

Fed-FSNet: Mitigating Non-I.I.D. Federated Learning via Fuzzy Synthesizing Network

Jingcai Guo, Song Guo, Jie Zhang, and Ziming Liu

Department of Computing, The Hong Kong Polytechnic University

{jingcai.guo, jieaa.zhang, ziming.liu}@connect.polyu.hk, song.guo@polyu.edu.hk

Abstract

Federated learning (FL) has emerged as a promising privacy-preserving distributed machine learning framework recently. It aims at collaboratively learning a shared global model by performing distributed training locally on edge devices and aggregating local models into a global one without centralized raw data sharing in the cloud server. However, due to the large local *data heterogeneities* (Non-I.I.D. data) across edge devices, the FL may easily obtain a global model that can produce more *shifted gradients* on local datasets, thereby degrading the model performance or even suffering from the non-convergence during training. In this paper, we propose a novel FL training framework, dubbed *Fed-FSNet*, using a properly designed Fuzzy Synthesizing Network (FSNet) to mitigate the Non-I.I.D. FL *at-the-source*. Concretely, we maintain an *edge-agnostic* hidden model in the cloud server to estimate a less-accurate while *direction-aware* inversion of the global model. The hidden model can then fuzzily synthesize several mimic I.I.D. data samples (sample features) conditioned on only the global model, which can be shared by edge devices to facilitate the FL training towards faster and better convergence. Moreover, since the synthesizing process involves neither access to the parameters/updates of local models nor analyzing individual local model outputs, our framework can still *ensure the privacy of FL*. Experimental results on several FL benchmarks demonstrate that our method can significantly mitigate the Non-I.I.D. issue and obtain better performance against other representative methods.

1 Introduction

Machine deep learning has made tremendous success in the past few years across multiple real-world applications [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30]. With the advances in computing capability of edge devices (e.g., mobile phones), various machine learning (ML) applications [23, 31, 32, 5, 33, 2, 34, 35, 13, 12] are usually deployed locally on user-sides to directly perform training or inference, rather than relying on a centralized cloud server-based training framework for better privacy preservation [36, 37, 38, 39, 19, 40, 29]. More specifically, users can either train an ML model individually upon their local data or download a pre-trained learned from external data from the cloud server. However, under such a paradigm, the former scenario may usually be quite stuck with local devices because the models cannot generalize well to massive new data, while the latter may result in a too mediocre model for each user by using only external data resources. For example, as one of the most ubiquitous edge devices in our daily life, the mobile phone can constantly generate privacy-sensitive data across multiple users, such as images, texts, voices, and so on. These data can represent sufficient diversity of various user patterns and thus are of great value to be utilized to train a better ML model for all participated users. However, due to privacy concerns, it is not applicable to upload and share these data in the cloud server for centralized model training, where the limitation is considered as the *isolated-data-island* that hinders the further improvement of the model [41]. To facilitate the usage of such privacy-sensitive data and

training a more generalized model, federated learning (FL) has emerged as a promising paradigm of privacy-preserving distributed ML framework [42, 43, 44, 45, 46, 47, 48, 49]. FL enables each device to train a model locally based on its own data, and communicates with other local models in the cloud server for aggregating a more generalized global model, without collecting and sharing any privacy-sensitive information from users.

However, unlike the standard cloud server-based distributed ML training framework, the FL is fundamentally posed with some challenges that limit its further development. Among them, due to the large heterogeneities across multiple devices, their data samples are usually not independent and identically distributed (Non-I.I.D.) [50, 51]. In such a scenario, the conventional FL is not able to efficiently train a shared global model and may thus degrade the performance or even encounter non-convergence during training. To mitigate the Non-I.I.D. FL, existing methods mostly follow two mainstreams. The first usually considers properly sampling several devices (or local training data) to participate in each round of FL to counterbalance the bias introduced by non-I.I.D. data [42, 52, 47]. Such treatment is based on the assumption that there may exist (partially) I.I.D. cases across devices. The second considers exploring a balanced aggregation strategy for each local model (or model updates) to offset the Non-I.I.D. issue [53, 54, 48, 55], of which each device is assumed to contribute differently to the overall model. Despite the progress made, these methods can only make use of limited statistical observations associated with local devices, e.g., local model parameters/updates or sketchy data distributions, which are not sufficient to address the data Non-I.I.D. issue *at-the-source*. Worse still, unlike the secure aggregation-based methods [56, 57] where the cloud server can only observe the *final sum* of model parameters/updates, the above-mentioned methods assign the priority of observing the parameters/updates of *each local model* to the cloud server, thereby making the FL training not fully privacy-preserved if the cloud server is adversarial-oriented and may potentially reconstruct raw data of local devices.

In this paper, we consider a novel strategy to synthesize mimic I.I.D. data samples (sample features) in a global view and offload them to participating local devices to directly mitigate the Non-I.I.D. FL *at-the-source*. Specifically, we propose a well-designed Fuzzy Synthesizing Network (*FSNet*) that maintains an *edge-agnostic* hidden model in the cloud server, by using an encoder-decoder architecture associated with the global model to estimate an inversion of it. This inversion is typically achieved by sampling Gaussian mixture with variational mixture coefficients, to perform encoding and reconstruction process by the global and the hidden model, respectively. Notably, since the sampling is not equivalent to the real data distribution, the obtained hidden model is not really accurate. However, it is expected that by inverting the parameters of the global model, the hidden model can more or less capture some statistical characteristics of the global data distribution, and can thus provide direction-aware information towards I.I.D. cases. This hidden model can then fuzzily synthesize several mimic global I.I.D. data samples (sample features) with sampled softmax distribution of classes (i.e., decoded by hidden model), which will be further utilized to regularize the training process of each local model by a well-designed simple yet efficient bias-rectified objective. To simplify, we denote our model as *Fed-FSNet*. Moreover, it should also be noted that since our proposed method is edge-agnostic and only synthesizes fuzzy mimic global I.I.D. data *conditioned on only the global model*, we thus can still ensure the FL privacy [56, 57].

In summary, the contributions of this paper are three-fold:

- We propose a novel FL training framework, dubbed *Fed-FSNet*, involving a well-designed *FSNet* to mitigate the Non-I.I.D. FL at-the-source. Our method can greatly narrow the gap between fully I.I.D. FL models.
- The proposed *FSNet* makes no inroads on the parameters/updates of local models nor analyzing individual local model outputs, thus can retain and ensure the privacy of the training process similar to conventional secure aggregation-based FL models.
- Experimental results on several FL benchmarks verified the effectiveness of our model on the superior mitigation of Non-I.I.D. FL.

2 Related Work

2.1 Non-I.I.D. Data in FL

Since different users usually have different habits and usage environments, the data samples generated on their devices are usually not independent and identically distributed (Non-I.I.D.) in real-world applications. Typically, during the FL training process, the Non-I.I.D. data means the differences between the data distributions $\mathcal{P}_i(x)$ and $\mathcal{P}_j(x)$ for any different devices $\mathcal{C}_i, \mathcal{C}_j \sim \mathcal{Q}$, where \mathcal{Q} denotes the entire participating devices. Ideally, an I.I.D. (or at least quasi-I.I.D.) assumption of the training data is of crucial importance to ensure that the stochastic gradient is an unbiased estimate of the full gradients and thus makes the entire training convergence. In practice, it is however not realistic to assume that the local data on each device is always I.I.D. To mitigate the Non-I.I.D. issue in FL, *MOCHA* [43] proposed a multi-task learning framework to handle the statistical Non-I.I.D. challenge which is robust to practical systems issues. But this solution differs significantly from the previous work on FL. *FedAvg* [42] has demonstrated that the iterative model averaging can work with certain non-I.I.D. data. *FedAtt* [54] proposed an attention mechanism in FL considering the contribution of each local model for aggregation and constructed a more generalized mobile keyboard suggestion model. More recently, *AFL* [52] proposed a simple and intuitive active device sampling scheme for FL training, with a probability conditioned on the current model and the data on the devices to offset the Non-I.I.D. Similarly, *FAVOR* [47] proposed a deep reinforcement learning-based experience-driven control framework that can intelligently sample devices to participate in each FL round to counterbalance the non-I.I.D. bias. Recently, knowledge distillation has been utilized to refine the server model aggregation using aggregated knowledge from heterogeneous users and obtained promising results, the representative methods include *FedDF* [48] and *FedGEN* [55].

2.2 Inference Attacks as a Preliminary

Recently, some generative models, e.g., auto-encoder (AE) [58, 28, 59, 8] and generative adversarial network (GAN) [60, 61, 62], and some optimization approaches are utilized for the inference attack on networks. With the ability to synthesize data samples of the target distribution [58, 60, 26], the generative models have proven to be efficient for inference attacks on ML models and thus can be used to recover privacy-sensitive data for certain purposes. For example, Hitaj et al. [63] proposed an inference attack model against collaborative deep learning that allows the attacker to reconstruct sensitive information on the victim’s device. *mGAN-AI* [64] incorporated a novel multi-task discrimination process on the reality, category, and identity of the target device for inferring user-level class representatives. Differently, *DLG* [65] applied an optimization method to match the gradients of the model on the malicious attacker and the gradients of the model trained on the target device to recover privacy-sensitive information.

Motivated by inference attacks, we can also have the idea to make use of the related tools, especially the generative models to synthesize mimic I.I.D. data for participated devices, and mitigates the Non-I.I.D. FL at-the-source. Specifically, it is noticed that the above inference attack models usually require *device-level information*, e.g., individual local model parameters/updates or outputs, to recover device-level data distribution, which however, naturally violates the FL privacy. Different from the above inference attack methods, we consider maintaining an edge-agnostic hidden model by training an encoder-decoder architecture conditioned on only the global model, to fuzzily synthesize mimic I.I.D. data in a global view. Our intuition is that although the synthesized mimic I.I.D. data is not fully accurate due to the absence of local model parameters/updates, it can still retain certain statistical characteristics of the overall global distribution that can provide direction-aware information towards I.I.D. cases. Such information is further utilized to regularize the FL training towards faster and better convergence.

3 Methodology

We first give the problem definition of the FL training. Next, we introduce our proposed method and formulation in detail. More specifically, our motivation is that since the Non-I.I.D. brings about the biased estimation of each local model, we thus can make use of the mimic I.I.D. data (synthesized by the proposed *FSNet*) to properly regularize the training process of each local model and offsets such bias at-the-source.

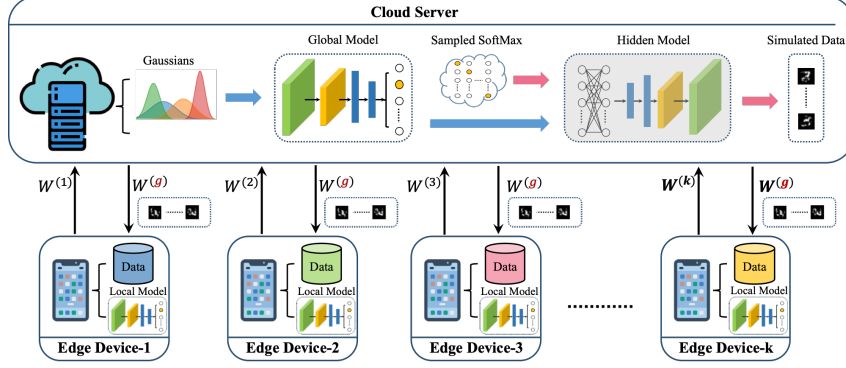


Figure 1: Illustration of the proposed FL training framework: an edge-agnostic hidden model is maintained in the cloud server to synthesize mimic I.I.D. data samples (sample features) conditioned on only the global model. These synthesized data will be offloaded along with the shared global model/parameters to each local device and further regularize the FL training.

3.1 Problem Definition

We start by formalizing the FL training and then introduce our proposed method. Given K edge devices (e.g., mobile phones) $\{C_1, C_2, \dots, C_K\} \sim \mathcal{Q}$, where \mathcal{Q} denotes the entire participated devices. Each C_i holds a private dataset belonging to $\{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K\}$, where the data distributions $\mathcal{P}_i(x)$ and $\mathcal{P}_j(x)$ for any different devices $C_i, C_j \sim \mathcal{Q}$ differ significantly. During the FL training, all participating devices are expected to collaboratively learn a shared global model. Without loss of generality, we consider the shared model to deal with a multi-class classification task, and thus each device trains a local model as:

$$\min_{\theta} \mathbb{E}_{(x,y) \in \mathcal{D}} [-y \log \mathcal{P}_{\theta}(\hat{y} | x)], \quad (1)$$

where $\mathcal{P}_{\theta}(\cdot)$ is the classifier with trainable parameter θ , and \hat{y} is the probability distribution (e.g., softmax outputs) that a data sample x is predicated on all classes compared with the ground-truth y . Considering K participated devices during the FL training, we can rewrite the objective (1) as:

$$\min_{\theta_1, \theta_2, \dots, \theta_K} \sum_{k=1}^K \mathbb{E}_{(x,y) \in \mathcal{D}_k} [-y \log \mathcal{P}_{\theta_k}(\hat{y} | x)], \quad (2)$$

where all K devices are trained collaboratively for a shared global model, i.e., $\mathcal{P}_{\theta^{(g)}}(\cdot)$. Specifically, $\mathcal{P}_{\theta^{(g)}}(\cdot)$ is obtained by iteratively performing local update and model aggregation in each communication round as:

- **Local Update:** Each participated device updates the local model on its private dataset using stochastic gradient descent (SGD) as $\theta_k \leftarrow \theta_k - \eta \nabla \mathcal{P}_{\theta_k}(\mathcal{B}_k)$, where \mathcal{B}_k denotes the batch samples for one local iteration.
- **Model Aggregation:** After the local update, the cloud server aggregates the updated local model parameters as $\theta^{(g)} = \sum_{k=1}^K \frac{N_k}{N} \theta_k$, where N_k and N denote the number of data samples in \mathcal{D}_k and the total number of all datasets, respectively. The updated global model is then distributed to all participating devices and repeats the steps.

3.2 Fuzzy Synthesizing Network (FSNet)

To synthesize mimic I.I.D. data, we introduce the fuzzy synthesizing network (*FSNet*). As demonstrated in Figure 1, our proposed FL training framework involving the *FSNet* can be decomposed into three major components: 1) the FL base training pipeline which is similar to *FedAvg* [42]; 2) the *FSNet* that maintains an edge-agnostic hidden model in the cloud server conditioned on only the global model, which can be used to synthesize mimic I.I.D. data samples (sample features); and 3) the joint training scheme of the proposed *Fed-FSNet* that offloads synthesized mimic I.I.D. samples (sample features) to each participated device and rectifies the bias caused by heterogeneous local data distribution, which can then mitigate the Non-I.I.D. FL at-the-source.

3.2.1 Training the *FSNet*

In our method, we use the encoder-decoder architecture to train the *FSNet*, of which the shared global model acts as the encoder that takes data samples as inputs and maps them to the label space (i.e., normalized probability distribution such as softmax outputs), and a newly designed hidden model acts as the decoder to reconstruct the data samples. The training process can be described as:

$$\min_{\omega} \mathbb{E}_{z \in G_s} \|\mathcal{H}_{\omega}(\mathcal{P}_{\theta^{(g)}}(z)) - z\|_2, \quad (3)$$

where $\mathcal{P}_{\theta^{(g)}}(\cdot)$ and $\mathcal{H}_{\omega}(\cdot)$ are the shared global model with fixed parameter $\theta^{(g)}$ and the hidden model with trainable parameter ω , respectively. $\|\cdot\|_2$ is the L_2 distance that measures the difference between the input sample and the reconstructed counterpart. Specifically, z denotes the samples sampled from a priori, e.g., Gaussian mixture with variational mixture coefficients $\sum_{k=1}^K \pi_k \mathcal{N}(\mu_k, \Sigma_k)$, where π_k denotes the variational mixture coefficients.

Notably, in our method, to better achieve the inversion of the shared global model, each local device is required to upload two simple scalars including its local mean μ_k and variance Σ_k to form the Gaussian mixture in the cloud server, which will not violate the FL privacy. Since there is no real data sample available in the cloud server for privacy preservation, we thus cannot have the observation on the real global distribution to train the hidden model $\mathcal{H}_{\omega}(\cdot)$ for an accurate reconstruction. However, we can expect that if the shared global model $\mathcal{P}_{\theta^{(g)}}(\cdot)$ is properly trained from all participated devices, it can certainly more or less retain some statistical characteristics of the global distribution in learned parameters $\theta^{(g)}$. Thus, in our method, when the fixed shared global model $\mathcal{P}_{\theta^{(g)}}(\cdot)$ and the hidden model $\mathcal{H}_{\omega}(\cdot)$ are trained in an *end-to-end* manner with variational samplings from $\sum_{k=1}^K \pi_k \mathcal{N}(\mu_k, \Sigma_k)$, we can still expect $\mathcal{H}_{\omega}(\cdot)$ to be a less accurate inversion of the shared model $\mathcal{P}_{\theta^{(g)}}(\cdot)$ that can provide direction-aware information, e.g., gradients that towards the global one, to facilitate the FL convergence.

Moreover, since the hidden model $\mathcal{H}_{\omega}(\cdot)$ is edge-agnostic and makes no inroads on the parameters/updates or outputs of each local model, it is thus can still ensure FL privacy. Hereby, we give a brief mathematical explanation of the privacy guarantee of our method. We define two scenarios S_1, S_2 as:

$$S_1 := \sum_{k=1}^N \alpha_k \mathcal{P}_{\theta_k}(\mathcal{D}_k), \quad (4)$$

$$S_2 := \{\mathcal{P}_{\theta_1}(\mathcal{D}_1), \mathcal{P}_{\theta_2}(\mathcal{D}_2), \dots, \mathcal{P}_{\theta_K}(\mathcal{D}_K)\}, \quad (5)$$

where S_1 can be presented as the form of the shared edge-agnostic global model $\mathcal{P}_{\theta^{(g)}}(\cdot)$ and S_2 denotes each local model separately. We assume $\mathcal{P}_{\theta_k} \in \mathbb{R}^d$ and we can rewrite S_1 as the following matrix expression as:

$$S_1 := \mathbf{M} \begin{bmatrix} \mathcal{P}_{\theta_1}(\mathcal{D}_1) \\ \mathcal{P}_{\theta_2}(\mathcal{D}_2) \\ \vdots \\ \mathcal{P}_{\theta_N}(\mathcal{D}_N) \end{bmatrix}, \quad (6)$$

where \mathbf{M} can be denoted as:

$$\mathbf{M} = \begin{bmatrix} \alpha_1 & & \alpha_2 & & \alpha_N & & \\ & \ddots & & \ddots & \cdots & & \ddots \\ & & \alpha_1 & & \alpha_2 & & \alpha_N \end{bmatrix}. \quad (7)$$

Note:

$$\begin{cases} \text{rank}(\mathbf{M}) = d \\ \dim(\ker(\mathbf{M})) = nd - d \end{cases}, \quad (8)$$

for $n > 1, nd - d > d$, it is obvious that $\mathbf{M}\vec{\mathcal{P}}_{\theta} = \sum_{k=1}^N \alpha_k \mathcal{P}_{\theta_k}(\mathcal{D}_k)$ has *non-unique* solution which means that an accurate estimation of S_2 is impossible from the observation on only S_1 . As a result, the proposed FL training framework can still preserve the privacy of each local device, which is similar to the secure aggregation-based FL models [56, 57].

3.2.2 Synthesizing Mimic I.I.D. Data

Next, we can use the decoder part of the *FSNet*, namely the hidden model $\mathcal{H}_\omega(\cdot)$, to synthesize mimic global I.I.D. data samples (sample features) and use them to regularize the training process of each local model. Specifically, we can know that the latent vector space of *FSNet* is the normalized probability distribution such as the softmax outputs that represent the classes. Essentially, the softmax distribution ideally acts as both the output and input of the shared global model $\mathcal{P}_{\theta(g)}(\cdot)$ and the hidden model $\mathcal{H}_\omega(\cdot)$, respectively, in an end-to-end manner. To synthesize mimic I.I.D. data samples (sample features), we can sample several representative softmax distributions from the latent vector space, and send them to the hidden model $\mathcal{H}_\omega(\cdot)$ to reconstruct them into the sample feature space. The synthesizing can be described as:

$$\mathcal{D}' = \mathcal{H}_\omega \left(\mathbf{S} = \begin{bmatrix} C_1 & \epsilon & \cdots & \epsilon \\ \epsilon & C_2 & \cdots & \epsilon \\ \vdots & \vdots & \ddots & \vdots \\ \epsilon & \epsilon & \cdots & C_n \end{bmatrix} \right), \quad (9)$$

where \mathcal{D}' is the synthesized dataset covering all global classes, and \mathbf{S} is the sampled softmax distributions to be used as the input of the hidden model $\mathcal{H}_\omega(\cdot)$. More specifically, each column of \mathbf{S} denotes a softmax output of a sample from target classes. For example, $\mathcal{H}_\omega((C_1, \epsilon, \dots, \epsilon)^\top)$ can condition the hidden model to synthesize a mimic data sample (sample features) for class “1” out of “ n ” classes, where $\sum (C_1, \epsilon, \dots, \epsilon)^\top = 1$ and $C_1 \gg \text{others}$. It should be noted that since $\mathcal{H}_\omega(\cdot)$ is conditioned on $\mathcal{P}_{\theta(g)}(\cdot)$ during training and conditioned on sampled \mathbf{S} during synthesizing, the obtained mimic I.I.D. data samples (sample features) are usually not fully accurate. Nevertheless, this fuzzily synthesized dataset \mathcal{D}' can still retain certain statistical characteristics of the overall global distribution that can provide direction-aware information, e.g., gradients that toward the global one, to facilitate the FL convergence. Such information can be used to regularize the training process and mitigates the Non-I.I.D. FL.

3.3 Joint Training on *Fed-FSNet*

With the mimic I.I.D. data samples $x' \in \mathcal{D}'$, we can then offload them (or extracted sample features) to each local device and train jointly with local data. Our intuition is that for each local dataset, i.e., $\{\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_K\}$, the local model is trained to classify its samples as accurate as possible and thus to retain and model the real local distribution. However, since the local data is Non-I.I.D. in each local device, the accurate modeling of each local dataset may eventually result in biased estimation of the global distribution and degrades the global model performance. To resolve the biased estimation, we can make use of the mimic global I.I.D. data, i.e., \mathcal{D}' , to force each local model to have the awareness of the global distribution as:

$$\min_{\theta} \underbrace{\mathbb{E}_{(x,y) \in \mathcal{D}} [-y \log \mathcal{P}_\theta(\hat{y} | x)]}_{(i)} + \underbrace{\beta \mathbb{E}_{x' \in \mathcal{D}'} [\mathcal{D}_{\mathcal{KL}}(\mathcal{U}(\tilde{y}) \| \mathcal{P}_\theta(\hat{y} | x'))]}_{(ii)}, \quad (10)$$

where term (i) corresponds to the classification loss that mainly guides the model to accurately classify samples of each local device. Term (ii) is the Non-I.I.D. regularizer that presents mimic global I.I.D. data, i.e., \mathcal{D}' , in each local model and forces them to shift the discrimination towards the global distribution. β is a hyper-parameter that controls the balance of term (ii). More specifically, $\mathcal{D}_{\mathcal{KL}}(\cdot)$ calculates the Kullback-Leibler (KL) divergence, a.k.a. relative entropy, between the predicted probability distribution of the mimic global I.I.D. data $\mathcal{P}_\theta(\hat{y} | x')$ and the uniform distribution $\mathcal{U}(\tilde{y})$. To minimize the KL divergence of term (ii), we can force the classification results, i.e., probability distributions of all mimic global I.I.D. data samples belonging to \mathcal{D}' , to approximate the uniform distributions. Thus, the obtained model will result in the flat probability distributions of them which means that the local model has endowed with the awareness of the global distribution while maintaining the discrimination of real local distribution at the same time. We call Eq. (10) the bias-rectified function for FL training, which resolves the biased estimation of local models and further makes the global aggregation more accurate in the cloud server. Similarly, considering K participated devices during the FL training, we can rewrite it as:

$$\min_{\theta_1, \theta_2, \dots, \theta_K} \mathbb{E}_{(x,y) \in \mathcal{D}_k} [-y \log \mathcal{P}_{\theta_k}(\hat{y} | x)] + \beta \mathbb{E}_{x' \in \mathcal{D}'} [\mathcal{D}_{\mathcal{KL}}(\mathcal{U}(\tilde{y}) \| \mathcal{P}_{\theta_k}(\hat{y} | x'))]. \quad (11)$$

Algorithm 1: Iterative Update

```
1 Initialize shared global model  $\mathcal{P}_{\theta^{(g)}}(\cdot)$  ;
2 Randomly select  $m$  devices, offload the shared global model weight  $\theta^{(g)}$  to each device;
3 Local training on selected  $m$  devices with Eq. (2), obtain  $\theta_1, \theta_2, \dots, \theta_m$ , and  $\mu_k / \Sigma_k$  pair ;
4 repeat
5   for  $round = 1, 2, \dots, T$  do
6     Upload  $\theta_1, \theta_2, \dots, \theta_m$ , and each  $\mu_k / \Sigma_k$  pair to the cloud server ;
7     Cloud server updates the shared global model weight as  $\theta^{(g)} = \sum_{k=1}^m \frac{N_k}{N} \theta_k$  ;
8     Cloud server trains the hidden model  $\mathcal{H}_\omega(\cdot)$  with Eq. (3) ;
9     Cloud server synthesizes the mimic I.I.D. data  $\mathcal{D}'$  with Eq. (9) ;
10    Randomly select  $m$  devices, offload the updated global model weight  $\theta^{(g)}$  and mimic
        I.I.D. data  $\mathcal{D}'$  (or extracted sample features) to each device ;
11    Local training on selected  $m$  devices with Eq. (11), obtain updated  $\theta_1, \theta_2, \dots, \theta_m$ , and
         $\mu_k / \Sigma_k$  pair.
12  end
13 until convergence;
```

The shared global model $\mathcal{P}_{\theta^{(g)}}(\cdot)$ is then obtained by iteratively performing local update and model aggregation in each communication round. To cooperatively optimize Eq. (11) and obtain the shared global model, we construct an iterative update algorithm in Algorithm 1.

4 Experiments

4.1 Task and Dataset

In our experiments, we implement the shared FL global model of the proposed *Fed-FSNet* as the image classification task, and we evaluate the performance on three benchmark datasets adapted in the “*edge device FL training environment*” including MNIST [66], Fashion-MNIST [67], and CIFAR-10 [68]. For each evaluation, we set a total of 100 edge devices with one cloud server for the FL training. In our experiments, we consider an extreme Non-I.I.D. data scenario for participated devices. To prepare the Non-I.I.D. data for each dataset, we first split the train-set into 200 sorted partitions and then disorder them randomly. Afterward, each edge device is offloaded with two random partitions of the above train-set. For each dataset, its test-set is used for a global test after each communication round. We designed three simple model architectures that fit the on-device training Scenario for MNIST [66], CIFAR-10 [68], and Fashion-MNIST [67], respectively.

4.2 Competitor and Metric

We compare the proposed *Fed-FSNet* with representative FL models including *FedAvg* [42], *AFL* [52], *FedDF* [48], and *FedGEN* [55]. Among them, the *FedAvg* [42] is the most standard FL training framework using averaging aggregation, which can be regarded as the baseline of FL. The other competitors adopt either data/devices sampling or improved aggregation strategies to mitigate the Non-I.I.D. issue. The selection criteria for our competitors are: 1) recent work: all of these competitors are released in the most recent years; 2) non-I.I.D. oriented models: all of these competitors address the data non-I.I.D. issue in federated learning; and 3) competitiveness: they clearly represent the state-of-the-art. The prediction accuracy is selected as the evaluation metric. Specifically, during each communication round, we record two predicted results including the averaged accuracy across all participated devices, i.e., denoted as Acc_{local} , and the global accuracy, i.e., denoted as Acc_{global} . It is noticed that, in the experiments, the “*upper-bound*” is set as the performance of *FedAvg* [42] (most standard FL training framework using averaging aggregation) in a *totally I.I.D. data setting* where each device shares full classes, and the “*lower-bound*” is set as its performance in the generated extreme Non-I.I.D. data setting.

4.3 Hyper-parameters

In our experiments, we use stochastic gradient descent (SGD) to optimize the training. Each evaluation is trained with only 50 global epochs, and each local model is updated with 10 local epochs, respectively. The mini-batch size is set to 60 with a learning rate of 0.01. The hyper-parameter β in Eqs. (10~11) is initially set to 1 and then decreased to one-tenth (1/10) in every 10 global rounds. During the training process, β controls the balance of the Non-I.I.D. regularizer, i.e., term (ii). The intuition is that at every beginning, each local device has very limited global information, thus a larger β can force local models to shift their discrimination towards the global one. As the training goes on, β trends to be decreased because more and more real global information is learned by the models. As to the hyper-parameter m in Algorithm 1 that denotes the number of participated edge devices during the FL training, we set a fraction $f=1.0$ and $f=0.1$ denoting two scenarios where all 100 devices and only 10 randomly selected devices of each epoch are trained during the FL training. Last, we force the sum of the variational mixture coefficients equals to 1 for the Gaussian mixture, i.e., $\sum_{k=1}^K \pi_k = 1$.

4.4 Evaluation

4.4.1 Prediction Results

The comparison of three benchmark datasets is demonstrated in Table 1 and Table 2. Specifically, Table 1 compares the proposed *Fed-FSNet* with the defined *upper-bound* and *lower-bound*, respectively. It can be observed from the results that our method surpasses the *lower-bound* by a significant margin in the extreme Non-I.I.D. data setting. The global accuracy and averaged local accuracy of the proposed *Fed-FSNet* achieve 88.2% / 85.3%, 61.7% / 60.4%, and 66.4% / 64.1% for MNIST [66], CIFAR-10 [68], and Fashion-MNIST [67] when $f=1.0$, respectively. Table 2 compares the proposed *Fed-FSNet* with selected representative FL methods in fully Non-I.I.D. ($f=0.1$) setting. We can observe that our method constantly obtains the best performance for global accuracy and averaged local accuracy, respectively.

In summary, the large margins with only 50 global epochs demonstrate the effectiveness of our method in two-fold: 1) our method can efficiently mitigate the Non-I.I.D. issue in FL training; and 2) our method is on-device friendly for the fewer training epochs and simple yet efficient architecture.

Method	MNIST		CIFAR-10		Fashion-MNIST	
	Acc_{global}	Acc_{local}	Acc_{global}	Acc_{local}	Acc_{global}	Acc_{local}
<i>Upper-bound</i> / I.I.D. ($f=0.1$)	97.6	97.3	72.6	71.9	76.4	74.2
<i>Upper-bound</i> / I.I.D. ($f=1.0$)	97.9	97.4	75.6	74.3	79.3	77.1
<i>Lower-bound</i> / Non-I.I.D. ($f=0.1$)	37.1	49.8	21.1	27.6	25.4	32.7
<i>Lower-bound</i> / Non-I.I.D. ($f=1.0$)	42.1	35.2	31.8	20.0	38.6	23.3
<i>Fed-FSNet (ours)</i> / Non-I.I.D. ($f=0.1$)	81.5	85.1	56.3	58.9	59.7	61.4
<i>Fed-FSNet (ours)</i> / Non-I.I.D. ($f=1.0$)	88.2	85.3	61.7	60.4	66.4	64.1

Table 1: Prediction results compared with “*upper-bound*” and “*lower-bound*”, respectively. The accuracy is reported in the averaged accuracy. Acc_{global} (%) denotes the averaged accuracy across all participated devices, and Acc_{local} (%) denotes the global accuracy. The *upper-bound* is set as the performance of *FedAvg* in a totally I.I.D. data setting, and the *lower-bound* is set as the performance of *FedAvg* in the generated extreme Non-I.I.D. data setting.

4.4.2 Training Analysis

We visualize the training process of the proposed *Fed-FSNet* with Non-I.I.D. settings on MNIST [66] in Figure 2. It can be observed from the results that the training process of our method (denoted as the red curve) is more stable and convergences much faster and better than the conventional Non-I.I.D. FL (denoted as green curve) by fully making use of the synthesized mimic I.I.D. data. It should be noted that, although our method still has a small margin from the totally I.I.D. FL (denoted as the blue curve), it greatly narrows the gap between non-I.I.D. and I.I.D., making it applicable for us to mitigate the Non-I.I.D. problem in FL and facilitates the implementation of practical applications.

Method	MNIST		CIFAR-10		Fashion-MNIST	
	Acc_{global}	Acc_{local}	Acc_{global}	Acc_{local}	Acc_{global}	Acc_{local}
<i>FedAvg</i> [42]	37.1	49.8	21.1	27.6	25.4	32.7
<i>AFL</i> [52]	54.2	59.7	33.8	42.2	37.9	49.8
<i>FedDF</i> [48]	74.5	83.3	49.7	52.5	51.7	57.3
<i>FedGEN</i> [55]	75.1	83.9	53.1	54.0	52.8	57.6
<i>Fed-FSNet</i> (ours)	81.5	85.1	56.3	58.9	59.7	61.4

Table 2: Prediction Results compared with representative models in Non-I.I.D. ($f=0.1$) setting. The accuracy is reported in the averaged accuracy. Acc_{global} (%) denotes the averaged accuracy across all participated devices, and Acc_{local} (%) denotes the global accuracy.

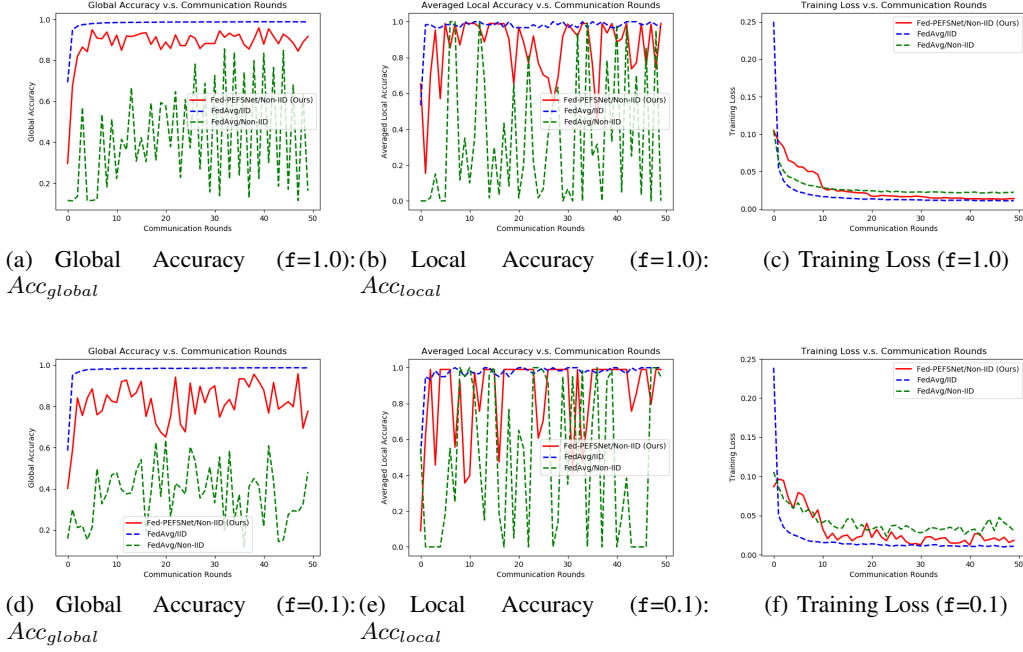


Figure 2: The training analysis of our proposed method. The red curve denotes our method with Non-I.I.D. settings, and the blue and green curves denote the upper bound (totally I.I.D.) and lower bound (Non-I.I.D.) of the baseline, respectively.

4.5 Discussion on Communication Overhead

In our experiments, since the mimic dataset \mathcal{D}' contains very few, i.e., only 60 synthesized samples (sample features) set in our experiment, it can result in a very limited impact on the communication overhead. Alternatively, if larger data, e.g., a CT volume, is utilized in real-world applications, the synthesized samples may increase exponentially. To solve this limitation, our solution is to deploy a pre-trained feature extractor in the cloud server and only transmit the extracted feature vector (far smaller) to each participating device, which can significantly reduce the communication overhead. Besides, since the proposed method requires fewer training epochs and simpler model architecture, thus it will be more suitable for on-device learning where the edge devices are usually equipped with less computing capabilities and fewer memories.

5 Conclusion

In this paper, we proposed a novel FL training framework involving a privacy ensured fuzzy synthesizing network (*Fed-FSNet*) to mitigate the Non-I.I.D. FL at-the-source. Our method narrows the gap

between the non-I.I.D. and I.I.D. FLs with a large margin. Experimental results on FL benchmarks verified the effectiveness of our model. In the future, we have two research directions to further improve the Non-I.I.D. FL. The first investigates more efficient and generalized methods to synthesize Non-I.I.D. data, and the second will focus on more privacy-preserved methods.

References

- [1] 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*, pages 248–255. Ieee, 2009.
- [2] Jingcai Guo and Song Guo. A novel perspective to zero-shot learning: Towards an alignment of manifold structures via semantic feature expansion. *IEEE Transactions on Multimedia*, 23:524–537, 2020.
- [3] Yisheng Lv, Yanjie Duan, Wenwen Kang, Zhengxi Li, and Fei-Yue Wang. Traffic flow prediction with big data: A deep learning approach. *IEEE Transactions on Intelligent Transportation Systems*, 16(2):865–873, 2014.
- [4] Thien Hai Nguyen and Kiyoaki Shirai. Topic modeling based sentiment analysis on social media for stock market prediction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1354–1364, 2015.
- [5] Jingcai Guo, Shiheng Ma, Jie Zhang, Qihua Zhou, and Song Guo. Dual-view attention networks for single image super-resolution. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2728–2736, 2020.
- [6] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. xdeepfm: Combining explicit and implicit feature interactions for recommender systems. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1754–1763, 2018.
- [7] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [8] Jingcai Guo and Song Guo. Adaptive adjustment with semantic feature space for zero-shot recognition. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3287–3291. IEEE, 2019.
- [9] Shuo Yang, Ping Luo, Chen-Change Loy, and Xiaoou Tang. From facial parts responses to face detection: A deep learning approach. In *Proceedings of the IEEE international conference on computer vision*, pages 3676–3684, 2015.
- [10] Nathan Shone, Tran Nguyen Ngoc, Vu Dinh Phai, and Qi Shi. A deep learning approach to network intrusion detection. *IEEE transactions on emerging topics in computational intelligence*, 2(1):41–50, 2018.
- [11] Jingcai Guo, Song Guo, Qihua Zhou, Ziming Liu, Xiaocheng Lu, and Fushuo Huo. Graph knows unknowns: Reformulate zero-shot learning as sample-level graph recognition. In *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI-23*, volume 2, 2023.
- [12] Ziming Liu, Song Guo, Jingcai Guo, Yuanyuan Xu, and Fushuo Huo. Towards unbiased multi-label zero-shot learning with pyramid and semantic attention. *arXiv preprint arXiv:2203.03483*, 2022.
- [13] Jingcai Guo, Song Guo, Shiheng Ma, Yuxia Sun, and Yuanyuan Xu. Conservative novelty synthesizing network for malware recognition in an open-set scenario. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [14] Xiaocheng Lu, Ziming Liu, Song Guo, and Jingcai Guo. Decomposed soft prompt guided fusion enhancing for compositional zero-shot learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [15] Yingchun Wang, Jingcai Guo, Song Guo, and Weizhan Zhang. Data quality-aware mixed-precision quantization via hybrid reinforcement learning. *arXiv preprint arXiv:2302.04453*, 2023.

- [16] Ziming Liu, Song Guo, Xiaocheng Lu, Jingcai Guo, Jiewei Zhang, Yue Zeng, and Fushuo Huo. (ml)2p-encoder: On exploration of channel-class correlation for multi-label zero-shot learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [17] Jingcai Guo and Ziming Liu. Application: Image-based visual perception. In *Machine Learning on Commodity Tiny Devices*, pages 123–144. CRC Press, 2023.
- [18] Fushuo Huo, Wenchao Xu, Jingcai Guo, Haozhao Wang, Yunfeng Fan, and Song Guo. Offline-online class-incremental continual learning via dual-prototype self-augment and refinement. *arXiv preprint arXiv:2303.10891*, 2023.
- [19] Qihua Zhou, Zhihao Qu, Song Guo, Boyuan Luo, Jingcai Guo, Zhenda Xu, and Rajendra Akerkar. On-device learning systems for edge intelligence: A software and hardware synergy perspective. *IEEE Internet of Things Journal*, 8(15):11916–11934, 2021.
- [20] Jingcai Guo. Learning robust visual-semantic mapping for zero-shot learning. *arXiv preprint arXiv:2104.05668*, 2021.
- [21] Yingchun Wang, Jingcai Guo, Song Guo, Weizhan Zhang, and Jie Zhang. Exploring optimal substructure for out-of-distribution generalization via feature-targeted model pruning. *arXiv preprint arXiv:2212.09458*, 2022.
- [22] Shiheng Ma, Jingcai Guo, Song Guo, and Minyi Guo. Position-aware convolutional networks for traffic prediction. *arXiv preprint arXiv:1904.06187*, 2019.
- [23] Jingcai Guo. An improved incremental training approach for large scaled dataset based on support vector machine. In *2016 IEEE/ACM 3rd International Conference on Big Data Computing Applications and Technologies (BDCAT)*, pages 149–157. IEEE, 2016.
- [24] Fushuo Huo, Wenchao Xu, Song Guo, Jingcai Guo, Haozhao Wang, and Ziming Liu. Procc: Progressive cross-primitive consistency for open-world compositional zero-shot learning. *arXiv preprint arXiv:2211.12417*, 2022.
- [25] Qihua Zhou, Ruibin Li, Song Guo, Yi Liu, Jingcai Guo, and Zhenda Xu. Cadm: Codec-aware diffusion modeling for neural-enhanced video streaming. *arXiv preprint arXiv:2211.08428*, 2022.
- [26] Jingcai Guo and Song Guo. Ams-sfe: Towards an alignment of manifold structures via semantic feature expansion for zero-shot learning. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, pages 73–78. IEEE, 2019.
- [27] Yingchun Wang, Song Guo, Jingcai Guo, Weizhan Zhang, Yida Xu, Jie Zhang, and Yi Liu. Efficient stein variational inference for reliable distribution-lossless network pruning. *arXiv preprint arXiv:2212.03537*, 2022.
- [28] Jingcai Guo and Song Guo. Ee-ae: An exclusivity enhanced unsupervised feature learning approach. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3517–3521. IEEE, 2019.
- [29] Qihua Zhou, Song Guo, Zhihao Qu, Jingcai Guo, Zhenda Xu, Jiewei Zhang, Tao Guo, Boyuan Luo, and Jingren Zhou. Octo:{INT8} training with loss-aware compensation and backward quantization for tiny on-device learning. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pages 177–191, 2021.
- [30] Jingcai Guo, Song Guo, Jie Zhang, and Ziming Liu. Fed-fsnet: Mitigating non-iid federated learning via fuzzy synthesizing network. *arXiv preprint arXiv:2208.12044*, 2022.
- [31] Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. Fastspeech: Fast, robust and controllable text to speech. *Advances in Neural Information Processing Systems*, 32, 2019.
- [32] Saeed Anwar and Nick Barnes. Real image denoising with feature attention. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3155–3164, 2019.
- [33] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9799–9808, 2020.
- [34] Ananth Balashankar and Lakshminarayanan Subramanian. Learning faithful representations of causal graphs. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 839–850, 2021.

- [35] Han Zhang, Jing Yu Koh, Jason Baldridge, Honglak Lee, and Yinfei Yang. Cross-modal contrastive learning for text-to-image generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 833–842, 2021.
- [36] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Slsgd: Secure and efficient distributed on-device machine learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 213–228. Springer, 2019.
- [37] Keiichi Ochiai, Kohei Senkawa, Naoki Yamamoto, Yuya Tanaka, and Yusuke Fukazawa. Real-time on-device troubleshooting recommendation for smartphones. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2783–2791, 2019.
- [38] Sujith Ravi. Efficient on-device models using neural projections. In *International Conference on Machine Learning*, pages 5370–5379, 2019.
- [39] Ilai Bistriz, Ariana Mann, and Nicholas Bambos. Distributed distillation for on-device learning. *Advances in Neural Information Processing Systems*, 33, 2020.
- [40] Wenhan Xian, Feihu Huang, and Heng Huang. Communication-efficient frank-wolfe algorithm for nonconvex decentralized distributed learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 10405–10413, 2021.
- [41] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- [42] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, pages 1273–1282. PMLR, 2017.
- [43] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. In *Advances in neural information processing systems*, pages 4424–4434, 2017.
- [44] Theodora S Brisimi, Ruidi Chen, Theofanie Mela, Alex Olshevsky, Ioannis Ch Paschalidis, and Wei Shi. Federated learning of predictive models from federated electronic health records. *International journal of medical informatics*, 112:59–67, 2018.
- [45] Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *International Conference on Machine Learning*, pages 634–643. PMLR, 2019.
- [46] Kang Loon Ng, Zichen Chen, Z Liu, H Yu, Y Liu, and Q Yang. A multi-player game for studying federated learning incentive schemes. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI 2020)*, pages 5179–5281, 2020.
- [47] Hao Wang, Zakhary Kaplan, Di Niu, and Baochun Li. Optimizing federated learning on non-iid data with reinforcement learning. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*, pages 1698–1707. IEEE, 2020.
- [48] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. In *NeurIPS*, 2020.
- [49] Xueyang Tang, Song Guo, and Jingcai Guo Guo. Personalized federated learning with contextualized generalization. In *IJCAI*, 2022.
- [50] Kevin Hsieh, Amar Phanishayee, Onur Mutlu, and Phillip Gibbons. The non-iid data quagmire of decentralized machine learning. In *International Conference on Machine Learning*, pages 4387–4398. PMLR, 2020.
- [51] Silvana Castano and Valeria De Antonellis. Global viewing of heterogeneous data sources. *IEEE Transactions on Knowledge and Data Engineering*, 13(2):277–297, 2001.
- [52] Jack Goetz, Kshitiz Malik, Duc Bui, Seungwhan Moon, Honglei Liu, and Anuj Kumar. Active federated learning. *arXiv preprint arXiv:1909.12641*, 2019.
- [53] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018.

- [54] S Ji, S Pan, G Long, X Li, J Jiang, and Z Huang. Learning private neural language modeling with attentive aggregation. In *The 2019 International Joint Conference on Neural Networks (IJCNN)*, 2019.
- [55] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *International Conference on Machine Learning*, pages 12878–12889. PMLR, 2021.
- [56] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1175–1191, 2017.
- [57] Jinhyun So, Başak Güler, and A Salman Avestimehr. Turbo-aggregate: Breaking the quadratic aggregation barrier in secure federated learning. *IEEE Journal on Selected Areas in Information Theory*, 2(1):479–489, 2021.
- [58] Pascal Vincent, Hugo Larochelle, Isabelle Lajoie, Yoshua Bengio, Pierre-Antoine Manzagol, and Léon Bottou. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of machine learning research*, 11(12), 2010.
- [59] Wenqing Chu and Deng Cai. Stacked similarity-aware autoencoders. In *IJCAI*, pages 1561–1567, 2017.
- [60] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27:2672–2680, 2014.
- [61] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [62] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020.
- [63] Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. Deep models under the gan: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 603–618, 2017.
- [64] Zhibo Wang, Mengkai Song, Zhifei Zhang, Yang Song, Qian Wang, and Hairong Qi. Beyond inferring class representatives: User-level privacy leakage from federated learning. In *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*, pages 2512–2520. IEEE, 2019.
- [65] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. In *Advances in Neural Information Processing Systems*, pages 14774–14784, 2019.
- [66] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [67] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- [68] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.