

FedCon: A Contrastive Framework for Federated Semi-Supervised Learning

Zewei Long⁺
Department of CS
UIUC
Champaign, USA
zeweil2@illinois.edu

Jiaqi Wang
College of IST
Penn State University
State College, USA
jqwangi@psu.edu

Yaqing Wang
School of ECE
Purdue University
West Lafayette, USA
wang5075@purdue.edu

Houping Xiao
Institute for Insight
Georgia State University
Atlanta, USA
hxiao@gsu.edu

Fenglong Ma
College of IST
Penn State University
State College, USA
fenglong@psu.edu

Abstract—Federated Semi-Supervised Learning (FedSSL) has gained rising attention from both academic and industrial researchers, due to its unique characteristics of co-training machine learning models with isolated yet unlabeled data. Most existing FedSSL methods focus on the classical scenario, i.e., the labeled and unlabeled data are stored at the client side. However, in real world applications, client users may not provide labels without any incentive. Thus, the scenario of labels at the server side is more practical. Since unlabeled data and labeled data are decoupled, most existing FedSSL approaches may fail to deal with such a scenario. To overcome this problem, in this paper, we propose FedCon, which introduces a new learning paradigm, i.e., contractive learning, to FedSSL. Experimental results on three datasets show that FedCon achieves the best performance with the contractive framework compared with state-of-the-art baselines under both IID and Non-IID settings. Besides, ablation studies demonstrate the characteristics of the proposed FedCon framework.

Index Terms—Federated Learning, Semi-Supervised Learning, Contrastive Learning.

I. INTRODUCTION

Federated Learning (FL), due to its unique characteristic of co-training machine learning models from fragmented data without leaking privacy [1], [2], [3], has been widely applied in different applications [4], [5], [6]. Most of existing FL studies [7], [8], [9] usually assume that the data stored in the local clients are fully annotated with ground-truth labels, but the server does not have any labeled data. However, this kind of assumption may be too strong for some real-world applications, since client users do not have enough incentives and efforts to label their generated data. Thus, a more practical and realistic scenario for federated learning is that the server holds all the labeled data, and clients have only unlabeled data as shown in Figure 1. This scenario belongs to *Federated Semi-Supervised Learning* (FedSSL). Thus, how to utilize unlabeled data residing on local clients to learn the global model is a new challenge for FL.

Recently, a few approaches [10] are proposed to tackle this challenge by integrating classical semi-supervised learning techniques into the federated learning framework, such as FedSem [11], FedMatch [12], and SSFL [13]. FedSem employs the pseudo-labeling to generate fake labels for unlabeled

⁺ This work was done when the first author remotely worked at Penn State University.

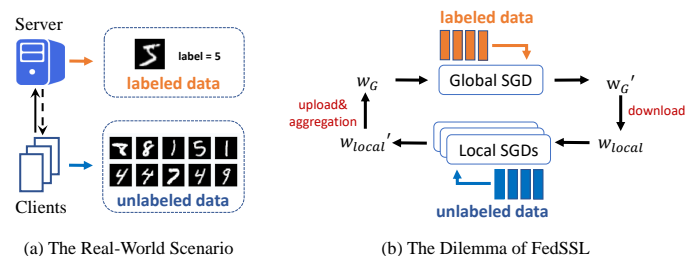


Fig. 1. Illustration of federated semi-supervised learning. (a) *The real-world scenario*. The labeled data are only available at the server, while unlabeled data are available at the local clients. (b) *The dilemma of FedSSL*. For each training round, the model cannot get access to unlabeled data at the clients and labeled data at the server simultaneously. In FedSSL, we firstly train the global model by global stochastic gradient descent (SGD) with labeled data at the server; then the local clients download the model and train with unlabeled data by local SGD; the local models are uploaded to the server and aggregated into a new global model for the next round.

data based on the trained FedAvg [1] model with labeled data. FedMatch introduces a new inter-client consistency loss and decomposition of the parameters to learn the labeled and unlabeled data separately. SSFL discusses the gradient diversity in FL and utilizes group normalization (GN) and group averaging (GA) to improve the performance. What these aforementioned methods have in common is that they focus on modifying semi-supervised methods to accommodate the new FedSSL setting. However, due to the decoupling of labeled data and unlabeled data, traditional semi-supervised methods will have a large performance loss [14]. Thus, it is urgent to design a new and specified semi-supervised learning framework for federated learning.

Towards this end, in this paper, we propose a novel and general federated semi-supervised learning framework, named FedCon, as shown in Figure 2. In particular, FedCon introduces a new **Contrastive** network into **Federated** learning to handle the challenge of the unlabeled data and further employs a unique two top-layer structure to solve the decoupling issue of labeled data and unlabeled data.

The contrastive network consists of two sub-networks: an online net and a target net. The online net keeps updating the parameters from the training data, and the target net updates slowly with the momentum mechanism and reserves the long-

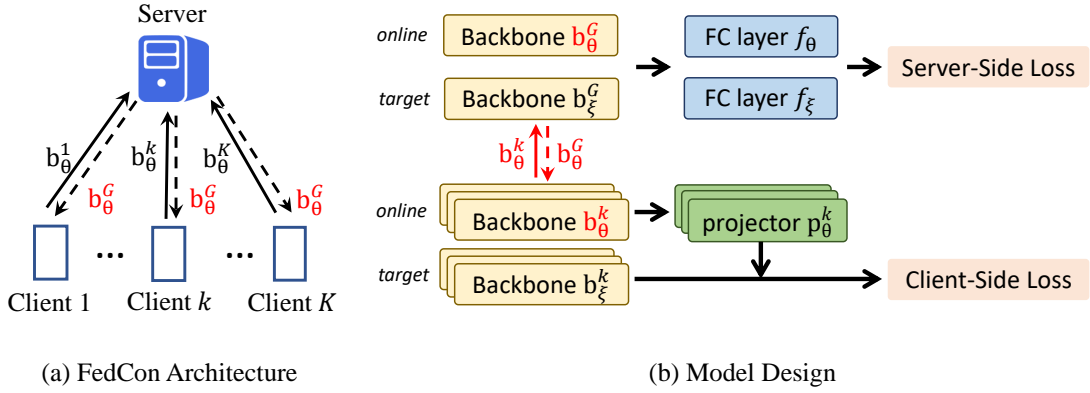


Fig. 2. Overview of the proposed FedCon Framework. (a) *FedCon architecture*. Each client k updates the contrastive network with its own data and uploads the model to the server. At the server side, the backbone of global model b_{θ}^G will be updated by averaging all the local parameters and further distributed to each client after the server update. (b) *Model design*. The training of labeled data and unlabeled data shares most architectures except two output layers and loss functions.

term information from the previous training. This division of labor enables the proposed FedCon framework to memorize the non-IID information and solve the data heterogeneity issue, which is one of key challenges in federated learning.

Except for utilizing a novel contrastive network, we also propose two different top-layer structures and loss functions for server-side and client-side updates. At the **server side**, we first train the model with labeled data. In particular, we use the designed contrastive network working with a backbone encoder b to learn pairs of representations. Then a fully-connected (FC) layer is used to make predictions. We employ the cross-entropy classification loss and consistency loss at the server side to train the model. The learned parameters will distribute to each client. At the **client side**, we use the similar strategy to learn pairs of representations. However, since all the data are unlabeled, which limits us to use the previous two losses. Intuitively, if a pair of presentations is from the same input, they should be close to each other even in the projection space. This motives us to calculate the mean-squared-error loss for unlabeled data stored in each client. After training at the client side, the client parameters will be uploaded to the server.

In summary, the main contributions of this work are summarized as follows:

- We recognize that the decoupling issue of labeled and unlabeled data may limit the power of existing federated semi-supervised learning approaches in real world scenarios.
- We propose a general, novel, and robust framework for federated semi-supervised learning, called FedCon¹, which handles the challenge of unlabeled data by introducing a contrastive network and two-output design.
- We show that FedCon outperforms state-of-the-art FedSSL baselines on three public datasets under both IID and non-IID scenarios. Moreover, ablation studies show that FedCon is more resilient to hyperparameter changes,

which further demonstrates the great robustness of the proposed FedCon.

II. PRELIMINARIES

In this section, we introduce some basic notations for FedSSL. FedSSL is a new collaborative learning paradigm, which aims to learn a global model from one server and several local clients using both labeled and unlabeled data. As mentioned in the Introduction section, in this paper, we focus on a more practical and realistic scenario – the labeled data only store at the server side, and for all the clients, they don't have any labeled data.

A. Input Data

Let $\mathcal{D}_L = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$ represent a set of labeled data stored in the server, where \mathbf{x}_i ($i \in \{1, \dots, n\}$) is a data instance, $y_i \in \{1, \dots, C\}$ is the corresponding label, and C is the number of label categories. For the k -th client, $\mathcal{D}_U^k = \{(\mathbf{x}_1^k), \dots, (\mathbf{x}_{n_k}^k)\}$ represents a set of unlabeled data, where n_k is the number of unlabeled data. Note that the unlabeled data distributions at different clients/users may not follow the IID distribution, i.e., the Non-IID issue.

B. Model Learning

Similar to the common federated learning setting, the training process is done by multiple rounds of exchanging and updating model parameters between the server and clients. In each round of communication, we allow a small number of clients denoted as B ($B \ll K$) to connect to the server and participate in the training process. Next, we demonstrate the FedSSL training pipeline, which can be slightly different from the standard FL setup.

At the beginning of each round, the global model parameters θ^G will be optimized by minimizing the loss function $\mathcal{L}_S(\mathcal{D}_L)$ with the labeled data \mathcal{D}_L in the server. Then, each selected client such as k will download the global model parameters θ^G to train its local model parameters θ^k . In the k -th client, the local model parameters θ^k will be updated by minimizing

¹The source code of the proposed FedCon framework is publicly available at <https://anonymous.4open.science/r/fedcon-pytorch-E151>

the loss function $\mathcal{L}_C(\mathcal{D}_{ij}^k)$. After the local update, the local parameters $\{\theta^1, \dots, \theta^B\}$ will be uploaded to the server for aggregating the new global model θ^G in the next round. This procedure is repeated until θ^G converges.

III. FEDCON FRAMEWORK

To alleviate the new dilemma caused by unlabeled data, we propose FedCon, an efficient FedSSL framework as shown in Figure 2. A contrastive network is employed as the general model to effectively handle the unlabeled data, which consists of two nets, i.e., an online net and a target net. We also design two different top layer structures and loss functions for server side and client side, respectively. Next, we will give the details of the proposed FedCon framework.

A. Contrastive network

The goal of FedCon is to learn a model from both labeled and unlabeled data, which are distributed in the server and clients. Therefore, our model is built for (1) learning high-dimensional representations for classification tasks from unlabeled data and (2) learning classification model from labeled data based on the unlabeled-pretraining model. To achieve this goal, we employ the contrastive network to unite the global model and local models, which uses two subnetworks for model learning, i.e., the online and target nets. The online net keeps updating the parameters from the training data, and the target net updates slowly with the momentum mechanism and reserves the long-term information from the previous training.

The training process of the contrastive network is introduced as follows. For the online net θ , we update its parameters by minimizing the loss with SGD. Then we define the target net parameters ξ_t at training step t as the exponential moving average (EMA) of successive θ_t . Specifically, given a target decay rate $\alpha \in [0, 1]$, after each training step we perform the following update, as:

$$\xi_t = \alpha \xi_{t-1} + (1 - \alpha) \theta_t. \quad (1)$$

This unique design has two obvious benefits: (1) The existing of two networks enables the proposed FedCon framework to decrease the model diversity by adding the EMA. (2) The contrastive network promises the performance gain for training the unlabeled data at the clients and labeled data at the server. It unites the global model and local models with minimum cost.

B. Server-side Design

As mentioned in the previous subsection, we use the contrastive network to extract the information from the server data. Given an input image x_i , FedCon produces two new views $x_i + \eta$ and $x_i + \eta'$ from x_i by applying image augmentation techniques. From the first augmented view $x_i + \eta$, the online net outputs a representation $y_\theta \triangleq b_\theta(x_i + \eta)$ and a prediction $c_\theta \triangleq f_\theta(y_\theta)$, where b_θ is the backbone such as convolution layers, and f_θ is the FC layer. The target network outputs $y_\xi \triangleq b_\xi(x_i + \eta')$ and the target prediction $c_\xi \triangleq f_\xi(y'_\theta)$ from the second augmented view $x + \eta'$ as shown in Figure 3.

We use the outputs of contrastive network, i.e., c_θ and c_ξ , and the ground-truth label y_i to calculate the loss for the server data. We apply the widely used **cross-entropy loss** as the classification loss for the labeled data, i.e.,

$$\mathcal{L}_\theta \triangleq \frac{1}{n} \sum_{i=1}^n \sum_{c=1}^C p(y_i = c) \log f_\theta(b_\theta(x_i + \eta)) \quad (2)$$

Besides, we introduce the **consistency loss** into model training. The benefit of using this loss is to increase the model generalizability to deal with a more common scenario that the server has unlabeled data. In particular, given two perturbed inputs $x_i + \eta$ and $x_i + \eta'$, the consistency loss disciplines the difference between the online net's predicted probabilities $c_\theta \triangleq f_\theta(b_\theta(x_i + \eta))$ and the target net's predicted probabilities $c_\xi \triangleq f_\xi(b_\xi(x_i + \eta'))$. The consistency loss is typically represented by the Mean Squared Error (MSE) as follows:

$$\mathcal{J}_{\theta, \xi} \triangleq \frac{1}{n+m} \sum_{i=1}^{n+m} \|c_\theta - c_\xi\|^2, \quad (3)$$

where n denotes the total number of labeled data, and m denotes the total number of unlabeled data in the server.

We symmetrize the loss L_θ in Eq. (2) and $\mathcal{J}_{\theta, \xi}$ in Eq. (3) by separately feeding $x + \eta'$ to the online net and $x + \eta$ to the target net to compute $\tilde{\mathcal{L}}_\xi$ and $\tilde{\mathcal{J}}_{\theta, \xi}$. In each training step, we perform a stochastic optimization step to minimize $\mathcal{L}_S = (\mathcal{L}_\theta + \mathcal{J}_{\theta, \xi} + \tilde{\mathcal{L}}_\xi + \tilde{\mathcal{J}}_{\theta, \xi})/2$ with respect to θ only. The optimization of new ξ is based on the moving average of θ and the current ξ using Eq. (1). The global model's dynamics are summarized as follows:

$$\begin{aligned} \theta &\leftarrow \text{optimizer}(\theta, \nabla_\theta \mathcal{L}_S, \eta), \\ \xi &\leftarrow \tau \xi + (1 - \tau) \theta, \end{aligned} \quad (4)$$

where *optimizer* is an optimizer, and η is the learning rate.

After training with labeled data, the backbone b_θ of the online net will be broadcast to selected clients while the fully-connected layer f_θ is stored in the server. When the local training is finished, the backbones b_θ will be uploaded to the server and aggregated into the global backbone, i.e. $b_\theta^G \leftarrow \frac{1}{B} \sum_{k=1}^B b_\theta^k$. The aggregated backbone and the fully-connected layer will be stitched together to form a new global online net θ for the next round learning.

C. Client-side Design

While the server-side model aims to learn the classification task, the goal of training client-side models is to learn the high-dimensional representations from the unlabeled data. Since the unlabeled data cannot provide the ground-truth labels as the labeled data do, there is no need to add an additional network architecture to learn the representations on the label space, i.e. c_θ . Therefore, we only keep the backbone b_θ of the global model and leave the fully-connected layer f_θ in the server, which only focuses on label space representation learning.

Similar to the server side, we output two representations $y_\theta \triangleq b_\theta(x_i + \eta)$ and $y_\xi \triangleq b_\theta(x_i + \eta')$ from two augmented views with the contrastive network. A naive solution is to

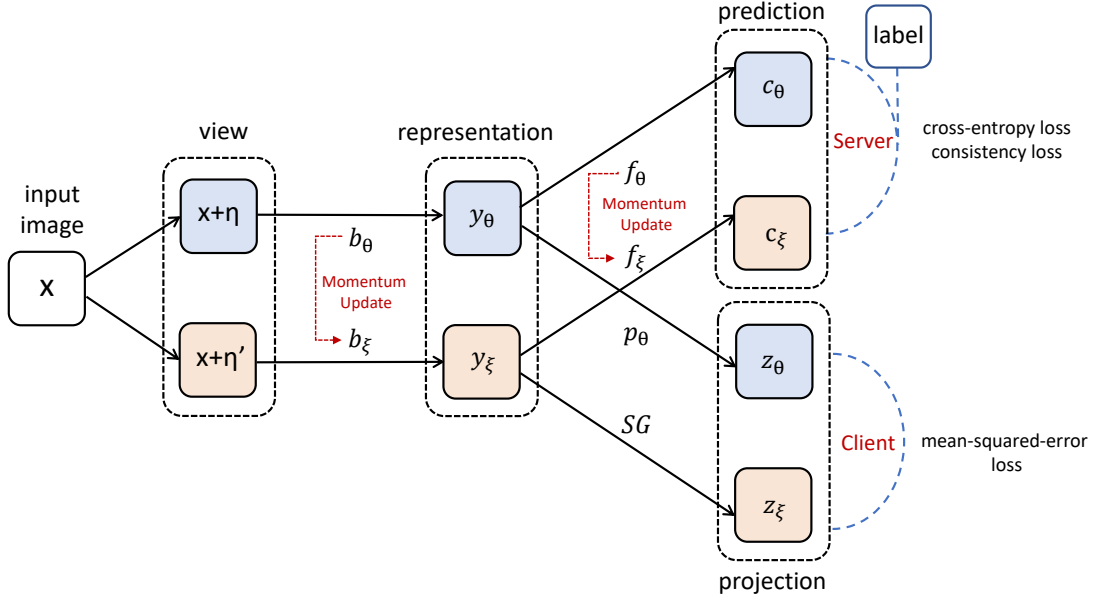


Fig. 3. The model’s architecture. FedCon shares the same backbone b_θ & b_ξ at both server and clients. In the client side, FedCon minimizes a mean-squared-error loss between z_θ and z_ξ , where the online net θ are the trained net, the target net ξ is updated by exponential moving average. In the server side, FedCon minimizes a cross-entropy loss and consistency loss between two model outputs and true label. “SG” is shorten for stop gradient.

calculate the consistency loss from these two representations. However, this loss can easily collapse into one constant solution without labeled data and cross-entropy loss. Therefore, as shown in Figure 3, a projector p_θ is added to θ , and the stop gradient (SG) is implemented in ξ to encourage encoding more information within the online projection and avoid collapsed solutions as [15] does. To utilize the unlabeled data, we use the output of online projector $z_\theta \triangleq p_\theta(y_\theta)$ and the output of target backbone $z_\xi \triangleq y_\xi$ to calculate the loss, i.e.,

$$\mathcal{L}_{\theta,\xi} \triangleq \|z_\theta - z'_\xi\|_2^2. \quad (5)$$

We also symmetrize the loss $\mathcal{L}_{\theta,\xi}$ in Eq. (5) by separately feeding $x + \eta'$ to the online net and $x + \eta$ to the target net to compute $\hat{\mathcal{L}}_{\theta,\xi}$. At each training step, we perform a stochastic optimization step to minimize $\mathcal{L}_C = (\mathcal{L}_{\theta,\xi} + \hat{\mathcal{L}}_{\theta,\xi})/2$ with respect to θ only, and then update ξ based on θ . The local model’s dynamics are summarized as

$$\begin{aligned} \theta &\leftarrow \text{optimizer}(\theta, \nabla_\theta \mathcal{L}_C, \eta), \\ \xi &\leftarrow \mu\xi + (1 - \mu)\theta. \end{aligned} \quad (6)$$

After the training of local unlabeled data, the backbone b_θ of the online work is uploaded to server while the projector p_θ is stored in the client. When this client is selected for the next local training, the backbone b_θ from the server and the projector p_θ in the client will be combined to build the new local online net θ_k .

IV. EXPERIMENTS

In this section, we evaluate the performance of the proposed FedCon framework, including both the IID and non-IID settings. We observe that the proposed FedCon achieves

the best performance compared with all existing baselines. Furthermore, we conduct ablation experiments to analyze the efficiency of our model. Note that the main experiments focus on the scenario that there is no unlabeled data on the server. However, we also conduct a preliminary experimental analysis to validate the performance changes of the proposed FedCon in a more general scenario, that is, incorporating unlabeled data on the sever. The results are shown in Section IV-E.

A. Experimental Setup

Datasets. In our experiments, we use three public datasets, including MNIST², CIFAR-10³, and SVHN⁴. The MNIST dataset is divided into a training set of 60,000 images and a test set of 10,000 images. There are 50,000 training samples and 10,000 testing samples in the CIFAR-10 dataset. In the SVHN dataset, 73,257 digits are used for training and 26,032 digits for testing. These three datasets are all used for the image classification task with 10 categories.

IID and Non-IID Settings. In the experiments, each dataset will be randomly shuffled into two parts, i.e., labeled and unlabeled data. Given γ as the ratio of the labeled data on the entire dataset, there are $|D| * \gamma$ labeled data at the server side, and $|D| * (1 - \gamma)$ unlabeled data are distributed to clients, where $|D|$ is the number of training data. For the **IID** setting, both labeled and unlabeled data all have C categories. We equally distribute the unlabeled data to each client, i.e., the number of unlabeled data in each category is the same. In the **Non-IID** setting, the labeled data on the server have all the 10

²<http://yann.lecun.com/exdb/mnist/>

³<https://www.cs.toronto.edu/~kriz/cifar.html>

⁴<http://ufldl.stanford.edu/housenumbers/>

TABLE I
MODEL ARCHITECTURES.

Dataset	ID	Operation		
MNIST	Backbone			
	1	Convolution ($10 \times 5 \times 5$) + Max Pooling (2×2)		
	2	Convolution ($20 \times 5 \times 5$) + Max Pooling (2×2)		
	Server-side		Client-side	
3	Fully Connected (320×50) + ReLU	MLP (320×320)		
4	Fully Connected (50×10) + Softmax	-		
CIFAR & SVHN	Backbone			
	1	Convolution ($32 \times 3 \times 3$) + BatchNorm + ReLU		
	2	Convolution ($64 \times 3 \times 3$) + ReLU + Max Pooling (2×2)		
	3	Convolution ($128 \times 3 \times 3$) + BatchNorm + ReLU		
	4	Convolution ($128 \times 3 \times 3$) + ReLU + Max Pooling (2×2) + dropout(0.05)		
	5	Convolution ($256 \times 3 \times 3$) + BatchNorm + ReLU		
	6	Convolution ($256 \times 3 \times 3$) + ReLU + Max Pooling (2×2)		
	Server-side		Client-side	
	7	Fully Connected (4096×1024) + ReLU + Dropout (0.1)	MLP (4096×4096)	
8	Fully Connected (1024×512) + ReLU + Dropout (0.1)	-		
9	Fully Connected (512×10) + Softmax	-		

TABLE II
PARAMETERS CHOSEN IN THE EXPERIMENTS.

Symbol	Value	Definition
R_G	150(SVHN)/200(others)	the number of global training round
K	100	the total number of clients
B	10	the number of active clients
R_L	1	the number of local epochs
BS_L	10	the labeled data training batch size
BS_U	50	the unlabeled data training batch size
BS_{test}	128	the testing batch size
γ	0.01/0.1	the fraction of labeled data

categories, but each client only contains 2 random categories of unlabeled data. We set γ as 0.01 or 0.1 in the following experiments.

Baselines. To fairly validate the proposed FedCon framework, we use the following state-of-the-art baselines: **1) FedAvg-FixMatch** [1], [16]: naive combinations of FedAvg with FixMatch. **2) FedProx-FixMatch** [17], [16]: naive combinations of FedProx with FixMatch. **3) FedAvg-UDA** [1], [18]: naive combinations of FedAvg with UDA. **4) FedProx-UDA** [17], [18]: naive combinations of FedAvg with UDA. **5) FedMatch** [12]: FedAvg-FixMatch with inter-client consistency and parameter decomposition. **6) SSFL** [13]: FedAvg-FixMatch with group normalization (GN) and group averaging (GA).

Image augmentations. FedCon adopts the weak data augmentation technique as in BYOL [15]. First, a random patch of the image is selected and resized to 224×224 with a random horizontal flip, followed by a color distortion, consisting of a random sequence of brightness, contrast, saturation, hue adjustments, and an optional grayscale conversion. After that, Gaussian blur and solarization are applied to the patches.

Architecture and Parameters. We use the same local model for all the baselines and FedCon on each dataset. For the MNIST dataset, we adopt a CNN [19] with two 5×5 convolution layers and two linear layers (21,840 parameters in total). For the CIFAR-10 & SVHN datasets, we apply a CNN

TABLE III
AVERAGE ACCURACY OF THREE RUNS ON THE THREE DATASETS UNDER THE IID SETTING WITH DIFFERENT RATIOS OF LABELED DATA.

Model	$\gamma = 0.01$			$\gamma = 0.1$		
	MNIST	CIFAR	SVHN	MNIST	CIFAR	SVHN
FedAvg-FixMatch	88.67%	49.75%	75.05%	95.97%	74.75%	92.10%
FedProx-FixMatch	89.08%	40.58%	56.07%	96.02%	75.76%	91.70%
FedAvg-UDA	88.67%	43.57%	57.98%	96.09%	75.32%	92.03%
FedProx-UDA	88.94%	41.23%	65.01%	96.40%	76.20%	91.65%
FedMatch	90.16%	52.64%	78.52%	96.09%	80.16%	92.09%
SSFL	90.64%	42.34%	55.70%	96.34%	76.00%	92.66%
FedCon	95.24%	54.84%	81.18%	98.08%	81.47%	93.19%

with six convolution layers and three linear layers (5,852,170 parameters in total). The details of each model architecture are shown in Table I. Other key parameters used in this paper are shown in Table II.

B. Performance Evaluation for the IID Setting

We first evaluate FedCon’s performance under the IID setting and report the mean accuracy on three datasets as shown in Table III. With the ratios of labeled data $\gamma = 0.01$, FedCon obtains 95.24% accuracy on the MNIST dataset (54.84% in CIFAR-10 and 81.18% in SVHN), which is a 4.60% (2.20% on CIFAR-10 and 2.66% on SVHN) improvement over the best FedSSL baseline. With $\gamma = 0.1$, FedCon still achieves the highest accuracy on all three datasets, but tightens the gap with respect to the classical semi-supervised baselines. This suggests that FedCon receives a significant performance gain when the labeled data in the server are extremely rare. This is because FedCon is built upon the contrastive framework, which has a better capacity to deal with unlabeled data.

We then provide analysis on baselines. FedAvg-FixMatch, FedProx-FixMatch, FedAvg-UDA, and FedProx-UDA are four baselines, which make a simple combination of classical SSL methods (FixMatch and UDA) and FL methods (FedAvg and FedProx). The experiments suggest that FedProx-based methods will receive a better performance than (or close to) FedAvg-based methods when the dataset is simple (MNIST) or the labeled data is plentiful ($\gamma = 0.1$). FixMatch-based methods outperform the UDA-based methods when the number of labeled data is extremely limited, and have a close performance to them in the contrary case. FedMatch, making a great improvement on aforementioned methods, achieves the highest performance among all the baselines. SSFL, with its unique group normalization, outperforms the simple combination methods.

C. Performance Evaluation for the Non-IID Settings

For the non-IID setting, the accuracy of almost all baselines in three datasets is lower than that of IID setting as shown in Table IV, even though they bring more new techniques into the basic SSL-based method. These results illustrate that we need to design an effective way of using unlabeled data while considering the effect of the non-IID setting. Otherwise, the data heterogeneity may degrade the model performance.

From Table IV, we can observe that FedCon outperforms all baselines under the non-IID setting. When $\gamma = 0.01$,

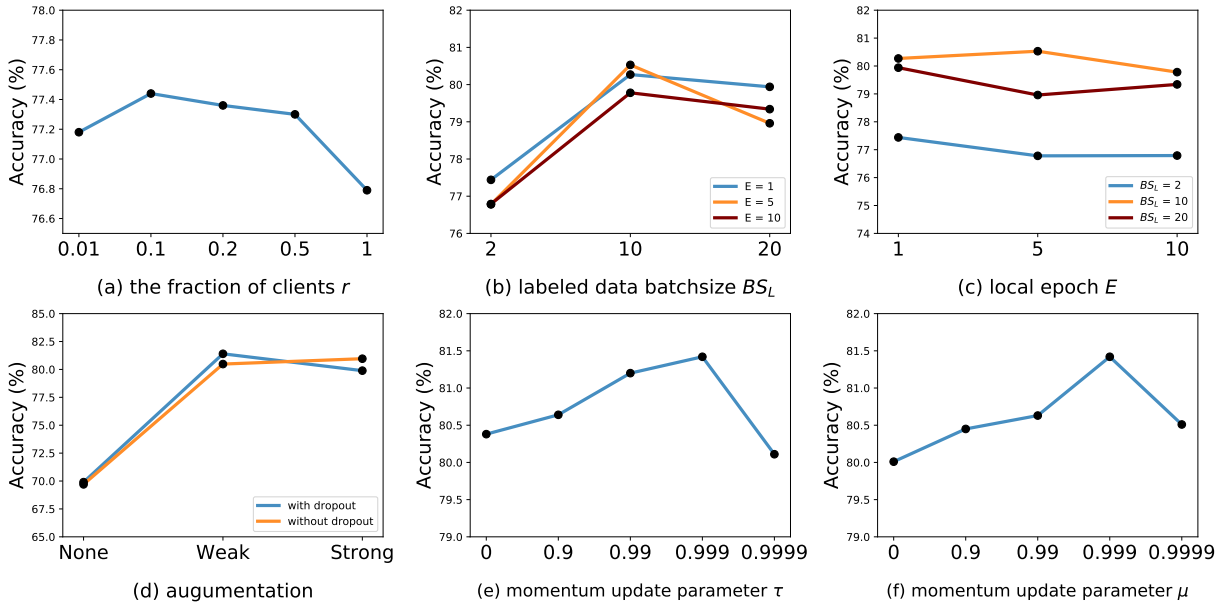


Fig. 4. Mean accuracy on CIFAR-10 with $\gamma = 0.1$ over four runs per hyperparameter setting. In each experiment, we vary one hyperparameter and fix other hyperparameters listed in Table II.

TABLE IV
AVERAGE ACCURACY OF THREE RUNS ON THE THREE DATASETS UNDER THE NON-IID SETTING WITH DIFFERENT RATIOS OF LABELED DATA.

Model	$\gamma = 0.01$			$\gamma = 0.1$		
	MNIST	CIFAR	SVHN	MNIST	CIFAR	SVHN
FedAvg-FixMatch	89.27%	48.93%	75.92%	96.15%	75.88%	92.08%
FedProx-FixMatch	91.98%	42.08%	62.67%	95.89%	75.36%	92.05%
FedAvg-UDA	89.98%	42.07%	56.64%	96.21%	75.31%	92.04%
FedProx-UDA	88.93%	43.66%	62.50%	95.59%	75.66%	92.26%
FedMatch	89.61%	52.65%	77.84%	95.70%	79.50%	91.46%
SSFL	92.22%	43.27%	67.33%	96.57%	75.16%	92.21%
FedCon	95.67%	53.25%	81.83%	98.22%	81.96%	92.67%

FedCon obtains 95.67% accuracy on the MNIST dataset (53.25% on CIFAR-10 and 81.83% on SVHN), which is a 3.45% (0.60% on CIFAR-10 and 3.99% on SVHN) improvement over the previous self-supervised state of the art. With $\gamma = 0.1$, FedCon still achieves the highest accuracy on all three datasets with an obvious performance gain. We can conclude that the contrastive network designed in FedCon has more advantages for dealing with the non-IID setting than any other baselines.

D. Ablation Studies

To assess the importance of various aspects of the model, we conduct experiments on CIFAR-10 with $\gamma = 0.1$, varying one or a few hyperparameters at a time while keeping the others fixed.

Client fraction r (Figure 4(a)). We firstly experiment with the client fraction r , which controls the amount of multi-client parallelism. Specifically, the client fraction r is defined as the fraction of the number of chosen clients B among all clients K , i.e., $\frac{B}{K}$. We report the accuracy value with total 200 rounds training with the IID case. With $BS_L = 10$, there

is no significant advantage in increasing the client fraction. This result demonstrates that the number of training unlabeled data has limited influence on the final performance, and the number of labeled data at the server domains the performance in our FedSSL setting. These results suggest us setting $r = 0.1$ for the following experiments, which strikes a good balance between computational efficiency and convergence rate.

Batch size BS_L (Figure 4(b)) & Local epoch E (Figure 4(c)). In this section, we fix $r = 0.1$ and add more computation per client on each round, either decreasing BS_L , increasing E , or both. Note that we fix $BS_U \triangleq 5 * BS_L$ to have a precise measurement of batch size. We firstly validate the model performance by test accuracy in three different batch size BS_L settings. In the real-world applications, the batch size BS_L has a close relationship to the client hardware. As BS_L is getting larger to take full advantage of available parallelism on the client hardware, the computation time will reduce dramatically in each client, which leads to higher efficiency. As the value of batch size BS_L grows, the test accuracy firstly increases rapidly before reaching the peak at $BS_L = 10$ and decreases slowly as the batch size keeps growing. Based on these observations, we fix $BS_L = 10$ in the experiments, which strikes a good balance between model performance and computation time.

Previous work [1] suggests that adding more local SGD updates per round, i.e., increasing local epoch E , can produce a dramatic decrease in communication costs when we fix the sum of local epoch. We conduct experiments to evaluate the impact of local epoch E under different batch size BS_L settings. As shown in Figure 4(c), the test accuracy for 200 communication rounds is relatively stable when local epoch E increases. Based on these results, we concludes that local

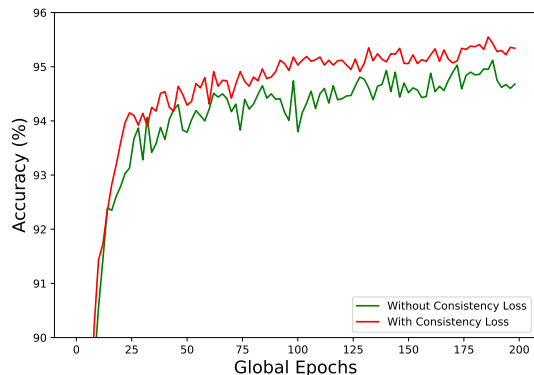


Fig. 5. Mean accuracy on MNIST ($\gamma = 0.01$ and $\beta = 0.05$) over four runs per hyperparameter setting, where γ represents the percentage of labeled data at server, and β means the percentage of unlabeled data at server.

epoch E has limited influence on model training in FedSSL, which is similar to the classical FL settings [1]. We choose $E = 1$ for most of our experiments.

Data augmentation & Dropout (Figure 4(d)). To confine the best augmentation technique under the FedSSL scenario, we conduct the ablation experiments on three different data augmentation strategy (None, Weak as BYOL [15], Strong as FixMatch [16]) and two settings (with or without dropout). We can observe that the model with weak augmentation on input images and using dropout receives the best accuracy on the test set, while the model with none or strong augmentation has a significant performance decrease. On the other hand, dropout has limited benefits when augmentation is absent. Without dropout, the strong augmentation model performs the best among other settings, but only has a small gap between the weak augmentation version. Based on these results, for most of our experiments, we choose weak augmentation and dropout as the model noise, which receives a good performance among all the settings.

Momentum update (Figure 4(e) and 4(f)). Two essential hyperparameters of FedCon are the EMA decay on server training τ and local μ training. We conduct experiments to select the best value of τ and μ and validate the sensitivity of our model to these values. We can see that in each case the good values span roughly an order of magnitude and outside these ranges the performance degrades quickly. Note that we use the ramp-up technique with an upper EMA decay during training. We choose this strategy because the online net improves quickly early in the training, and thus the target net should forget the old, inaccurate, online weights quickly. According to these results, we use $\tau = 0.999$ and $\mu = 0.999$ in each training run, which receives the best performance among all the settings.

E. Importance of Consistency Loss for Unlabeled Server Data

In this section, we conduct a preliminary experiment to validate the performance of FedCon in a more general scenario, where the server has both the labeled data and the unlabeled

data. The designed consistency loss in Eq. (3) is able to extract information from the unlabeled data in the server. Ideally, injecting unlabeled data to the model is able to increase the performance. We report the test-set accuracy and accuracy curve with total 200 rounds training in the IID case. From Figure 5, we can observe that the model with consistency loss receives a significant performance increase during the training. When $\gamma = 0.01$ and $\beta = 0.05$, FedCon with consistency loss obtains 95.55% accuracy on MNIST dataset, compared with the one without consistency loss obtains only 94.68% accuracy. This result demonstrates that the consistency loss in Eq. (3) has a great influence on the final performance when the server has both the labeled data and unlabeled data.

V. RELATED WORK

Federated learning (FL) aims to collaboratively train a joint model using data from different parties or clients. Most algorithms of FL focus on the supervised setting and mainly solve three challenges: statistical heterogeneity [20], [21], [22], system constraints [23], [24], [25], and trustworthiness [26], [27], [28]. In this paper, we only discuss the challenge of statistical heterogeneity, i.e., the Non-IID setting. To address this challenge, various algorithms have been proposed like sharing a some part of data [20], training personal model for each client [21], or adjusting the SGD convergence of FL [22]. However, relative little attention has been paid to solve the data heterogeneity in federated semi-supervised learning [11].

Federated semi-supervised learning (FedSSL), which introduces unlabeled data into federated learning, significantly increases the difficulty of the analysis of model training. Several approaches are proposed to tackle the FedSSL problem by integrating classical semi-supervised learning into the federated learning framework, such as FedSem [11], FedMatch [12], and SSFL [13]. FedSem [11] is a simple two-phase training with pseudo labeling. A study on inter-client consistency suggests that a simple application of SSL methods might not perform well in FL, and the inter-client level consistency might improve the performance [12]. The group normalization (GN) and group averaging (GA) techniques are proposed to decrease gradient diversity and further improve the Non-IID problem. However, these efforts do not propose a more general and practical algorithm and validate their potentials on new challenges of federated semi-supervised learning.

Semi-supervised learning (SSL) mitigates the requirement for labeled data by providing a way of leveraging unlabeled data [29]. The recent studies in SSL are diverse but one trend of unity. Pseudo labeling, which converts unlabeled data to labeled data, utilizes unlabeled data by labeling the data with a dynamic threshold [30]. A nature and well-working idea on consistency regularization has been widely adopted in SSL [31], [14], [32], [33], [34]. A further discussion on how loss geometry interacts with training procedures suggests that the flat platform of SGD leads to the convergence dilemma of consistency-based SSL [35]. By exploring further or mixing many practical methods, UDA [18], MixMatch [36], ReMix-Match [37], and Fixmatch [16] are proposed. In our work, we

mainly focus on utilizing the pure consistency-based methods working with federated learning.

VI. CONCLUSION

In this work, we focus on a practical and challenging setting in federated semi-supervised learning (FedSSL), i.e., all the labeled data are stored in the sever and unlabeled data are in the clients. To fully consider the new fundamental challenges caused by unlabeled data, we propose a novel yet general framework, called FedCon, which is not only effective and robust for several real-world FedSSL scenarios but also takes data heterogeneity into consideration. Experiments on three image datasets under both IID and non-IID settings demonstrate the effectiveness of the proposed FedCon framework compared with state-of-the-art baselines for the federated semi-supervised learning task. Moreover, we conduct ablation experiments to analyze the insights and illustrate the generality and characteristic of our model.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial Intelligence and Statistics*. PMLR, 2017, pp. 1273–1282.
- [2] Q. Yang, Y. Liu, T. Chen, and Y. Tong, "Federated machine learning: Concept and applications," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 10, no. 2, pp. 1–19, 2019.
- [3] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings *et al.*, "Advances and open problems in federated learning," *arXiv preprint arXiv:1912.04977*, 2019.
- [4] A. Hard, K. Rao, R. Mathews, S. Ramaswamy, F. Beaufays, S. Augenstein, H. Eichner, C. Kiddon, and D. Ramage, "Federated learning for mobile keyboard prediction," *arXiv preprint arXiv:1811.03604*, 2018.
- [5] W. Yang, Y. Zhang, K. Ye, L. Li, and C.-Z. Xu, "Ffd: A federated learning based method for credit card fraud detection," in *International Conference on Big Data*. Springer, 2019, pp. 18–32.
- [6] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi, "Federated learning of predictive models from federated electronic health records," *International journal of medical informatics*, vol. 112, pp. 59–67, 2018.
- [7] A. K. Sahu, T. Li, M. Sanjabi, M. Zaheer, A. Talwalkar, and V. Smith, "On the convergence of federated optimization in heterogeneous networks," *arXiv preprint arXiv:1812.06127*, vol. 3, 2018.
- [8] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, "On the convergence of fedavg on non-iid data," *ICLR*, 2020.
- [9] Y. Han and X. Zhang, "Robust federated learning via collaborative machine teaching," in *AAAI*, 2020, pp. 4075–4082.
- [10] Y. Jin, X. Wei, Y. Liu, and Q. Yang, "Towards utilizing unlabeled data in federated learning: A survey and prospective," 2020.
- [11] A. Albaseer, B. S. Ciftler, M. Abdallah, and A. Al-Fuqaha, "Exploiting unlabeled data in smart cities using federated learning," *2020 International Wireless Communications and Mobile Computing (IWCMC)*, 2020.
- [12] W. Jeong, J. Yoon, E. Yang, and S. J. Hwang, "Federated semi-supervised learning with inter-client consistency & disjoint learning," *ICLR*, 2021.
- [13] Z. Zhang, Z. Yao, Y. Yang, Y. Yan, J. E. Gonzalez, and M. W. Mahoney, "Benchmarking semi-supervised federated learning," *arXiv preprint arXiv:2008.11364*, 2020.
- [14] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Advances in neural information processing systems*, 2017, pp. 1195–1204.
- [15] J.-B. Grill, F. Strub, F. Altché, C. Tallec, P. H. Richemond, E. Buchatskaya, C. Doersch, B. A. Pires, Z. D. Guo, M. G. Azar *et al.*, "Bootstrap your own latent: A new approach to self-supervised learning," *arXiv preprint arXiv:2006.07733*, 2020.
- [16] K. Sohn, D. Berthelot, C.-L. Li, Z. Zhang, N. Carlini, E. D. Cubuk, A. Kurakin, H. Zhang, and C. Raffel, "Fixmatch: Simplifying semi-supervised learning with consistency and confidence," *arXiv preprint: 2001.07685*, 2020.
- [17] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *arXiv preprint arXiv:1812.06127*, 2018.
- [18] Q. Xie, Z. Dai, E. Hovy, M.-T. Luong, and Q. V. Le, "Unsupervised data augmentation for consistency training," *arXiv preprint arXiv:1904.12848*, 2019.
- [19] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [20] Y. Zhao, M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra, "Federated learning with non-iid data," *arXiv preprint arXiv:1806.00582*, 2018.
- [21] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *MLSys 2020*, 2018.
- [22] L. Huang, Y. Yin, Z. Fu, S. Zhang, H. Deng, and D. Liu, "Loadabooost: Loss-based adaboost federated machine learning on medical data," *arXiv preprint: 1811.12629*, 2018.
- [23] S. Caldas, J. Konečný, H. B. McMahan, and A. Talwalkar, "Expanding the reach of federated learning by reducing client resource requirements," *arXiv preprint arXiv:1812.07210*, 2018.
- [24] W. Luping, W. Wei, and L. Bo, "Cmfl: Mitigating communication overhead for federated learning," in *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2019, pp. 954–964.
- [25] F. Chen, M. Luo, Z. Dong, Z. Li, and X. He, "Federated meta-learning with fast convergence and efficient communication," *arXiv preprint arXiv:1802.07876*, 2018.
- [26] A. Bhowmick, J. Duchi, J. Freudiger, G. Kapoor, and R. Rogers, "Protection against reconstruction and its applications in private federated learning," *arXiv preprint arXiv:1812.00984*, 2018.
- [27] R. C. Geyer, T. Klein, and M. Nabi, "Differentially private federated learning: A client level perspective," *arXiv preprint arXiv:1712.07557*, 2017.
- [28] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical secure aggregation for federated learning on user-held data," *arXiv preprint arXiv:1611.04482*, 2016.
- [29] O. Chapelle, B. Scholkopf, and A. Zien, "Semi-supervised learning (chapelle, o. *et al.*, eds.; 2006)[book reviews]," *IEEE Transactions on Neural Networks*, vol. 20, no. 3, pp. 542–542, 2009.
- [30] D.-H. Lee, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *Workshop on challenges in representation learning, ICML*, vol. 3, no. 2, 2013.
- [31] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko, "Semi-supervised learning with ladder networks," in *Advances in neural information processing systems*, 2015, pp. 3546–3554.
- [32] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," in *ICLR*, *arXiv:1610.02242*, 2017.
- [33] T. Miyato, S.-i. Maeda, M. Koyama, and S. Ishii, "Virtual adversarial training: a regularization method for supervised and semi-supervised learning," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 8, pp. 1979–1993, 2018.
- [34] S. Park, J.-K. Park, S.-J. Shin, and I.-C. Moon, "Adversarial dropout for supervised and semi-supervised learning," *AAAI*, 2018.
- [35] B. Athiwaratkun, M. Finzi, P. Izmailov, and A. G. Wilson, "There are many consistent explanations of unlabeled data: Why you should average," *ICLR*, 2019.
- [36] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. A. Raffel, "Mixmatch: A holistic approach to semi-supervised learning," in *Advances in Neural Information Processing Systems*, 2019, pp. 5049–5059.
- [37] D. Berthelot, N. Carlini, E. D. Cubuk, A. Kurakin, K. Sohn, H. Zhang, and C. Raffel, "Remixmatch: Semi-supervised learning with distribution matching and augmentation anchoring," in *ICLR*, 2019.