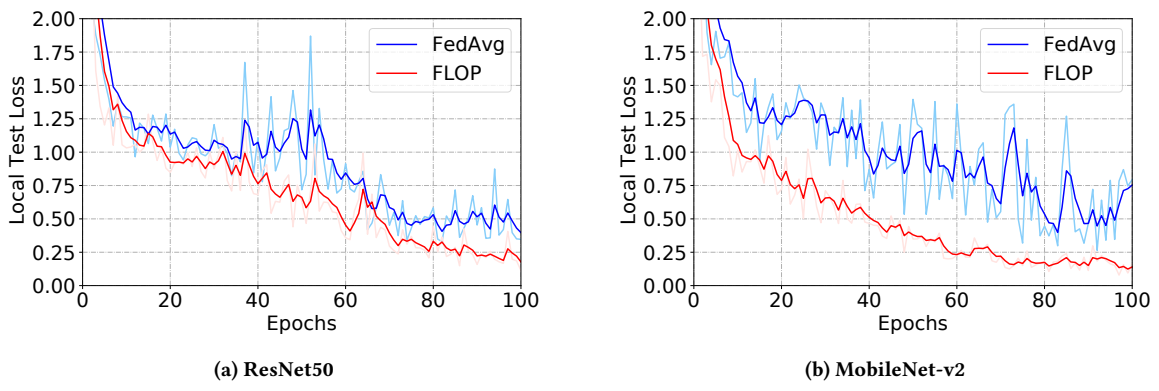


Figure 5: Local testing loss on the Kvasir dataset. Our FLOP framework outperforms compared to the classical FedAvg framework. X-axis: epoch; Y-axis: averaged local testing loss.

set (70%) and local test set (30%) and present local testing losses in Figure 2. We observe in Figure 2 that the local testing loss of our FLOP algorithm converges faster than the classical FedAvg algorithm and our FLOP obtains the better solutions than FedAvg.

We repeat the experiments with different random seeds and record the best local testing accuracy. Their average and Standard

Deviation are reported in Table 1. We find that across all models investigated, FLOP achieves higher local testing accuracy than FedAvg. In particular, FLOP on the CovidNet model improves the classical FedAvg by over 2%. Among the four models, ResNet50 achieves the best results for both the classical FedAvg framework and our FLOP framework. This is also evident in Figure 3, in which

we compare the local testing loss for the four models altogether under either FedAvg or FLOP. The performance of COVID-Net is similar to MobileNet-V2 under both frameworks. The local testing losses of ResNet50 and ResNeXt decrease faster than the lighter MobileNet-V2 and COVID-Net models. It is worth noting that ResNet50 achieves the best performance with respect to both metrics of local testing accuracy and local testing loss.

We additionally note from Figure 2 that the local test loss for FLOP decreases more stably than that of FedAvg as the models train. The curve for classical FedAvg, shown in blue, fluctuates dramatically while the curve for FLOP, shown in red, becomes stable. This improvement can be observed clearer in Figure 3. All curves in Figure 3b are more stable than the curves in 3a. We will further analyze the reason that our FLOP method achieves better results in Section 4.4.

To make the experimental results more comprehensive, we also explore the sensitivity of the models to each label. By presenting the confusion matrices for the models shown in Figure 4, we further demonstrate that the accuracy for each label in our FLOP is comparable to or even better than the accuracy in FedAvg. Specifically, the accuracy of Covid-19 label turns to be higher than the one in FedAvg, and our FLOP does not sacrifice the privacy.

Ablation Study. As mentioned earlier in Section 4.2.3, FLOP only shares the model \mathcal{M}_s^u between clients and the server. In the ablation study, we simulate FedAvg [22] by averaging the sum of $\tilde{\mathcal{M}}^{u_i}$ over all the clients and test it on the full test dataset, in order to obtain a similar and comparable **global** testing accuracy as tested by FedAvg. Table 2 shows the global testing accuracy on the COVIDx dataset. For all four models, the results of FLOP are comparable with those of classical FedAvg, outperforming FedAvg for three of the four models examined. With this performance and the increased privacy/security afforded by sharing only a partial model, we advocate for FLOP as a superior federated learning framework over FedAvg.

4.2.5 Results on Kvasir. As for the Kvasir dataset, we also split into **local** training sets (80%) and **local** test sets (20%) for each client. We train two models (ResNet50 and MobileNet-V2) to verify the effectiveness of the FLOP.

Table 3: Local testing accuracy for Kvasir

Framework	ResNet50	MobileNet-v2
FedAvg	88.85±2.39	91.08±1.37
FLOP	95.05±1.26	97.44±0.30

Table 4: Global testing accuracy for Kvasir

Framework	ResNet50	MobileNet-v2
FedAvg	82.46±2.13	84.15±1.04
FLOP	82.83±2.73	84.01±1.37

Local testing loss is shown in Figure 5. Again, we observe that the local testing loss of FLOP converges faster than the classical




Figure 6: Local testing loss for Fashion-MNIST. X-axis: epoch; Y-axis: averaged local testing loss.

FedAvg. Furthermore, we also see FLOP outperform FedAvg with respect to local testing accuracy, shown in Table 3. Our FLOP increases the accuracy by ~ 6% for both models.

An interesting result is that in this case, MobileNet-V2 achieves better results than ResNet50, despite MobileNet-V2 being lighter neural network than ResNet50. In the case of the Kvasir dataset, our experiments suggest using MobileNet-V2 over ResNet50 to achieve the best results for Federated learning.

Ablation Study. Similar to the experiments on COVIDx, we also conduct an ablation study on Kvasir. Shown in Table 4, we present the **global** test accuracy for these two models. Our framework is still comparable to FedAvg while not sacrificing client privacy.

4.3 Experiments on Benchmark Datasets

To further verify the effectiveness of FLOP, we conduct experiments on two benchmark datasets that are also publicly available: Fashion-MNIST and CIFAR-10.

Fashion-MNIST. Fashion-MNIST consists of a training set of 60,000 images and a test set of 10,000 images. Each example is a 28×28 grayscale image, associated with a label from 10 classes. We use a 3-layer CNN model: two convolutional layers followed by one linear layer.

Following the similar settings of the experiments on COVIDx and Kvasir, the simulation on Fashion-MNIST is for 100 clients and terminates after 50 rounds. For every round, we randomly select 15 clients from the 100 clients and perform 10 local epochs for training on each client. We optimize using stochastic gradient descent (SGD) with batch size 60. After each round, we record the local testing accuracy and the local testing loss on each client. Then we average them over all the clients and show the results in Figure 6.

CIFAR-10. CIFAR-10 consists of 60,000 32×32 colour images in 10 classes, with 50,000 training images and 10,000 test images. The simulation on CIFAR-10 is for 50 clients in total and terminates after 100 rounds. For every round, we randomly select 20 clients from the 50 clients and perform 5 local epochs on each client. The model is VGG-11, and the last linear layer is not shared. The optimizer is stochastic gradient descent (SGD) with batch size 100. After each round, we record the local testing accuracy and the local testing loss on each client. Then we average them over all the clients and show the results in Figure 7.

Table 5: Local testing accuracy on COVIDx when the datasets of the clients are IID. We expect a possible tradeoff between protecting privacy and model performance; however, we find FLOP improves local testing accuracy over FedAvg for all four models tested.

Framework	COVID-Net	MobileNet-v2	ResNet50	ResNeXt
FedAvg	91.33±0.10	91.87±2.98	94.01±0.19	94.03±0.71
FLOP	91.52±0.13	91.90±0.63	94.51±0.27	94.15±1.06

Table 6: Global testing accuracy on COVIDx when the datasets of the clients are IID.

Framework	COVID-Net	MobileNet-v2	ResNet50	ResNeXt
FedAvg	88.84±0.34	89.11±3.95	91.49±0.29	90.98±0.59
FLOP	88.51±0.26	89.65±3.63	91.45±0.25	91.20±0.81




Figure 7: Local testing loss for CIFAR-10. X-axis: epoch; Y-axis: averaged local testing loss.

Table 7: Local testing accuracy for Kvasir when the datasets of the clients are IID.

Framework	ResNet50	MobileNet-v2
FedAvg	97.82±0.87	98.38±0.31
FLOP	98±0.46	98.45±0.17

Table 8: Global testing accuracy for Kvasir when the datasets of the clients are IID. Note that while FedAvg slightly outperforms FLOP in this case, it is likely at the expense of model privacy

Framework	ResNet50	MobileNet-v2
FedAvg	97.24±1.07	98.28±0.19
FLOP	96.45±2.64	98.04±0.29

4.4 Discussion

In this subsection, we discuss why our FLOP algorithm leads to the improved results when compared with FedAvg.

The component of the networks we do not share is \mathcal{M}_p^u , which in our experiments is the classifier. It is typically composed of several linear layers. The classifier whose output is the predicted labels, is the most related component to the data distribution of the ground-truth labels. Hence, we believe this component carries much more information of clients’ datasets. This motivates keeping the classifier separate from server in the FLOP algorithm.

Because this component will be affected by clients’ data, it is highly personalized for each client. In the classical FedAvg framework, the clients share the full model and receive a new one after each round from the sever. The new classifier is less personalized for the clients’ data than the former one. Thus, the local testing loss is less stable than ours, and we similarly see that our framework can achieve superior local testing accuracy.

4.5 Results of the IID cases

In this subsection, we report the training accuracy for the IID cases on medical datasets. The hyperparameters and other experimental settings of the IID cases are the same as the Non-IID cases.

COVIDx. While we initially expect that there may be a tradeoff between model performance and privacy protection, we actually find that our FLOP framework outperforms classical FedAvg on four models (COVID-Net, MobileNet-v2, ResNet50, ResNeXt), though the improvement is less for the IID case than for the non-IID case. The local testing accuracy on the COVIDx dataset for the IID case are shown in Table 5. We believe that the less dramatic improvement than for the non-IID case is in line with our discussion in Section 4.4. Specifically, in the IID case, the non-shared classifier for each client is slightly less personalized, since the data distributions on the clients are more similar than in the non-IID case. Again, we see that the performance of the ResNet50 model is the best among the four models tested.

As for the global testing accuracy in Table 6, the accuracy of our FLOP method is again comparable to the accuracy of FedAvg, again strengthening our argument that FLOP is a strong method of preserving privacy and training effective models in a federated setting.

Kvasir. With respect to local testing accuracy, shown in Table 7, FLOP again outperforms FedAvg on the Kvasir dataset in the IID case, and again MobileNet-v2 outperforms the ResNet50, which is consistent across the results for non-IID and IID cases.

With respect to global testing accuracy, shown in Table 8, our FLOP method performs competitively, but does not outperform, the FedAvg scheme. We expect that this slight underperformance is in tradeoff to the additional privacy afforded to FLOP by sharing only a partial model between server and clients.

5 CONCLUSION

We have proposed a Federated Learning method in which only a partial model is shared between clients and server – FLOP – and demonstrated its use particularly for applications with medical data. Our proposed algorithm reduces privacy and security risks by sequestering client data on their local devices. Experimental results on both real-world medical datasets and benchmark datasets demonstrate the advantages of our algorithm. In future work, we intend to accelerate the training of the models following the techniques in [13, 35] and apply our algorithm to other tasks. Overall, we believe that our research makes an important step for improving the performance of deep learning models on data-scarce healthcare tasks, as our algorithm allows different hospitals to collaboratively train models without sharing local patients’ data.

REFERENCES

- [1] 2020. CORONAVIRUS. “<https://coronavirus.jhu.edu/map.html>”.
- [2] Ossama Abdel-Hamid, Abdel-rahman Mohamed, Hui Jiang, Li Deng, Gerald Penn, and Dong Yu. 2014. Convolutional neural networks for speech recognition. *IEEE/ACM Transactions on audio, speech, and language processing* 22, 10 (2014), 1533–1545.
- [3] Yoshinori Aono, Takuya Hayashi, Lihua Wang, Shihō Moriai, et al. 2017. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security* 13, 5 (2017), 1333–1345.
- [4] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. 2017. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*. 1175–1191.
- [5] Theodora S Brisimi, Ruidi Chen, Theofanie Mela, Alex Olshevsky, Ioannis Ch Paschalidis, and Wei Shi. 2018. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics* 112 (2018), 59–67.
- [6] Fei Chen, Zhenhua Dong, Zhenguo Li, and Xiuqiang He. 2018. Federated meta-learning for recommendation. *arXiv preprint arXiv:1802.07876* (2018).
- [7] Nanshan Chen, Min Zhou, Xuan Dong, Jieming Qu, Fengyun Gong, Yang Han, Yang Qiu, Jingli Wang, Ying Liu, Yuan Wei, et al. 2020. Epidemiological and clinical characteristics of 99 cases of 2019 novel coronavirus pneumonia in Wuhan, China: a descriptive study. *The Lancet* 395, 10223 (2020), 507–513.
- [8] Joseph Paul Cohen, Paul Morrison, and Lan Dao. 2020. COVID-19 image data collection. *arXiv 2003.11597* (2020). <https://github.com/iee8023/covid-chestxray-dataset>
- [9] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. 2020. Inverting Gradients—How easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053* (2020).
- [10] Robin C Geyer, Tassilo Klein, and Moin Nabi. 2017. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557* (2017).
- [11] Varun Gulshan, Lily Peng, Marc Coram, Martin C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Tom Madams, Jorge Cudros, et al. 2016. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Jama* 316, 22 (2016), 2402–2410.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [13] Zhouyuan Huo, Bin Gu, Heng Huang, et al. 2018. Decoupled parallel backpropagation with convergence guarantee. In *International Conference on Machine Learning*. PMLR, 2098–2106.
- [14] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. 2019. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977* (2019).
- [15] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882* (2014).
- [16] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [17] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [18] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [19] Rajesh Kumar, Abdullah Aman Khan, Sinmin Zhang, WenYong Wang, Yousif Abuidris, Waqas Amin, and Jay Kumar. 2020. Blockchain-federated-learning and deep learning models for covid-19 detection using ct imaging. *arXiv preprint arXiv:2007.06537* (2020).
- [20] Qun Li, Xuhua Guan, Peng Wu, Xiaoye Wang, Lei Zhou, Yeqing Tong, Ruiqi Ren, Kathy SM Leung, Eric HY Lau, Jessica Y Wong, et al. 2020. Early transmission dynamics in Wuhan, China, of novel coronavirus-infected pneumonia. *New England Journal of Medicine* (2020).
- [21] Boyi Liu, Bingjie Yan, Yize Zhou, Yifan Yang, and Yixian Zhang. 2020. Experiments of federated learning for covid-19 chest x-ray images. *arXiv preprint arXiv:2007.05592* (2020).
- [22] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. 2017. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *Artificial Intelligence and Statistics*. 1273–1282.
- [23] Dimitri Palaz, Mathew Magimai Doss, and Ronan Collobert. 2015. Convolutional neural networks-based continuous speech recognition using raw speech signal. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 4295–4299.
- [24] Konstantin Pogorelov, Kristin Ranheim Randel, Carsten Griwodz, Sigrun Losada Eskeland, Thomas de Lange, Dag Johansen, Concetto Spampinato, Duc-Tien Dang-Nguyen, Mathias Lux, Peter Thelin Schmidt, Michael Riegler, and Pål Halvorsen. 2017. KVASIR: A Multi-Class Image Dataset for Computer Aided Gastrointestinal Disease Detection. In *Proceedings of the 8th ACM on Multimedia Systems Conference (MMSys’17)*. ACM, New York, NY, USA, 164–169. <https://doi.org/10.1145/3083187.3083212>
- [25] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4510–4520.
- [26] Reza Shokri and Vitaly Shmatikov. 2015. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*. 1310–1321.
- [27] Santiago Silva, Boris A Gutman, Eduardo Romero, Paul M Thompson, Andre Altmann, and Marco Lorenzi. 2019. Federated learning in distributed medical databases: Meta-analysis of large-scale subcortical brain data. In *2019 IEEE 16th international symposium on biomedical imaging (ISBI 2019)*. IEEE, 270–274.
- [28] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. 2017. Federated multi-task learning. In *Advances in Neural Information Processing Systems*. 4424–4434.
- [29] Abdul Waheed, Muskan Goyal, Deepak Gupta, Ashish Khanna, Fadi Al-Turjman, and Plácido Rogério Pinheiro. 2020. Covidgan: Data augmentation using auxiliary classifier gan for improved covid-19 detection. *IEEE Access* 8 (2020), 91916–91923.
- [30] Dawei Wang, Bo Hu, Chang Hu, Fangfang Zhu, Xing Liu, Jing Zhang, Binbin Wang, Hui Xiang, Zhenshun Cheng, Yong Xiong, et al. 2020. Clinical characteristics of 138 hospitalized patients with 2019 novel coronavirus-infected pneumonia in Wuhan, China. *Jama* 323, 11 (2020), 1061–1069.
- [31] Linda Wang and Alexander Wong. 2020. COVID-Net: A Tailored Deep Convolutional Neural Network Design for Detection of COVID-19 Cases from Chest X-Ray Images. *arXiv preprint arXiv:2003.09871* (2020).
- [32] Wenlin Wang, Hongteng Xu, Zhe Gan, Bai Li, Guoyin Wang, Liqun Chen, Qian Yang, Wenqi Wang, and Lawrence Carin. 2020. Graph-driven generative models for heterogeneous multi-task learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 979–988.
- [33] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017).
- [34] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. 2017. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1492–1500.
- [35] Qian Yang, Zhouyuan Huo, Wenlin Wang, Heng Huang, and Lawrence Carin. 2019. Ouroboros: On Accelerating Training of Transformer-Based Language Models. *arXiv preprint arXiv:1909.06695* (2019).
- [36] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. 2019. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 2 (2019), 1–19.
- [37] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*. 649–657.
- [38] Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. 2020. iDLG: Improved Deep Leakage from Gradients. *arXiv preprint arXiv:2001.02610* (2020).
- [39] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandrara. 2018. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582* (2018).
- [40] Chuansheng Zheng, Xianbo Deng, Qing Fu, Qiang Zhou, Jiawei Feng, Hui Ma, Wenyu Liu, and Xinggang Wang. 2020. Deep learning-based detection for COVID-19 from chest CT using weak label. *medRxiv* (2020).
- [41] Ligeng Zhu, Zhijian Liu, and Song Han. 2019. Deep leakage from gradients. In *Advances in Neural Information Processing Systems*. 14747–14756.