

# Federated Learning without Full Labels: A Survey

Yilun Jin<sup>†</sup>      Yang Liu<sup>‡</sup>      Kai Chen<sup>†</sup>      Qiang Yang<sup>†</sup>

<sup>†</sup> Department of CSE, HKUST, Hong Kong, China

yilun.jin@connect.ust.hk, {qyang, kaichen}@cse.ust.hk

<sup>‡</sup> Institute for AI Industry Research, Tsinghua University, Beijing, China

liuy03@air.tsinghua.edu.cn

## Abstract

*Data privacy has become an increasingly important concern in real-world big data applications such as machine learning. To address the problem, federated learning (FL) has been a promising solution to building effective machine learning models from decentralized and private data. Existing federated learning algorithms mainly tackle the supervised learning problem, where data are assumed to be fully labeled. However, in practice, fully labeled data is often hard to obtain, as the participants may not have sufficient domain expertise, or they lack the motivation and tools to label data. Therefore, the problem of federated learning without full labels is important in real-world FL applications. In this paper, we discuss how the problem can be solved with machine learning techniques that leverage unlabeled data. We present a survey of methods that combine FL with semi-supervised learning, self-supervised learning, and transfer learning methods. We also summarize the datasets used to evaluate FL methods without full labels. Finally, we highlight future directions in the context of FL without full labels.*

## 1 Introduction

Deep learning (DL) algorithms have achieved great success in the past decade. Powered by large-scale data such as ImageNet [1], ActivityNet [2], BookCorpus [3], and WikiText [4], deep learning models have been successfully applied to image classification [5], object detection [6], and natural language understanding [7]. However, the success of DL relies on large-scale, high-quality data, which is not always available in practice for two reasons. On one hand, collecting and labeling data is costly, making it difficult for a single organization to accumulate and store large-scale data. On the other hand, it is also infeasible to share data across organizations to build large-scale datasets, as doing so leads to potential leakage of data privacy. In recent years, a series of laws and regulations have been enacted, such as the General Data Protection Regulation (GDPR) [8] and the California Consumer Privacy Act (CCPA) [9], imposing constraints on data sharing. Therefore, how to jointly leverage the knowledge encoded in decentralized data while protecting data privacy becomes a critical problem.

*Federated Learning* (FL) [10, 11] is a promising solution to the problem and has received great attention from both the industry and the research community. The key idea of FL is that *participants* (also known as *clients* or *parties*) exchange intermediate results, such as model parameters and gradients, instead of raw data, to jointly train machine learning models. As the raw data never leave their owners during model training, FL becomes an

---

Copyright 2023 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

**Bulletin of the IEEE Computer Society Technical Committee on Data Engineering**

attractive privacy-preserving solution to the problem of decentralized machine learning. Up to now, a plethora of FL techniques has been proposed, focusing primarily on addressing the issues of data heterogeneity [13, 15], system heterogeneity [14, 17], data privacy and security [16, 18], and communication efficiency [12, 19].

Despite the significant research efforts, there is still one important yet under-explored topic in FL, which is how to effectively leverage unlabeled data to learn better federated models. In existing efforts of FL [12, 13, 14], it is assumed that all data held by all participants are fully labeled, and that a supervised learning problem is to be solved. However, the assumption may not hold in practice for two reasons. First, participants may not be sufficiently motivated to label their data. For example, suppose a sentiment classification model is to be trained with FL, smartphone users would be unwilling to spend time and effort to label all sentences typed in the phone. Second, participants may not have sufficient expertise to label their data. For example, wearable devices record various data (e.g. heart rate, breath rate, etc.) about the user’s physical conditions, labeling which would require domain expertise in medical science and cannot be done by ordinary users. Based on the above observations, we argue that unlabeled data widely exist in real-world FL applications, and that the problem of *Federated Learning without Full Labels* is an important problem to study.

There are generally three learning paradigms in centralized machine learning (ML) that tackle the problem of learning without full labels, *semi-supervised learning* [20, 21], *self-supervised learning* [24, 23], and *transfer learning* [22], all of which have drawn much attention from researchers. Among them, semi-supervised learning aims to leverage unlabeled data to assist the limited labeled data [25, 26, 27]. Self-supervised learning aims to learn indicative feature representations from unlabeled data, which are then used to assist downstream supervised learning tasks [28, 29, 30]. Transfer learning aims to use sufficient data from a *source* domain to assist learning in a *target* domain with insufficient data [31, 32, 33], where the target domain commonly contains unlabeled data. However, despite the large number of existing works in these areas, it is not straightforward to apply them in FL due to the following challenges.

- **Isolation of labeled and unlabeled data.** In traditional semi-supervised learning and transfer learning, the server has access to both labeled and unlabeled data. However, in FL without full labels, it is common for a participant to have unlabeled data only. For example, a medical institute may not have the expertise to diagnose a complex illness, leaving all its data unlabeled. Moreover, it is not allowed in FL to exchange labeled data to solve the problem. The isolation of labeled and unlabeled data may compromise the overall performance. As observed in [34, 35], training with only unlabeled data leads to forgetting the knowledge learned from labeled data, which negatively impacts the overall performance. Therefore, it is important to bridge the knowledge between labeled and unlabeled data, without data exchange.
- **Privacy of labeled data.** In the problem of FL without full labels, the number of labeled data is often limited. Therefore, participants have to repetitively access and exchange information about them to exploit the knowledge in the labels. This leads to risks of privacy leakage of the labeled data. For example, semi-honest participants can learn to reconstruct the labeled data via gradient inversion attacks [36].
- **Data heterogeneity.** Data heterogeneity, i.e. the local data held by different participants have different data distributions, is an important property in FL that causes accuracy degradation [13, 14]. Similarly, data heterogeneity also poses challenges in the problem of FL without full labels. For example, as the number of labeled data is limited, local models tend to overfit the local data more easily, which causes a greater amount of weight divergence [37] and performance degradation.
- **Balancing performance and efficiency.** The large-scale unlabeled data in the problem creates a tradeoff between performance and efficiency. Specifically, while large-scale unlabeled data is available for training, their impacts on the model performance may be marginal, and the overall efficiency can be improved by sampling a fraction of unlabeled data without compromising model performance.

In this paper, we present a survey of the problem of FL without full labels and its existing solutions. The rest of the paper is organized as follows. Section 2 presents necessary backgrounds about FL as well as machine learning paradigms without full labels, including semi-supervised learning, self-supervised learning, and transfer learning. Sections 3, 4, 5 then review methods on federated semi-supervised learning, federated self-supervised learning, and federated transfer learning, respectively. Section 6 summarizes the datasets used for evaluating FL methods without full labels. Section 7 analyzes the similarities and differences between our work and related surveys. Finally, Section 8 presents an outlook on potential directions in the context of FL without full labels.

## 2 Preliminaries

In this section, we formally introduce backgrounds about FL, as well as the machine learning paradigms leveraging unlabeled data, semi-supervised learning, self-supervised learning, and transfer learning.

### 2.1 Federated Learning (FL)

Federated Learning aims to virtually unify decentralized data held by different participants to train machine learning models while protecting data privacy. Depending on how the data is split across participants, FL can be divided into horizontal federated learning (HFL) and vertical federated learning (VFL) [10]. In HFL, participants own data with the same feature space (e.g. participants own image data from different users), while in VFL, participants own data with the same user space but different feature spaces (e.g. a financial institute owns transaction records of a user, while an e-commerce corporation owns purchase records). In this paper, following the majority of existing research efforts, we primarily focus on HFL<sup>1</sup>, i.e. all participants share the same feature space. Formally, we consider an FL scenario with  $C$  participants, denoted as  $1, \dots, C$ . Each participant  $i$  owns a dataset  $\mathcal{D}_i = \{\mathbf{X}_{ij}, y_{ij}\}_{j=1}^{N_i}$ , where  $N_i = |\mathcal{D}_i|$  is the number of data held by participant  $i$ , and  $\mathbf{X}_{ij}, y_{ij}$  denote the features and the label of the  $j$ -th sample from client  $i$ , respectively. We use  $p_i(\mathbf{X}, y), p_i(\mathbf{X}), p_i(y|\mathbf{X})$  to denote the joint distribution, marginal distribution, and conditional distribution of client  $i$ , respectively. Denoting the model parameters as  $\theta \in \mathbb{R}^d$ , the overall optimization objective of FL is as follows,

$$\min_{\theta} f_{fl}(\theta) = \frac{1}{C} \sum_{i=1}^C f_{fl,i}(\theta), \text{ where } f_{fl,i}(\theta) = \frac{1}{N_i} \sum_{j=1}^{N_i} l(\mathbf{X}_{ij}, y_{ij}; \theta), \text{ s.t. } M_p(\theta) < \varepsilon_p, \quad (2)$$

where  $f_{fl,i}(\theta)$  is the local optimization objective of participant  $i$ , and  $l(\mathbf{X}, y; \theta)$  is a loss function, such as the cross-entropy loss for classification problems. In addition,  $M_p(\theta)$  denotes a metric measuring the privacy leakage of  $\theta$  (e.g. the budget in differential privacy (DP) [72]), and  $\varepsilon_p$  is a privacy constraint.

The training process of FL generally involves multiple communication rounds, each of which contains two steps, local training, and server aggregation.

- In the local training stage, a subset of all participants is selected. They are given the latest global model and will train the model with their local data for several epochs.
- In the server aggregation stage, participants upload their updated parameters to the server. The server aggregates received parameters via weighted averaging to obtain the global model for the next round.

Depending on the properties of participants, FL can be categorized into *cross-device* FL and *cross-silo* FL [11]. Participants of cross-device FL are commonly smart devices (e.g. phones, sensors, wearables) connected with wireless networks, while participants of cross-silo FL are commonly large organizations with connected datacenters, implying the following differences:

<sup>1</sup>Unless otherwise specified, we will use FL to refer to HFL throughout this paper.

- *Computation/communication capability.* Participants in cross-device FL commonly have limited computation (e.g. small memory, limited power supply) and communication capability (e.g. wireless network).
- *Stability.* Participants in cross-device FL are not stable and may drop out due to network breakdown.
- *Participant states.* In general, participants in cross-device FL cannot carry state vectors, in that they may only participate in one round of FL, and then drop out indefinitely.

## 2.2 Machine Learning with Unlabeled Data

### 2.2.1 Semi-supervised Learning

In semi-supervised learning, there are two datasets, a labeled dataset  $\mathcal{L} = \{\mathbf{X}_j, y_j\}_{j=1}^{|\mathcal{L}|}$ , and an unlabeled dataset  $\mathcal{U} = \{\mathbf{X}_k\}_{k=1}^{|\mathcal{U}|}$ ,  $|\mathcal{U}| \ll |\mathcal{L}|$ . In addition, the marginal distributions of  $\mathcal{L}, \mathcal{U}$  are the same, i.e.  $p_{\mathcal{L}}(\mathbf{X}) = p_{\mathcal{U}}(\mathbf{X})$ . The goal of semi-supervised learning, involving both labeled and unlabeled data, is as follows,

$$\min_{\theta} f_{semi}(\theta) = \frac{1}{|\mathcal{L}|} \sum_{j=1}^{|\mathcal{L}|} l_s(\mathbf{X}_j, y_j; \theta) + \frac{1}{|\mathcal{U}|} \sum_{k=1}^{|\mathcal{U}|} l_u(\mathbf{X}_k; \theta), \quad (3)$$

where  $l_s, l_u$  denotes the loss for labeled (supervised) and unlabeled data, respectively.

We then introduce some widely adopted techniques in semi-supervised learning.

**Pseudo-Labeling** [79]. Pseudo-labeling is a simple but effective trick for semi-supervised learning. Specifically, for each unlabeled data sample, its pseudo-label is taken as the class with the highest predicted probability,

$$\hat{y}_k = \arg \max_c g_{\theta}(\mathbf{X}_k)_c, \quad (4)$$

where  $g_{\theta}$  is the model with parameter  $\theta$ , and  $g_{\theta}(\mathbf{X}_k)_c$  denotes the predicted probability of class  $c$  for  $\mathbf{X}_k$ . There is often a confidence threshold  $\tau$ , such that pseudo-labels are only taken on confident samples with  $\hat{y}_k > \tau$ . After that, the pseudo-labels are used to supervise learning on unlabeled data, i.e.

$$l_u(\mathbf{X}_k; \theta) = l_s(\mathbf{X}_k, \hat{y}_k; \theta). \quad (5)$$

**Teacher-student Models** [25]. Teacher-student models in semi-supervised learning leverage two networks, a teacher model  $\theta_{tea}$  and a student model  $\theta_{stu}$ . On one hand, the student model is trained to be consistent with the teacher model to enhance its robustness

$$l_u(\mathbf{X}_k; \theta) = d(g_{\theta_{stu}}(\mathbf{X}_k), g_{\theta_{tea}}(\mathbf{X}_k)), \quad (6)$$

where  $d(\cdot, \cdot)$  is a distance metric. On the other hand, the teacher model is updated with moving averaging (parameterized by  $\alpha$ ) over the student model after each iteration

$$\theta_{tea} = (1 - \alpha)\theta_{tea} + \alpha\theta_{stu}. \quad (7)$$

### 2.2.2 Self-supervised Learning

Self-supervised learning aims to learn good feature representations from unlabeled data to facilitate downstream machine learning tasks. There are in general two ways to perform self-supervised learning, generative learning, and contrastive learning [24]. Generative learning trains the model to reconstruct the original data  $\mathbf{X}$  from masked data to learn the internal semantics within  $\mathbf{X}$ , while contrastive learning trains the model to distinguish between

Machine Learning Paradigm	Assumptions		Application	Limitations
	Train-test i.i.d.	Labeled Data		
Supervised Learning	✓	✓	Sufficient labeled data	Labeled data is hard to obtain.
Semi-supervised Learning	✓	Insufficient	A few labeled data + large-scale unlabeled data	Labeled and unlabeled data should have the same distribution.
Self-supervised Learning	✓	×	Large-scale unlabeled data	Cannot directly perform supervised tasks.
Transfer Learning	×	From another domain	Unlabeled data + labeled data <i>from another domain</i>	Hard to select a helpful source domain. Potential negative transfer.

Table 3: A comparison between supervised learning, semi-supervised learning, self-supervised learning, and transfer learning.

‘positive’ and ‘negative’ samples. In this survey, we primarily focus on contrastive learning, whose objective is given as follows.

$$\min_{\theta} f_{ctr}(\theta) = \sum_{\mathbf{X} \in \mathcal{U}} [d(g_{\theta}(\mathbf{X}), g_{\theta}(\mathbf{X}_+)) - \lambda \cdot d(g_{\theta}(\mathbf{X}), g_{\theta}(\mathbf{X}_-))], \quad (8)$$

where  $\mathcal{U}$  is the unlabeled dataset,  $g_{\theta}$  is a neural network parameterized by  $\theta$ ,  $\mathbf{X}_+$ ,  $\mathbf{X}_-$  are positive and negative samples sampled for data  $\mathbf{X}$ ,  $\lambda$  is the weight for negative samples, and  $d(\cdot, \cdot)$  is a distance metric. By minimizing  $f_{ctr}$ , the model  $g_{\theta}$  learns to minimize the distance between positive samples in the feature space, while maximizing the distance between negative ones. Some representative contrastive learning methods include SimCLR [45], MoCo [46], BYOL [48], and SimSiam [47]. We briefly explain their similarities and differences.

**Similarities.** All four methods employ a Siamese structure – two networks with the same architecture. One of them is called the online network  $\theta_o$  and the other is called the target network  $\theta_{tar}$ . The main difference is that the online network is directly updated via gradient descent, while the target network is generally not.

**Differences.** The differences between existing self-supervised learning methods are generally three-fold.

1. **Architecture.** In SimCLR and MoCo, the online and the target networks have the same architecture. On the contrary, for SimSiam and BYOL, the online network contains an additional predictor, i.e.  $\theta_o = (\theta_o^f, \theta_o^p)$ . The predictor aims to transform features between different views, enabling additional diversity.
2. **Target Network Parameter.** For SimCLR and SimSiam, the target network shares the same parameters as the online network  $\theta_o = \theta_{tar}$ , while for BYOL and MoCo, the target network is updated with an exponential moving average similar to Eqn. 7.
3. **Negative Samples.** On one hand, SimCLR and MoCo require negative samples  $\mathbf{X}_-$ . MoCo generates negative samples from previous batches, while SimCLR takes all other samples in the same batch as negative samples. On the other hand, SimSiam and BYOL do not require negative samples (i.e.  $\lambda = 0$ ).

### 2.2.3 Transfer Learning

Both semi-supervised learning and self-supervised learning assume that the training and test data are independent and identically distributed (i.i.d.), regardless of whether labels are present. However, transfer learning [22] does not require the assumption. Specifically, transfer learning deals with multiple data distributions (also called *domains*)  $p_i(\mathbf{X}, y), i = 1, 2, \dots, T$ , where the model is trained on one, and tested on another. Without loss of generality, we assume that  $T = 2$ . We denote  $\mathcal{L}_1 = \{\mathbf{X}_{1i}, y_{1i}\}_{i=1}^{|\mathcal{L}_1|} \sim p_1(\mathbf{X}, y)$  as the *source* dataset, and  $\mathcal{U}_2 = \{\mathbf{X}_{2j}\}_{j=1}^{|\mathcal{U}_2|} \sim p_2(\mathbf{X})$  as the *target* dataset. The overall goal is to minimize the error on the target dataset. However, as there are no labeled target data, we resort to the abundant source data to learn a model that generalizes

well to the target dataset. A commonly studied optimization objective is as follows,

$$\min_{\theta_f, \theta_c} \underbrace{\sum_{i=1}^{|\mathcal{L}_1|} l_s(\mathbf{X}_{1i}, y_{1i}; \theta_f, \theta_c)}_{f_{cls}(\mathcal{L}_1; \theta_f, \theta_c)} + \lambda \cdot \underbrace{d(g_{\theta_f}(\mathcal{L}_1), g_{\theta_f}(\mathcal{U}_2))}_{f_{dom}(\mathcal{L}_1, \mathcal{U}_2; \theta_f)}, \quad (9)$$

where  $\theta_f, \theta_c$ , are parameters of the feature extractor and the classifier, respectively,  $d(\cdot, \cdot)$  is a distance metric,  $g_{\theta_f}(\mathcal{L}_1) = \{g_{\theta_f}(\mathbf{X}_{1i})\}_{i=1}^{|\mathcal{L}_1|}$  denotes the set of source features extracted by  $\theta_f$ , and  $f_{cls}, f_{dom}$  denote the classifier loss on the source domain and the domain distance between domains, respectively. Intuitively, Eqn. 9 aims to minimize the classification error on the source domain, while minimizing the distance between source domain features and target domain features. In this way, the feature extractor  $\theta_f$  is considered to extract domain-invariant features, and the classifier can be reused in the target domain. Commonly used distance metrics  $d(\cdot, \cdot)$  include  $L_2$  distance, maximum mean discrepancy (MMD) [32] and adversarial domain discriminator [33].

In addition, if an additional labeled target dataset  $\mathcal{L}_2$  is available,  $\theta_f, \theta_c$  can be further fine-tuned with  $\mathcal{L}_2$ .

Transfer learning can generally be categorized into *homogeneous transfer learning* and *heterogeneous transfer learning* [22]. Homogeneous transfer learning assumes that domains share the same feature and label space, while heterogeneous transfer learning does not make such an assumption. For example, consider a movie recommender system that would like to borrow relevant knowledge from a book recommender system. If both systems rely on text reviews and ratings for recommendation, then a homogeneous transfer learning is to be solved, with the shared feature space being texts, and the shared label space being the ratings. However, if the movie recommender wants to leverage additional video clips, then the problem becomes a heterogeneous transfer learning problem, as the book recommender does not have video features. Heterogeneous transfer learning generally requires explicit cross-domain links to better bridge heterogeneous features and labels. For example, a novel and its related movie products should have similar feature representations.

## 2.2.4 Summary and Discussion

We summarize the three learning paradigms involving unlabeled data in Table 3. As shown, supervised learning has two key assumptions, the i.i.d. property between training and test data, and sufficient labeled data. Therefore, supervised learning is not applicable when either the labeled data is insufficient, or the training and test data come from different distributions. To address the drawback, semi-supervised learning, self-supervised learning, and transfer learning are proposed to relax the two key assumptions.

- Semi-supervised learning relaxes the assumption of sufficient labeled data. With limited labeled data, semi-supervised learning aims to exploit large-scale unlabeled data that have the same distribution as labeled data with techniques such as pseudo-labeling or teacher-student models. The main limitation of semi-supervised learning is the difficulty to obtain i.i.d. unlabeled data. For example, for the task of medical imaging, the images taken from multiple hospitals may follow different distributions due to device differences, demographic shifts, etc.
- Self-supervised learning further relaxes the assumption of labeled data. It aims to learn meaningful feature representations from the internal structures of unlabeled data, such as patches, rotations, and coloring in images. The main limitation of self-supervised learning is that, although it does not require labels to learn feature representations, they cannot be directly used to perform supervised tasks (e.g. classification).
- Transfer learning further relaxes the assumption of i.i.d. train and test data. Given unlabeled data in a domain, it aims to learn from a different but related domain with sufficient labeled data, and to transfer helpful knowledge to the unlabeled data. The main limitation of transfer learning is that it commonly requires trial-and-errors to select an adequate source domain. When inadequate source domains are chosen, negative transfer [83] may happen which compromises model accuracy.

Setting	Method	Label Isolation	Data Privacy	Data Heterogeneity	Efficiency Tradeoff
Label-at-client	RSCFed [38]	Teacher-student model	×	Sub-consensus models & distance-weighted aggregation	×
	FedSSL [39]	Pseudo-labeling	Differential privacy (DP)	Global generative model	×
	FedMatch [35]	Pseudo-labeling	×	Inter-client consistency	Disjoint & sparse learning
	FedPU [41]	Negative labels from other clients	×	×	×
	AdaFedSemi [40]	Pseudo-labeling	×	×	Tuning confidence threshold and participation rate.
	DS-FL [42]	Ensemble pseudo-labeling	×	Entropy reduction averaging	Transmit logits, not parameters
Label-at-server	SemiFL [34]	Alternate training & Pseudo-labeling	×	×	×
	FedMatch [35]	Pseudo-labeling & Disjoint learning	×	Inter-client consistency loss	Disjoint & sparse learning

Table 4: Summary of techniques for federated semi-supervised learning. × indicates that the proposed method does not focus on this issue.

### 3 Federated Semi-supervised Learning

In this section, we present an overview of federated semi-supervised learning, whose main goal is to jointly use both labeled and unlabeled data owned by participants to improve FL. Before introducing detailed techniques, we first categorize federated semi-supervised learning into two settings following [35]:

- **Label-at-client**, where the labeled data are located at the clients, while the server only has access to unlabeled data. For example, when a company would like to train an FL model for object detection using images taken from smartphones, the company has no access to the local data of users, and labeling can only be done by users. However, users are generally unwilling to label every picture taken from their smartphones, creating a label-at-client setting for federated semi-supervised learning. Formally, the objective function of this setting is as follows,

$$\min_{\theta} \frac{1}{C} \sum_{i=1}^C f_{semi,i}(\theta), \text{ s.t. } M_p(\theta) < \varepsilon_p \quad (10)$$

where  $f_{semi,i}(\theta)$  denotes the semi supervised learning loss (Eqn. 3) evaluated on the dataset of participant  $i$ , and  $M_p, \varepsilon_p$  follow Eqn. 2.

- **Label-at-server**, where the labeled data are located at the server, while clients have only unlabeled data. For example, consider a company of wearable devices that would like to train a health condition monitoring model with FL. In this case, users generally do not have the expertise to label data related to health conditions, leaving the data at clients unlabeled. The objective can be similarly formulated as

$$\min_{\theta} \frac{1}{|\mathcal{L}|} \sum_{j=1}^{|\mathcal{L}|} l_s(\mathbf{X}_j, y_j; \theta) + \frac{1}{C} \sum_{i=1}^C \left( \frac{1}{|\mathcal{U}_i|} \sum_{k=1}^{|\mathcal{U}_i|} l_u(\mathbf{X}_{ik}; \theta) \right), \text{ s.t. } M_p(\theta) < \varepsilon_p. \quad (11)$$

Methods for each federated semi-supervised learning setting are discussed in the following sections. We also summarize existing methods in Table 4.

### 3.1 The Label-at-client Setting

The label-at-client setting of federated semi-supervised learning is similar to conventional FL (Eqn. 2), in that clients can train local models with their labeled data, and the updated parameters are aggregated by the server. Therefore, the label-at-client setting inherits the challenges of data heterogeneity, data privacy, and efficiency tradeoff from conventional FL. In addition, some clients may not have labeled data to train their local models, causing the label isolation problem. We introduce how existing works address these problems in this section.

RSCFed [38] primarily focuses on the label isolation problem and the data heterogeneity problem in federated semi-supervised learning. For local training, the teacher-student model (introduced in Section 2.2.1) is adopted for training on unlabeled data. To further address the data heterogeneity problem, RSCFed proposes a sub-consensus sampling method and a distance-weighted aggregation method. In each round, several sub-consensus models are aggregated by independently sampling multiple subsets of all participants, such that each sub-consensus model is expected to contain participants with labeled data. Moreover, the local models are weighted according to their distance to sub-consensus models, such that deviating models receive low weights and their impacts are minimized.

FedSSL [39] tackles the label isolation problem, the data privacy problem, and the data heterogeneity problem. To facilitate local training of unlabeled clients, FedSSL leverages the technique of pseudo-labeling. Further, to tackle the data heterogeneity problem, FedSSL learns a global generative model to generate data from a unified feature space, such that the data heterogeneity is mitigated by the generated data. Finally, to prevent privacy leakage caused by the generative model, FedSSL leverages differential privacy (DP) to limit the information leakage of the training data in the generative model.

FedMatch [35] proposes an inter-client consistency loss to address the data heterogeneity problem. Specifically, top- $k$  nearest clients are sampled for each client, and on each data sample, the output of the local model is regularized with those of the top- $k$  client models to ensure consistency. In addition, FedMatch proposes disjoint learning that splits the parameters for labeled and unlabeled data, and the parameters for unlabeled data are sparse. Upon updates, clients with only unlabeled data upload sparse tensors, reducing the communication cost.

FedPU [41] studies a more challenging setting within semi-supervised learning, positive and unlabeled learning, in which each client has only labels in a subset of classes. In this setting, a client has only information about a part of all classes, leading to a severe label isolation problem. To tackle the problem, FedPU derives a novel objective function, such that the task of learning the negative classes of a client is relegated to other clients who have labeled data in the negative class. In this way, each client is only responsible for learning the positive classes and can do local training by itself. Empirically, the proposed FedPU outperforms FedMatch [35] in the positive-and-unlabeled learning setting.

AdaFedSemi [40] proposes a system to achieve the tradeoff between efficiency and model accuracy in federated semi-supervised learning with server-side unlabeled data. For every round, the model is trained with labeled data at clients and aggregated at the server. The server-side unlabeled data are incorporated into the training process via pseudo-labeling. AdaFedSemi [40] identifies two key parameters to balance the tradeoff between efficiency and performance, the client participation rate  $P$ , and the confidence threshold of pseudo-labels  $\tau$ . A lower  $P$  reduces both the communication cost and the model accuracy, while a high  $\tau$  reduces the server-side computation cost while also limiting the usage of unlabeled data. Therefore, AdaFedSemi designs a tuning method based on multi-armed bandits (MAB) to tune both parameters as training proceeds. Experiments show that AdaFedSemi achieves a good balance between efficiency and accuracy by dynamically adjusting  $P$  and  $\tau$  in different training phases.

DS-FL [42] tackles a similar problem to AdaFedSemi, where clients own labeled data while the server owns unlabeled data. It proposes an ensemble pseudo-label solution to leverage the server-side unlabeled data. Specifically, instead of a single pseudo-label  $\hat{y}_k$  for a data sample  $\mathbf{X}_k$ , it averages the pseudo-labels generated by all clients, i.e.  $\hat{y}_k = \text{MEAN}_{c=1}^C g_{\theta_c}(\mathbf{X}_k)$ . This creates an ensemble of client models and offers better performance. Moreover, as only pseudo-labels are transmitted instead of model parameters, the communication cost can be



significantly saved. In addition, DS-FL observes that training on pseudo-labels leads to a high prediction entropy. It then proposes an entropy-reduced aggregation, which sharpens the local outputs  $g_{\theta_c}(\mathbf{X}_k)$  before aggregation.

### 3.2 The Label-at-server Setting

The label-at-server setting, where clients do not have any labeled data, is more challenging than the label-at-client setting. The reason is that all clients own unlabeled data only and cannot provide additional supervision signals to the FL model. As shown in [35] and [34], training with only unlabeled data may lead to catastrophic forgetting of the knowledge learned from labeled data, and thus compromises the model performance.

To address the isolation between labeled data and unlabeled data, FedMatch [35] proposes a disjoint learning scheme that involves two sets of parameters for labeled and unlabeled data, respectively. The parameters for labeled data are fixed when training on unlabeled data, and vice versa, to prevent the knowledge from being overwritten. Disjoint learning brings additional benefits in communication efficiency, in that the parameters for unlabeled data, which are transmitted between participants and the server, are set to be sparse. In addition, to address the heterogeneous data held by different clients, FedMatch proposes an inter-client consistency loss, such that local models from different participants generate similar outputs on the same data.

SemiFL [34] takes another approach to solving the challenges. It proposes to fine-tune the global model with labeled data to enhance its quality and to alleviate the forgetting caused by unsupervised training at clients. Furthermore, instead of regularizing model outputs across clients, SemiFL proposes to maximize the consistency between client models and the global model. Specifically, the global model generates pseudo-labels for client-side unlabeled data, and the local models of clients are trained to fit the pseudo-labels. Empirical results show that SemiFL yields more competitive results than FedMatch.

## 4 Federated Self-supervised Learning

In this section, we introduce how self-supervised learning can be combined with FL to learn with decentralized and purely unlabeled data. Although there are two types of self-supervised learning, generative and contrastive learning, so far only contrastive methods have been studied in the FL setting, and thus we limit the discussions within federated contrastive self-supervised learning. The objective function can be formalized as

$$\min_{\theta} \frac{1}{C} \sum_{i=1}^C f_{ctr,i}(\theta), \text{ s.t. } M_p(\theta) < \varepsilon_p, \quad (12)$$

where  $f_{ctr,i}$  denotes  $f_{ctr}$  (Eqn. 8) evaluated at participant  $i$ . Compared to FL with full supervision, federated contrastive learning does not have globally consistent labels, and thus, the local contrastive objectives may deviate from one another to a greater extent. Therefore, heterogeneous data poses a greater challenge to federated contrastive learning. Table 5 summarizes existing works in federated contrastive self-supervised learning.

FedCA [49], as one of the earliest works to study federated self-supervised learning, proposes a dictionary module and an alignment module to solve the feature misalignment problem caused by data heterogeneity. Extending SimCLR, the dictionary module in FedCA aims to use the global model to generate consistent negative samples across clients, while the alignment module uses a set of public data to align the representations generated by local models. However, the alignment module of FedCA requires sharing a public dataset, which compromises data privacy.

SSFL [52] addresses the data heterogeneity problem in federated self-supervised learning with a personalized FL framework [53, 54], in which each participant trains a unique local model instead of training a shared global model. The drawback of SSFL is that the adopted self-supervised learning method requires a large batch size, which is hard to achieve on resource-limited edge devices.

Method	Label Isolation	Data Privacy	Data Heterogeneity	Efficiency Tradeoff
FedCA [49]	SimCLR	×	Dictionary & Alignment module	×
SSFL [52]	SimSiam	×	Personalized models	×
FedU [44]	BYOL	×	Selective divergence-aware update	×
FedEMA [50]	BYOL	×	Moving average client update	×
FedX [51]	Local relation loss	×	Global contrastive & relation loss	×
Orchestra [43]	Rotation prediction & clustering	Sending local centroids instead of all representations	Bi-level clustering	×

Table 5: Summary of techniques for federated self-supervised learning. × indicates that the proposed method does not focus on this issue.

FedU [44] designs a heterogeneity-aware aggregation scheme to address data heterogeneity in federated self-supervised learning. As discussed in Section 2.2.2, there are generally two networks in contrastive learning, an online network and a target network. Therefore, how to aggregate and update the two networks in FL with data heterogeneity becomes an important research question. With empirical experiments, FedU discovers that aggregating and updating only the online network yields better performances. Moreover, as FedU extends BYOL with an additional predictor model, it is also necessary to design an update rule for it. FedU designs a divergence-aware predictor update rule, which updates the local predictor only when its deviation from the global predictor is low. These rules ensure that data heterogeneity is well captured by local and global models.

Extending FedU, FedEMA [50] presents an extensive empirical study on the design components of federated contrastive learning. It performs experiments combining FL with MoCo, SimCLR, SimSiam, and BYOL, and identifies BYOL as the best base method. Based on the results, FedRMA with a divergence-aware moving average update rule is proposed. The difference between FedEMA and FedU is that, FedU overwrites the local online model with the global online model

$$\theta_{o,c}^r = \theta_o^{r-1}, \quad (13)$$

where  $\theta_{o,c}^r$  denotes the local online model at client  $c$  and round  $r$ , and  $\theta_o^{r-1}$  denotes the global online model aggregated at the previous round. On the contrary, FedEMA updates the local online model by interpolating between the global and local online models to adaptively incorporate global knowledge, i.e.

$$\theta_{o,c}^r = (1 - \mu)\theta_o^{r-1} + \mu\theta_{o,c}^{r-1}, \quad (14)$$

where  $\mu$  is a parameter based on weight divergence,

$$\mu = \min(\lambda\|\theta_{o,c}^{r-1} - \theta_o^{r-1}\|, 1). \quad (15)$$

While FedU and FedEMA are simple and effective, they both require *stateful* clients to keep track of the divergence between local and global models, and are thus not applicable in cross-device FL.

Contrary to FedU and FedEMA, Orchestra [43] proposes a theoretically guided federated self-supervised learning method that works with cross-device FL. Orchestra is based on the theory that feature representations with good clustering properties yield low classification errors. Therefore, in addition to contrastive learning, Orchestra aims to simultaneously enhance the clustering properties of all data representations. However, sharing all data representations for clustering may cause the problem of privacy leakage. Orchestra addresses the problem with a bi-level clustering method, in which clients first cluster their data representations, and send only the local centroids to the server. The server performs a second clustering on the local centroids to obtain global clustering centroids, which are sent back to clients to compute cluster assignments. As local centroids reveal less information than all data representations, this bi-level clustering method better preserves data privacy.

Orthogonal to the above methods, FedX [51] proposes a versatile add-on module for federated self-supervised learning methods. FedX consists of both local and global relational loss terms that can be added to various contrastive learning modules. The local relational loss aims to ensure that under the local model, two augmentations of the same data sample have similar relations (similarities) to samples within the same batch  $\mathcal{B}$ , i.e.

$$\mathbf{r}_i^j = \frac{\exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_j))}{\sum_{k \in \mathcal{B}} \exp(\text{sim}(\mathbf{z}_i, \mathbf{z}_k))}, \mathbf{r}_{i+}^j = \frac{\exp(\text{sim}(\mathbf{z}_{i+}, \mathbf{z}_j))}{\sum_{k \in \mathcal{B}} \exp(\text{sim}(\mathbf{z}_{i+}, \mathbf{z}_k))}, \quad (16)$$

$$L_{rel} = \text{JS}(\mathbf{r}_i, \mathbf{r}_{i+}), \quad (17)$$

where  $\mathbf{z}_i, \mathbf{z}_{i+}$  denotes the feature representation of the  $i$ -th sample (with different augmentations) in the batch,  $\mathbf{r}_i$  denotes the (normalized) similarity between the  $i$ -th sample and all other samples, and JS denotes the Jensen-Shannon divergence. The global relational loss is similarly defined such that under the local model, two augmentations of the same data sample have similar relations to the global representations. Empirical results show that FedX is versatile and can improve the performance of various contrastive learning methods in FL, including FedSimCLR, FedMoCo, and FedU.

## 5 Federated Transfer Learning

In this section, we summarize efforts that combine FL with transfer learning (FTL). We categorize existing works in FTL into homogeneous FTL and heterogeneous FTL, whose differences are introduced in Section 2.2.3.

### 5.1 Homogeneous FTL

In this section, we introduce research works on homogeneous FTL. Assuming that there are  $S$  source domains and  $T$  target domains, each of which is held by one participant, the objective of homogeneous FTL is as follows,

$$\min_{\theta_f, \theta_c} \sum_{i=1}^S f_{cls}(\mathcal{L}_i; \theta_f, \theta_c) + \sum_{i=1}^S \sum_{j=1}^T \lambda_{ij} f_{dom}(\mathcal{L}_i, \mathcal{U}_j; \theta_f), \text{ s.t. } M_p(\theta_f, \theta_c) < \varepsilon_p, \quad (18)$$

where  $\mathcal{L}_i, \mathcal{U}_j$  denote the labeled/unlabeled dataset held by the  $i$ -th source/target domain, respectively,  $f_{cls}, f_{dom}$  follow Eqn. 9, and  $\lambda_{ij}$  are hyperparameters used to select helpful source domains. If source domain  $i$  and target domain  $j$  are similar, we can assign a high  $\lambda_{ij}$ , and vice versa. Depending on how many source or target domains are involved, we can categorize existing works into two settings, single-source, and multi-source. In the multi-source setting, selecting the most appropriate source domain poses an additional challenge compared to the single-source setting. We introduce related works in both settings in the following sections.

#### 5.1.1 Single-source Setting

The single-source setting in federated transfer learning commonly involves one server with labeled data, and multiple clients with unlabeled data. As the clients themselves may have different data distributions, each client creates a unique target domain, which requires a flexible adaptation method to tackle multiple targets.

To our knowledge, DualAdapt [56] is the first work to tackle the single-source, multi-target federated transfer learning problem. DualAdapt extends from the maximum classifier discrepancy (MCD) method [61]. Specifically, MCD involves a feature extractor  $\theta_f$  and two classifiers  $\theta_{c1}, \theta_{c2}$ , trained with the following steps iteratively:

- First,  $\theta_f, \theta_{c1}, \theta_{c2}$  are trained to minimize the error on the source domain  $\mathcal{L}_1$ .
- Second, given a target domain sample  $\mathbf{X}_t$ , we fix the feature extractor  $\theta_f$  and maximize the discrepancy between the classifiers, i.e.  $\max_{\theta_{c1}, \theta_{c2}} L_{cd} = d(g_{\theta_{c1}, \theta_{c2}}(\mathbf{X}_t), g_{\theta_f, \theta_{c2}}(\mathbf{X}_t))$ . This step aims to find target samples that are dissimilar to the source domain.

Method	Homogeneous	# Source	# Target	Label Isolation	Data Privacy	Data Heterogeneity	Efficiency Tradeoff
DualAdapt [56]	✓	1	> 1	MCD & Pseudo-labeling& MixUp approximation	×	GMM weighting	×
FRuDA [58]	✓	1	> 1	DANN [33] & Optimal collaborator selection	×	Optimal collaborator selection	Lazy update
FADA [55]	✓	> 1	1	DANN [33] & Representation sharing	×	Gap statistics weighting	×
FADE [57]	✓	> 1	1	DANN	No representation sharing	CDAN	×
EfficientFDA [60]	✓	> 1	1	Max. mean discrepancy (MMD)	Homomorphic encryption (HE)	×	Optimized HE operation
PrADA [86]	✓	2	1	Grouped DANN	Homomorphic encryption (HE)	×	×
SFTL [85]	×	1	1	Sample alignment loss	HE & Secret sharing (SS)	Sample alignment loss	×
SFHTL [84]	×	1	>1	Label propagation	Split learning [87]	Unified feature space	×

Table 6: Summary of techniques for (unsupervised) federated transfer learning. × indicates that the proposed method does not focus on this issue. 'Homogeneous' indicates whether the work focuses on homogeneous FTL (✓) or heterogeneous FTL (×). # Source, # Target denote the number of source and target domains considered in the work, respectively.

- Third, the classifiers are fixed, and the feature extractor  $\theta_f$  is trained to minimize  $L_{cd}$  to generate domain invariant features.

The FL setting creates two challenges for MCD. First, Step 2 should be taken at clients, yet as no labels are available, Step 2 may result in naive non-discriminative solutions. To address the problem, DualAdapt proposes client self-training, where pseudo-labels generated by the server model are used to train the classifiers in addition to  $L_{cd}$ . Second, to maintain a single feature extractor  $\theta_f$ , Step 3 is done at the server, which has no access to target samples  $\mathbf{X}_t$ . DualAdapt proposes to use mixup [62] to approximate target samples  $\mathbf{X}_t$ . To further mitigate the impact of domain discrepancy, DualAdapt proposes to fit Gaussian mixture models (GMM) at each participant. At each participant, samples from other participants are re-weighted via the fit GMMs, such that impacts of highly dissimilar samples are mitigated.

FRuDA [58] proposes a system for single-source, multi-target federated transfer learning with DANN [33]. Similar to DualAdapt, it also considers the setting with multiple unlabeled target domains, for which it proposes an optimal collaboration selection (OCS) method. The intuition of OCS is that, for a new target domain, instead of always transferring from the only source domain, it is also possible to transfer from an existing target domain that is closer to the new domain. To implement the intuition, OCS derives an upper bound for the transfer learning error from one domain to another,

$$\varepsilon_{CE, D_2}(h, h') \leq \theta_{CE}(\varepsilon_{L_1, D_1}(h, h') + 2\theta W(D_1, D_2)), \quad (19)$$

where  $\varepsilon_{M, D}(h, l)$  denotes the error, measured by metric  $M$ , of the hypothesis  $h$  on data distribution  $D$  with the label function  $l$ ,  $\theta_{CE}, \theta$  are constants,  $h, h'$  are source and target hypotheses,  $D_1, D_2$  are source and target data distributions, respectively, and  $W(D_1, D_2)$  denotes the Wasserstein distance between  $D_1, D_2$ . With Eqn. 19, the optimal collaborator of each target domain can be selected by minimizing the right-hand side. To further improve efficiency, a lazy update scheme, exchanging discriminator gradients every  $p$  iteration, is further proposed.

### 5.1.2 Multi-source Setting

A more challenging setting of federated transfer learning is the multi-source setting, where multiple source domains with labeled data are available to transfer knowledge to a single unlabeled target domain. In this setting, it is necessary to select a source domain with helpful knowledge without directly observing source data.

To our knowledge, FADA [55] is the first work to tackle the multi-source federated transfer learning problem. FADA extends the adversarial domain adaptation [33] method, with a domain discriminator between each source domain and the target domain. The domain discriminator aims to tell whether each feature representation

belongs to the source and the target domain, and the feature extractor is then trained to fool the domain discriminator to learn domain invariant features. To train the domain discriminator, FADA directly exchanges feature representations from both domains, which may lead to potential privacy threats. In addition, to select the most relevant source domain to transfer from, FADA proposes a source domain weighting method based on gap statistics. Gap statistics [63] measures how well the feature representations are clustered,

$$I = \sum_{r=1}^k \frac{1}{2n_r} \sum_{i,j \in C_r} \|\mathbf{z}_i - \mathbf{z}_j\|_2, \quad (20)$$

where  $\mathbf{f}_i$  denotes the feature representation of the  $i$ -th sample,  $C_1 \dots C_k$  denote the index set of  $k$  clusters, and  $n_r$  is the number of samples in cluster  $r$ . A low  $I$  indicates that the feature representations can be clustered with low intra-cluster variance, which usually indicates good features. FADA then computes how the gap statistics of the target domain drop after learning with each source domain, i.e.

$$I_i^{gain} = I_i^{r-1} - I_i^r, \quad (21)$$

where  $r$  denotes the communication round, and  $i$  denotes the source domain index. Finally, FADA applies weights on source domains via with  $\text{Softmax}(I_1^{gain}, I_2^{gain}, \dots)$ .

FADE [57] improves over FADA by not sharing representations to learn the domain discriminator, thus better protecting data privacy. Instead, the domain discriminator is kept local at each client, and is trained locally and updated via parameter aggregation. FADE theoretically shows that the design leads to the same optimal values as FADA, but empirically leads to negative impacts. The issues of the design are that the trained discriminator may have low sensitivity (and thus takes longer to converge) and user mode collapse (and thus fail to represent heterogeneous data). To address the drawbacks, FADA presents two tricks. To tackle the low sensitivity issue, FADE squares the adversarial loss such that it is more reactive under large loss values. To tackle the user mode collapse issue, FADE proposes to maximize the mutual information between users (related to classes) and representations, and implements the idea with conditional adversarial domain adaptation (CDAN) [80].

EfficientFDA [60] is another improvement over FADA in that source and target domain feature representations are encrypted with homomorphic encryption (HE) [64], and the maximum mean discrepancy (MMD) [32] is computed over ciphertexts. As homomorphic encryption incurs large computation and communication costs, EfficientFDA further proposes two ciphertext optimizations. First, ciphertexts in each batch of samples are aggregated to reduce communication overhead. Second, for computing gradients with ciphertexts, the chain rule is applied to replace ciphertext computations with plaintexts to improve computational efficiency. Experiments show that EfficientFDA achieves privacy in federated transfer learning, while being 10-100x more efficient than naive HE-based implementations.

While the above works tackle the problem with multiple source domains with the same feature space, PrADA [86] tackles a different problem, involving two source domains with different feature spaces. PrADA considers a partially labeled target domain A  $\{\mathbf{X}_l^A, y_l^A\} \cup \{\mathbf{X}_u^A\}$ , a labeled source domain B  $\{\mathbf{X}^B \in \mathbb{R}^{N_B \times D}, y^B\}$ , and a feature source domain C  $\{\mathbf{X}_C^A \in \mathbb{R}^{N_A \times D_C}\} \cup \{\mathbf{X}_C^B \in \mathbb{R}^{N_B \times D_C}\}$ . Domains A and B share the same feature space with different distributions, while domain C aims to provide rich auxiliary features for samples in both A and B. PrADA presents a fine-grained domain adaptation technique, in which features from domain C are first manually grouped into  $g$  tightly relevant feature groups. Each feature group is then assigned a feature extractor and a domain discriminator to perform fine-grained, group-level domain adaptation. In addition, to protect data privacy, the whole training process is protected with homomorphic encryption. Experiments show that with the grouped domain adaptation, PrADA achieves better transferability and interpretability.

## 5.2 Heterogeneous FTL

In this section, we introduce existing works about heterogeneous FTL. Compared to homogeneous FTL, the main difference of heterogeneous FTL is that it commonly requires cross-domain links between data (e.g. different

features of the same user ID, the same features from different users, etc.) to bridge the heterogeneous feature spaces. Formally, assuming a heterogeneous FTL setting with two parties, A and B, with data  $\mathcal{D}_A, \mathcal{D}_B$ , with  $\mathcal{D}_{AB} = \mathcal{D}_A \cap \mathcal{D}_B$  being the overlapping dataset (i.e. cross-domain links), the objective of heterogeneous FTL is

$$\min_{\theta_A, \theta_B} L_A(\mathcal{D}_A; \theta_A) + L_B(\mathcal{D}_B; \theta_B) + \lambda L_{\text{algn}}(\mathcal{D}_{AB}; \theta_A, \theta_B), \text{ s.t. } M_p(\theta_A, \theta_B) < \varepsilon_p, \quad (22)$$

where  $L_A, L_B$  are loss functions on dataset  $\mathcal{D}_A, \mathcal{D}_B$ , respectively, and  $L_{\text{algn}}$  is an alignment loss that aims to align the overlapping dataset  $\mathcal{D}_{AB}$  between domains. However, in FL, sharing sample features or labels pose potential privacy threats. How to leverage the cross-domain sample links to transfer knowledge while preserving privacy thus becomes a key challenge to solve.

To our knowledge, SFTL [85] is the first work to tackle the heterogeneous FTL problem. It considers a two-party setting and assumes that some user IDs  $\mathcal{I}_{AB}$  exist in both parties (with different features). SFTL proposes an alignment loss to minimize the difference between features of the same users to achieve knowledge transfer,

$$L_{\text{algn}} = \sum_{i \in \mathcal{I}_{AB}} d(g_{\theta_A}(\mathbf{X}_i^A), g_{\theta_B}(\mathbf{X}_i^B)), \quad (23)$$

where  $\mathcal{I}_{AB}$  denotes the overlapping user ID set,  $g_{\theta_A}, g_{\theta_B}$  denote neural network models of party A and B, and  $\mathbf{X}_i^A, \mathbf{X}_i^B$  denote the features of user  $i$  held by party A and B, respectively. In addition, SFTL addresses the data privacy problem by designing two secure protocols for SFTL, one based on homomorphic encryption, and the other based on secret sharing (SS).

The drawbacks of SFTL are that it is limited to the two-party setting, and both A and B have only partial models and cannot perform independent inference. To address these drawbacks, SFHFL [84] proposes an improved framework that supports multiple parties. The main difficulty in the multi-party heterogeneous FTL is the lack of overlapping samples and labels. To address the lack of overlapping samples, SFHFL proposes a feature reconstruction technique to complement the missing non-overlapping features. Specifically, all parties are trained to project their features into a unified latent feature space. Then, each party learns a reconstruction function that projects the unified features to raw features. With the reconstruction functions, each party can expand the feature spaces of non-overlapping samples, thus enlarging the training dataset. In addition, SFHFL proposes a pseudo-labeling method based on label propagation [20] to address the lack of labels. Specifically, a nearest neighbor graph based on feature proximity in the unified feature space is constructed, and the labels are propagated from labeled samples to unlabeled samples via the graph. Finally, to protect the privacy of labels, SFHFL is trained with split learning, such that labels are not directly shared with other parties.

## 6 Datasets and Evaluations

Benchmarking datasets are important for the development of machine learning research. In this section, we introduce commonly used datasets and benchmarks for the problem of FL without full labels in the existing literature. A summary of datasets can be found in Table 7. We find out that for both federated semi-supervised and unsupervised learning, existing works mainly partition (e.g. according to Dirichlet distributions) datasets for centralized machine learning (e.g. CIFAR-10, CIFAR-100, SVHN) manually, and manually sample a subset of labels. On the contrary, for federated transfer learning, datasets generally form natural partitions (e.g. city in GTA5, product types in AmazonReview, etc.) based on different domains. We thus conclude that real-world datasets representing realistic data heterogeneity and label isolation problems are still needed to credibly evaluate federated semi-supervised and self-supervised methods.

Dataset	FL Methods without Full Labels			Application	# Domains	# Samples	Partition
	Semi	Self	Trans.				
CIFAR-10	✓[38, 39, 35, 40, 34]	✓[52, 44, 43, 51, 50]	×	CV	1	60000	Dirichlet & Uniform
CIFAR-100	✓[38, 34]	✓[43, 44, 50]	×	CV	1	60000	Dirichlet & Uniform
SVHN	✓[40, 34]	✓[51]	×	CV	1	73257	Dirichlet & Uniform
Sent140	✓[39]	×	×	NLP	1	1600498	Natural (Twitter User)
Reuters	✓[42]	×	×	NLP	1	11228	Dirichlet
IMDb	✓[42]	×	×	NLP	1	50000	Dirichlet
Landmark-23K	×	✓[52]	×	CV	1	1600000	Natural (Location)
Digit-Five	×	×	✓[55, 58]	CV	5	107348	Natural (Style)
Office-Caltech10	×	×	✓[55, 58, 60]	CV	4	2533	Natural (Style)
DomainNet	×	×	✓[55, 58]	CV	6	416401	Natural (Style)
AmazonReview	×	×	✓[55]	NLP	4	8000	Natural (Product Category)
Mic2Mic	×	×	✓[58]	Speech	4	65000	Natural (Device Type)
GTA5	×	×	✓[56]	CV	4	25000	Natural (Location)

Table 7: Commonly used datasets for evaluating FL methods without full labels. ✓ and × indicate that the dataset has or has not been used for evaluating an FL setting without full labels, respectively. # domains, # samples denote the number of domains and the total number of samples in the dataset. Datasets with multiple domains are more commonly used for unsupervised federated transfer learning.

## 7 Related Surveys

Federated learning has attracted the attention of researchers worldwide. Therefore, there have been many survey papers that cover various aspects of FL. In this section, we summarize and analyze existing survey papers compared to our work. Table 8 shows a summary of comparisons between related surveys and ours.

First, our work differs from general surveys on FL [11, 10, 88] in that they provide comprehensive reviews on a wide range of FL aspects, including privacy preservation, communication reduction, straggler mitigation, incentive mechanisms, etc. Among them, communication and privacy are also important issues in the problem of FL without full labels and are covered in our survey. On the contrary, our survey is focused on a specific aspect, namely how to deal with unlabeled data. Second, our work also differs from surveys on semi-supervised learning [21], self-supervised learning [24], and transfer learning [22] in the centralized setting, in that while they extensively summarize machine learning techniques for unlabeled data, they fail to cover FL-specific challenges, such as label isolation, data privacy, etc. Finally, compared to surveys that focus on FL algorithms on non-i.i.d. data [89, 90, 91], our work focuses on leveraging unlabeled data to assist FL, while these surveys focus on FL with fully labeled data, but are not independent and identically distributed. Nonetheless, these surveys are related to our work in that non-i.i.d. data is an important challenge in all FL settings, and we also summarize how existing works address the challenge in the problem of FL without full labels.

The most related survey to our work is [59], which surveyed FL techniques to tackle data space, statistical, and system heterogeneity. Our work is similar to [59] in two ways. On one hand, statistical heterogeneity is a key challenge in FL, and we also summarize how existing works address the challenge in FL without full labels. On the other hand, homogeneous and heterogeneous FTL (Section 5) are powerful tools to solve statistical and data space heterogeneity, respectively, which are also covered in Sections 3 and 4 in [59]. Nonetheless, the main focus of [59] lies in supervised FL with labeled data, which is different from our work which additionally covers federated semi-supervised and self-supervised methods.

Survey Papers	Similarities	Differences
[10, 11, 88]	Similar to our survey, these papers review existing solutions to protect data privacy and reduce communication/computation overhead.	These papers cover a wide range of aspects in general FL, while our survey focuses on a specific problem of leveraging unlabeled data.
[22, 23, 24, 21, 92, 93]	Similar to our survey, these papers review machine learning methods for unlabeled data, including semi-supervised, self-supervised, and transfer learning.	These papers do not cover FL specific challenges, such as labeled data isolation, data heterogeneity, data privacy, etc.
[90, 89, 91]	Similar to our survey, these papers review methods in FL that address the problem of non-i.i.d. data (i.e. data heterogeneity).	These papers primarily focus on optimization algorithms for fully supervised FL, while our work focuses specifically on leveraging unlabeled data.
[59]	Similar to our survey, [59] covers methods to tackle data heterogeneity. Also, [59] reviews existing works on homogeneous and heterogeneous FTL.	[59] primarily focuses on heterogeneity in supervised FL, while our work focuses on leveraging unlabeled data and covers federated semi-supervised and self-supervised learning.

Table 8: Comparative analysis between our survey and related surveys.

Learning Paradigm	Main Techniques	Advantages	Disadvantages
Federated Semi-supervised Learning	Enhancing methods in centralized settings with 1. <b>Label isolation</b> : Pseudo-labeling, domain alignment, etc.	Similar formulation to conventional FL. Can directly perform supervised tasks.	Data heterogeneity inherently violates i.i.d. assumption. Large-scale unlabeled data creates an efficiency tradeoff.
Federated Self-supervised Learning	2. <b>Privacy</b> : DP, HE, etc. 3. <b>Data heterogeneity</b> : Source domain selection, divergence-aware update, etc.	Full utilization of client data. Suitable for unsupervised tasks like retrieval, clustering, etc.	Data heterogeneity inherently violates i.i.d. assumption. Need labels for supervised tasks.
Federated Transfer Learning	4. <b>Efficiency tradeoff</b> : Sample selection, communication reduction, HE optimization, etc.	Models data heterogeneity, which is a key challenge in FL. Flexible formulation (heterogeneous FTL).	Source domain selection requires intricate design or manual effort.

Table 9: A summary of techniques, advantages, and disadvantages of learning paradigms reviewed in this paper.



## 8 Conclusion and Future Directions

### 8.1 Summary of the Survey

In this paper, we present a survey about the problem of federated learning without full labels. We introduce three learning paradigms to solve the problem, federated semi-supervised learning, federated self-supervised learning, and federated transfer learning. We further review existing works in these paradigms and discuss how they address the crucial challenges, i.e. label isolation, privacy protection, data heterogeneity, and efficiency tradeoff. Table 9 shows a summary of the main techniques, advantages, and disadvantages of learning paradigms discussed in this paper. We finally present a summary of the datasets and benchmarks used to evaluate FL methods without full labels.

### 8.2 Future Directions

Compared to general FL with full supervision, the problem of FL without full labels is still under-explored. We highlight the following future directions in the context of FL without full labels.

#### 8.2.1 Trustworthiness

Trustworthiness is an important aspect in real-world machine learning systems like FL. Generally speaking, users of machine learning systems would expect a system to be private, secure, robust, fair, and interpretable, which is what trustworthiness mean in the context of FL.

Unlabeled data can play an important role in enhancing trustworthiness from multiple aspects.

- **Robustness:** A robust system requires that its output should be insensitive to small noises added to the input. A machine learning system that is not robust can significantly compromise its security in real-world applications. For example, studies [69] have shown that it is possible to tweak physical objects to fool an object detection model. In applications like autonomous driving, this property becomes a security threat. Many research works have studied how to enhance robustness with unlabeled data [68, 67]. For example, Carmon et al. and Uesato et al. [67, 70] show that pseudo-labeling, one of the most common semi-supervised learning techniques, can boost the robustness by 3-5% over state-of-the-art defense models. Deng et al. [68] additionally find out that even out-of-distribution unlabeled data helps enhance robustness. Therefore, how these techniques can be adapted in the FL setting with heterogeneous data is an interesting future direction. Also, as common methods of learning robust models (i.e. adversarial training [81]) are inefficient, it is promising to study whether FL methods without full labels can be an efficient substitute.
- **Privacy:** In real-world machine learning applications, labeling data itself is a compromise of data privacy, as domain experts have to directly observe the data. Therefore, solving the FL problem without full labels inherently leads to better data privacy. In addition, unlabeled data provides a better way of navigating through the privacy-utility tradeoff in differential privacy (DP) [72]. For example, PATE [71] shows that with an additional set of unlabeled data, it simultaneously achieves a higher model accuracy and a tighter privacy bound compared to the state-of-the-art DPSGD method [73]. Therefore, how to select and leverage unlabeled data to aggregate client knowledge privately while maintaining good model accuracy is also a promising direction.
- **Interpretability:** Interpretability indicates that a machine learning system should be able to make sense of its decision, which generally creates trust between users and system developers. There are many ways to instill interpretability in machine learning, among which disentangled representation learning [74] is a popular direction. Informally speaking, disentangled representation aims to map the inputs to latent representations where high-level factors in the input data are organized in a structured manner in the

representations (e.g. brightness, human pose, facial expressions, etc.). Thus, disentangled representations provide intuitive ways to manipulate and understand deep learning models and features.

Much progress has been made in unsupervised disentangled representation learning. For example, InfoGAN [75] learns disentangled representations by maximizing the mutual information between the features and the output. Beta-VAE [76] disentangles features by adding an independence regularization on the feature groups. Therefore, it is promising to instill interpretability in FL via unlabeled data with disentangled representations. In FL, the participants commonly hold data with varying data distributions. Therefore, how to stably disentangle the heterogeneous feature distributions from multiple participants is a challenge for interpretable FL without full labels.

- **Fairness:** As machine learning models are increasingly involved in decision-making in the daily lives of people, the models should not discriminate one group of users against another (e.g. gender, race, etc.). Informally speaking, the fairness of a machine learning model  $g_\theta$  over a sensitive attribute  $s$  can be described as the difference between the model performances given different values of  $s$ ,

$$\Delta_{s,\theta} = \|m(g_\theta|s = 1) - m(g_\theta|s = 0)\|, \quad (24)$$

where  $\|\cdot\|$  is a distance, and  $m(g_\theta|s = 1)$  is a performance metric stating how well the model performs when the sensitive attribute  $s = 1$ .

When  $m(g_\theta|s)$  does not involve labels (e.g. some groups have a higher probability to be predicted positive), FADE [57] provides a good solution to ensure group fairness. However, when  $m(g_\theta|s = 1)$  requires labeled data (e.g. classification accuracy is lower for under-represented groups), enforcing fairness with unlabeled data remains an open problem, both for general machine learning and FL.

## 8.2.2 Generalization to Unseen Domains

All the introduced techniques in this paper require at least observing the test domain such that it can work well on it. Even for federated transfer learning, some unlabeled samples in the target domain are still needed for successful adaptation. However, in real-world applications, it is often required to adapt to completely unseen domains. For example, FL models should try to adapt to new users that constantly join mobile applications, who, at the time of joining, have no interaction data available. The problem setting triggers research in federated domain generalization (FedDG). However, existing works in FedDG [77, 78] assume that all domains are fully labeled, which, as stated in this survey, is not realistic. It is thus important to study the FedDG problem under limited labeled data and large-scale unlabeled data.

## 8.2.3 Automatic FL without Full Labels

Automatic machine learning (AutoML) [82] is a class of methods that aim to achieve good model performances without manual tuning (e.g. architecture, hyperparameters, etc.). In FL without full labels, as different participants may hold heterogeneous labeled or unlabeled data, it may not be optimal for them to share the same model architecture. Integrating AutoML to FL without full labels thus enables participants to find personalized architectures to achieve the performance-efficiency tradeoff. However, participants with only unlabeled data cannot independently evaluate the performance of the model, creating challenges to automatic FL without full labels.

## References

- [1] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-scale Hierarchical Image Database. 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 248–255, 2009.
- [2] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. ActivityNet: A Large-scale Video Benchmark for Human Activity Understanding. 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 961–970, 2015.
- [3] Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books. 2015 IEEE International Conference on Computer vision (ICCV), 19–27, 2015.
- [4] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer Sentinel Mixture Models. International Conference on Learning Representations, 2017.
- [5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778, 2016.
- [6] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. 2017 IEEE International Conference on Computer Vision (ICCV), 2961–2969, 2017.
- [7] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT), 4171–4186, 2019.
- [8] The European Parliament. General Data Protection Regulation. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:02016R0679-20160504>, 27 April 2016.
- [9] State of California Department of Justice. California Consumer Privacy Act. <https://oag.ca.gov/privacy/ccpa>, 2018.
- [10] 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 (2019): 1-19.
- [11] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, and Others. Advances and Open Problems in Federated Learning. Foundations and Trends® in Machine Learning, 14.1–2 (2021), 1–210.
- [12] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient Learning of Deep Networks from Decentralized Data. International Conference on Artificial Intelligence and Statistics (AISTATS), 1273–1282, 2017.
- [13] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated Optimization in Heterogeneous Networks. Proceedings of Machine Learning and Systems (MLSys), 429–450, 2020.
- [14] Fan Lai, Xiangfeng Zhu, Harsha V. Madhyastha, and Mosharaf Chowdhury. Oort: Efficient Federated Learning via Guided Participant Selection. USENIX Symposium on Operating Systems Design and Implementation (OSDI), 19–35, 2021.

- [15] Sai Praneeth Karimireddy, Satyen Kale, Mahryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. SCAFFOLD: Stochastic Controlled Averaging for Federated Learning. International Conference on Machine Learning (ICML), 5132–5143, 2020.
- [16] 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. ACM SIGSAC Conference on Computer and Communications Security (CCS), 1175–1191, 2017.
- [17] Enmao Diao, Jie Ding, and Vahid Tarokh. HeteroFL: Computation and Communication Efficient Federated Learning for Heterogeneous Clients. International Conference on Learning Representations (ICLR), 2021.
- [18] Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. BatchCrypt: Efficient Homomorphic Encryption for Cross-Silo Federated Learning. USENIX Annual Technical Conference (ATC), 493–506, 2020.
- [19] Amirhossein Reisizadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. Fed-PAQ: A Communication-efficient Federated Learning Method with Periodic Averaging and Quantization. International Conference on Artificial Intelligence and Statistics (AISTATS), 2021–2031, 2020.
- [20] Xiaojin Jerry Zhu. Semi-supervised Learning Literature Survey. University of Wisconsin-Madison, 2005.
- [21] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised Learning. IEEE Transactions on Neural Networks, 20.3 (2009), 542–542.
- [22] Sinno Jialin Pan and Qiang Yang. A Survey on Transfer Learning. IEEE Transactions on Knowledge Discovery and Data Engineering, 22.10 (2010), 1345–1359.
- [23] Longlong Jing and Yingli Tian. Self-supervised Visual Feature Learning with Deep Neural Networks: A Survey. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 43.11 (2020), 4037–4058.
- [24] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. Self-supervised Learning: Generative or Contrastive. IEEE Transactions on Knowledge and Data Engineering (TKDE), 35.1 (2021), 857–876.
- [25] Antti Tarvainen and Harri Valpola. Mean Teachers are Better Role Models: Weight-averaged Consistency Targets Improve Semi-supervised Deep Learning Results. Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [26] Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual Adversarial Training: A Regularization Method for Supervised and Semi-supervised Learning. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), 41.8 (2018), 1979–1993.
- [27] David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicholas Papernot, Avital Oliver, and Colin A Raffel. MixMatch: A Holistic Approach to Semi-supervised Learning. Advances in Neural Information Processing Systems (NeurIPS), 2019.
- [28] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised Visual Representation Learning by Context Prediction. IEEE International Conference on Computer Vision (ICCV), 1422–1430, 2015.
- [29] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised Representation Learning by Predicting Image Rotations. International Conference on Learning Representations, 2018.

- [30] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked Autoencoders are Scalable Vision Learners. IEEE/CVF Conference on Computer Vision and Pattern Recognition, 16000–16009, 2022.
- [31] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. How Transferable are Features in Deep Neural Networks?. Advances in Neural Information Processing Systems (NIPS), 2014.
- [32] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning Transferable Features with Deep Adaptation Networks. International Conference on Machine Learning (ICML), 97–105, 2015.
- [33] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial Training of Neural Networks. Journal of Machine Learning Research (JMLR), 17.1 (2016), 2096–2130.
- [34] Enmao Diao, Jie Ding, and Vahid Tarokh. SemiFL: Semi-supervised Federated Learning for Unlabeled Clients with Alternate Training. Advances in Neural Information Processing Systems (NeurIPS), 2022.
- [35] Wonyong Jeong, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. Federated Semi-supervised Learning with Inter-client Consistency & Disjoint Learning. International Conference on Learning Representations (ICLR), 2021.
- [36] Jonas Geiping, Hartmut Bauermeister, Hannah Droöge, and Michael Moeller. Inverting Gradients-How Easy Is It to Break Privacy in Federated Learning? Advances in Neural Information Processing Systems (NeurIPS), 16937–16947, 2020.
- [37] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated Learning with Non-iid Data. arXiv preprint arXiv:1806.00582, 2018.
- [38] Xiaoxiao Liang, Yiqun Lin, Huazhu Fu, Lei Zhu, and Xiaomeng Li. RSCFed: Random Sampling Consensus Federated Semi-supervised Learning. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 10154–10163, 2022.
- [39] Chenyou Fan, Junjie Hu, and Jianwei Huang. Private Semi-supervised Federated Learning. International Joint Conference on Artificial Intelligence (IJCAI), 2009–2015, 2022.
- [40] Lun Wang, Yang Xu, Hongli Xu, Jianchun Liu, Zhiyuan Wang, and Liusheng Huang. Enhancing Federated Learning with In-cloud Unlabeled Data. IEEE International Conference on Data Engineering (ICDE), 136–149, 2022.
- [41] Xinyang Lin, Hanting Chen, Yixing Xu, Chao Xu, Xiaolin Gui, Yiping Deng, and Yunhe Wang. Federated Learning with Positive and Unlabeled Data. International Conference on Machine Learning (ICML), 13344–13355, 2022.
- [42] Sohei Itahara, Takayuki Nishio, Yusuke Koda, Masahiro Morikura, and Koji Yamamoto. Distillation-based Semi-supervised Federated Learning for Communication-efficient Collaborative Training with Non-iid Private Data. IEEE Transactions on Mobile Computing (TMC), 22.1 (2021), 191–205.
- [43] Ekdeep Lubana, Chi Ian Tang, Fahim Kawsar, Robert Dick, and Akhil Mathur. Orchestra: Unsupervised Federated Learning via Globally Consistent Clustering. International Conference on Machine Learning (ICML), 14461–14484, 2022.
- [44] Weiming Zhuang, Xin Gan, Yonggang Wen, Shuai Zhang, and Shuai Yi. Collaborative Unsupervised Visual Representation Learning from Decentralized Data. IEEE/CVF International Conference on Computer Vision (ICCV), 4912–4921, 2021.

- [45] Ting Chen, Simon Kornblith, Mohammad Norouzi, Geoffrey Hinton. A Simple Framework for Contrastive Learning of Visual Representations. International Conference on Machine Learning (ICML), 1597–1607, 2020.
- [46] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Representation Learning. IEEE/CVF Conference on Computer Vision and Pattern Recognition, 9729–9738, 2020.
- [47] Xinlei Chen and Kaiming He, Exploring Simple Siamese Representation Learning. IEEE/CVF Conference on Computer Vision and Pattern Recognition, 15750–15758, 2021.
- [48] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, and Others. Bootstrap Your Own Latent-A New Approach to Self-supervised Learning. Advances in Neural Information Processing Systems, 21271–21284, 2020.
- [49] Fengda Zhang, Kun Kuang, Zhaoyang You, Tao Shen, Jun Xiao, Yin Zhang, Chao Wu, Yueting Zhuang, and Xiaolin Li. Federated Unsupervised Representation Learning. arXiv preprint arXiv:2010.08982, 2020.
- [50] Weiming Zhuang, Yonggang Wen, and Shuai Zhang. Divergence-aware Federated Self-supervised Learning. International Conference on Learning Representations (ICLR), 2022.
- [51] Sungwon Han, Sungwon Park, Fangzhao Wu, Sundong Kim, Chuhan Wu, Xing Xie, and Meeyoung Cha. FedX: Unsupervised Federated Learning with Cross Knowledge Distillation. European Conference on Computer Vision (ECCV), 691–707, 2022.
- [52] Chaoyang He, Zhengyu Yang, Erum Mushtaq, Sunwoo Lee, Mahdi Soltanolkotabi, and Salman Avestimehr. SSFL: Tackling Label Deficiency in Federated Learning via Personalized Self-supervision. arXiv preprint arXiv:2110.02470, 2021.
- [53] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards Personalized Federated Learning. IEEE Transactions on Neural Networks and Learning Systems, Early Access, 2022.
- [54] Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and Robust Federated Learning through Personalization. International Conference on Machine Learning (ICML), 6357–6368, 2021.
- [55] Xingchao Peng, Zijun Huang, Yizhe Zhu, and Kate Saenko. Federated Adversarial Domain Adaptation. International Conference on Learning Representations (ICLR), 2020.
- [56] Chun-Han Yao, Boqing Gong, Hang Qi, Yin Cui, Yukun Zhu, and Ming-Hsuan Yang. Federated Multi-Target Domain Adaptation. IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), 1424–1433, 2022.
- [57] Junyuan Hong, Zhuangdi Zhu, Shuyang Yu, Zhangyang Wang, Hiroko H. Dodge, and Jiayu Zhou. Federated Adversarial Debiasing for Fair and Transferable Representations. ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD), 617–627, 2021.
- [58] Shaoduo Gan, Akhil Mathur, Anton Isopoussu, Fahim Kawsar, Nadia Berthouze, and Nicholas D. Lane. FRuDA: Framework for Distributed Adversarial Domain Adaptation. IEEE Transactions on Parallel and Distributed Systems (TPDS), 33.11 (2022), 3153–3164.
- [59] Dashan Gao, Xin Yao, and Qiang Yang. A Survey on Heterogeneous Federated Transfer Learning. arXiv preprint arXiv:2210.04505, 2022.

- [60] Hua Kang, Zhiyang Li, and Qian Zhang. Communicational and Computational Efficient Federated Domain Adaptation. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 33.12 (2022), 3678–3689.
- [61] Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum Classifier Discrepancy for Unsupervised Domain Adaptation. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3723–3732, 2018.
- [62] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. MixUp: Beyond Empirical Risk Minimization. *International Conference on Learning Representations (ICLR)*, 2017.
- [63] Robert Tibshirani, Guenther Walther, and Trevor Hastie. Estimating the Number of Clusters in a Data Set via the Gap Statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63.2 (2001), 411–423.
- [64] Yoshinori Aono, Takuya Hayashi, Lihua Wang, Shiho Moriai, and Others. Privacy-preserving Deep Learning via Additively Homomorphic Encryption. *IEEE Transactions on Information Forensics and Security*, 13.5 (2017), 1333–1345.
- [65] Fan Lai, Yinwei Dai, Sanjay Singapuram, Jiachen Liu, Xiangfeng Zhu, Harsha Madhyastha, and Mosharaf Chowdhury. FedScale: Benchmarking Model and System Performance of Federated Learning at Scale. *International Conference on Machine Learning*, 11814–11827, 2022.
- [66] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H. Brendan McMahan, Virginia Smith, and Ameet Talwalkar. LEAF: A Benchmark for Federated Settings *arXiv preprint arXiv:1812.01097*, 2018.
- [67] Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, John C. Duchi, and Percy S. Liang. Unlabeled Data Improves Adversarial Robustness. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [68] Zhun Deng, Linjun Zhang, Amirata Ghorbani, and James Zou. Improving Adversarial Robustness via Unlabeled Out-of-domain Data. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2845–2853, 2021.
- [69] Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. Adversarial Examples in the Physical World. *Artificial Intelligence Safety and Security*, 99–112, 2018.
- [70] Jean-Baptiste Alayrac, Jonathan Uesato, Po-Sen Huang, Alhussein Fawzi, Robert Stanforth, and Pushmeet Kohli. Are Labels Required for Improving Adversarial Robustness? *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [71] Nicholas Papernot, Martin Abadi, Ulfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised Knowledge Transfer for Deep Learning from Private Training Data. *International Conference on Learning Representations (ICLR)*, 2017.
- [72] Cynthia Dwork and Aaron Roth. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends® in Theoretical Computer Science*, 9.3–4 (2014), 211–407.
- [73] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep Learning with Differential Privacy. *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 308–318, 2016.
- [74] Hsin-Ying Lee, Hung-Yu Tseng, Jia-Bin Huang, Maneesh Singh, and Ming-Hsuan Yang. Diverse Image-to-image Translation via Disentangled Representations. *European Conference on Computer Vision (ECCV)*, 35–51, 2018.

- [75] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. Advances in Neural Information Processing Systems (NeurIPS), 2016.
- [76] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. Beta-VAE: Learning Basic Visual Concepts with a Constrained Variational Framework. International Conference on Learning Representations (ICLR), 2017.
- [77] Quande Liu, Cheng Chen, Jing Qin, Qi Dou, and Pheng-Ann Heng. FedDG: Federated Domain Generalization on Medical Image Segmentation via Episodic Learning in Continuous Frequency Space. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 1013–1023, 2021.
- [78] A. Tuan Nguyen, Philip Torr, and Ser-Nam Lim. FedSR: A Simple and Effective Domain Generalization Method for Federated Learning. Advances in Neural Information Processing Systems (NeurIPS), 2022.
- [79] Dong-Hyun Lee. Pseudo-Label: The Simple and Efficient Semi-supervised Learning Method for Deep Neural Networks. Workshop on Challenges in Representation Learning, ICML, 2013.
- [80] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I. Jordan. Conditional Adversarial Domain Adaptation. Advances in Neural Information Processing System (NIPS), 2018.
- [81] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks. International Conference on Learning Representations (ICLR), 2018.
- [82] Xin He, Kaiyong Zhao, Xiaowen Chu. AutoML: A Survey of the State-of-the-art. Knowledge-Based Systems, 212 (2021), 106622.
- [83] Michael T. Rosenstein, Zvika Marx, Leslie Pack Kaelbling, and Thomas G. Dietterich. To Transfer or Not to Transfer. NIPS 2005 Workshop on Transfer Learning, 2005.
- [84] Siwei Feng, Boyang Li, Han Yu, Yang Liu, and Qiang Yang. Semi-Supervised Federated Heterogeneous Transfer Learning. Knowledge-Based Systems, 252 (2022), 109384.
- [85] Yang Liu, Yan Kang, Chaoping Xing, Tianjian Chen, and Qiang Yang. A Secure Federated Transfer Learning Framework. IEEE Intelligent Systems, 35.4 (2020), 70–82.
- [86] Yan Kang, Yuanqin He, Jiahuan Luo, Tao Fan, Yang Liu, and Qiang Yang. Privacy-Preserving Federated Adversarial Domain Adaptation over Feature Groups for Interpretability. IEEE Transactions on Big Data, 2022.
- [87] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. Split Learning for Health: Distributed Deep Learning Without Sharing Raw Patient Data. arXiv preprint arXiv:1812.00564, 2018.
- [88] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated Learning: Challenges, Methods, and Future Directions. IEEE Signal Processing Magazine, 37.3 (2020), 50–60.
- [89] Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H Brendan McMahan, Maruan Al-Shedivat, and Others. A Field Guide to Federated Optimization. arXiv preprint arXiv:2107.06917, 2021.
- [90] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated Learning on Non-iid Data Silos: An Experimental Study. IEEE International Conference on Data Engineering (ICDE), 965–978, 2022.



- [91] Hangyu Zhu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. Federated Learning on Non-iid Data: A Survey. Neurocomputing, 465 (2021), 371–390.
- [92] Jesper E. Van Engelen and Holger H. Hoos. A Survey on Semi-supervised Learning. Machine Learning, 109.2 (2020), 373–440.
- [93] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A Comprehensive Survey on Transfer Learning. Proceedings of the IEEE, 109.1 (2020), 43–76.