A Systematic Survey on Federated Semi-supervised Learning

Zixing Song[1], **Xiangli Yang**[2], **Yifei Zhang**[1], **Xinyu Fu**[1], **Zenglin Xu**[3,4] and **Irwin King**[1]
[1]The Chinese University of Hong Kong, Hong Kong SAR, China
[2]University of Electronic Science and Technology of China, Chengdu, China
[3]Harbin Institute of Technology, Shenzhen, China
[4]Peng Cheng Laboratory, Shenzhen, China
{zxsong, yfzhang, xyfu, king}@cse.cuhk.edu.hk, xlyang@std.uestc.edu.cn, xuzenglin@hit.edu.cn

Abstract

Federated learning (FL) revolutionizes distributed machine learning by enabling devices to collaboratively learn a model while maintaining data privacy. However, FL usually faces a critical challenge with limited labeled data, making semi-supervised learning (SSL) crucial for utilizing abundant unlabeled data. The integration of SSL within the federated framework gives rise to federated semi-supervised learning (FSSL), a novel approach that exploits unlabeled data across devices without compromising privacy. This paper systematically explores FSSL, shedding light on its four basic problem settings that commonly appear in real-world scenarios. By examining the unique challenges, generic solutions, and representative methods tailored for each setting of FSSL, we aim to provide a cohesive overview of the current state of the art and pave the way for future research directions in this promising field.

1 Introduction

Federated learning (FL) [Li *et al.*, 2023b] is an advanced framework that enables multiple clients to collaboratively train a shared global model with help of a server while keeping their data local, thus prioritizing data privacy and security [Yin *et al.*, 2022; Zhang *et al.*, 2023d]. This distributed learning paradigm is increasingly recognized for its capability to mitigate challenges across various sectors, including edge computing [Nguyen *et al.*, 2021] and healthcare [Nguyen *et al.*, 2023], where data privacy is paramount. However, traditional FL approaches often assume the availability of fully labeled datasets, an assumption that rarely holds in practice. Given real-world constraints such as lack of time, expertise, or motivation, users often leave large portions of data unlabeled [Jin *et al.*, 2023].

Addressing these practical challenges necessitates a shift towards federated semi-supervised learning (FSSL). Semi-supervised learning (SSL), successful in centralized settings, employs both labeled and unlabeled data by exploiting the latent information within the latter to enhance model performance [Yang *et al.*, 2023]. Transferring these SSL strategies to the federated paradigm yields FSSL, which aims to leverage both labeled and unlabeled data in a federated context. FSSL can thus enhance model efficacy in real-world applications where labeled data is scarce, maximizing the use of abundant yet underexploited unlabeled data in federated learning environments.

Integrating SSL into the FL framework introduces multifaceted challenges. First, FL's decentralized architecture, which safeguards data privacy by keeping data on multiple local clients, impedes the direct application of conventional SSL methods due to restricted access to global data. Second, the non-IID (independently and identically distributed) nature of FL data complicates the effective use of SSL, as biased data distributions can negatively affect label predictions. Third, the distribution of labeled and unlabeled data across the server and clients introduces varying scenarios, each necessitating tailored algorithmic approaches. Different configurations of labeled/unlabeled data locations demand distinct algorithmic solutions to address the specific challenges. Thus, the incorporation of SSL into FL motivates the development of novel strategies to overcome these obstacles.

Despite the accelerated progress recently observed in FSSL, there is still a noticeable absence of a comprehensive survey. Therefore, to elucidate the intersection of these two impactful paradigms, we conduct a systematic examination of FSSL methodologies, aiming to integrate the privacy-preserving features of FL with the efficiency of SSL in utilizing unlabeled data. First, we outline several common and basic problem settings in FSSL with clear illustrations. Second, we delve into each scenario, analyzing its unique challenges, presenting a general solution framework, and reviewing some representative works published in prestigious journals and conferences. Third, we identify existing research gaps and propose potential trajectories for future exploration in the promising domain of FSSL. We recognize that a recent survey [Jin *et al.*, 2023] investigates FL from a broader perspective of label efficient learning. However, our work stands apart by narrowly concentrating on FL specifically within the SSL framework, offering a more detailed analysis.

2 Preliminaries

2.1 Federated Learning

Federated learning is designed to train machine learning models using data distributed across various participants, with the primary goal of safeguarding data privacy. In a classic federated learning setting, a shared global model is trained across multiple devices, often referred to as *clients*, ensuring the privacy of their local data. Instead of sending data to a centralized *server*, the learning happens directly at the client level. Depending on how the data is split across participants, FL can be categorized into horizontal federated learning (HFL) and vertical federated learning (VFL). In HFL, all clients possess data with identical feature spaces, such as images from different users. In VFL, however, clients hold data pertaining to the same users but with different feature spaces, exemplified by a bank and a retail company holding transaction records and purchase history of the same users, respectively. In this paper, following the majority of existing research efforts, we primarily focus on HFL.

Formally, we assume a traditional FL scenario with N clients, and each client i has its own local fully-labeled dataset $\mathcal{D}_i = \{(x_i^j, y_i^j)\}_{j=1}^{n_i}$ with $n_i = |\mathcal{D}_i|$. Then, the global training objective is defined as follows,

$$\min_{\theta} \sum_{i=1}^{N} w_i \ell_i(\mathcal{D}_i; \theta). \tag{1}$$

Here, local loss $\ell_i(\mathcal{D}_i; \theta)$ represents how well the model fits the local data $\mathcal{D}_i$, and the global loss is a weighted average of all local losses. w_i is the weight assigned to the i-th client, which can be set as $w_i = \frac{n_i}{\sum_{i=1}^{N} n_i}$ in FedAvg [McMahan *et al.*, 2017]. It ensures clients with more samples have a proportionately greater influence on the global model.

The training procedure of FL typically involves multiple communication rounds, each comprising two key phases: local training and server aggregation. In the local training phase, a subset of participants is chosen to receive the current global model, which they then refine using their own data. In the following server aggregation phase, these participants transmit their updated model parameters back to the server. The server aggregates these parameters and obtains an updated global model for the next communication round.

2.2 Semi-supervised Learning

Semi-supervised learning (SSL) leverages both a small amount of labeled data and a large amount of unlabeled data during training. It is an effective way to improve learning accuracy without the need for fully labeled datasets. Given the training dataset $\mathcal{D} = \mathcal{L} \cup \mathcal{U}$ with limited labeled samples $\mathcal{L} = \{(x_i, y_i)\}_{i=1}^{n}$ and abundant unlabeled samples $\mathcal{U} = \{x_i\}_{i=1}^{m}$ ($m \gg n$). Then the overall training objective is defined as follows,

$$\min_{\theta} \sum_{(x_i, y_i) \in \mathcal{L}} \ell_s(x_i, y_i; \theta) + \lambda \sum_{x_i \in \mathcal{D}} \ell_u(x_i; \theta). \tag{2}$$

Here, ℓ_s denotes the per-example supervised loss, e.g., cross-entropy for classification, ℓ_u denotes the per-example unsupervised loss, and λ is the trade-off parameter.

Several SSL methods are widely adopted in the literature. Pseudo-label [Lee and others, 2013] employs the class prediction with the highest probability that surpasses a predefined confidence threshold as the pseudo-label $\hat{y}_i$ for unlabeled samples. These pseudo-labels are subsequently utilized for supervised training on the unlabeled data, formulated as $\ell_u(x_i; \theta) = \ell_l(x_i, \hat{y}_i; \theta)$. Consistency regularization enforces the predictions from the augmented examples and original instances to output the same class label [Xie *et al.*, 2020].

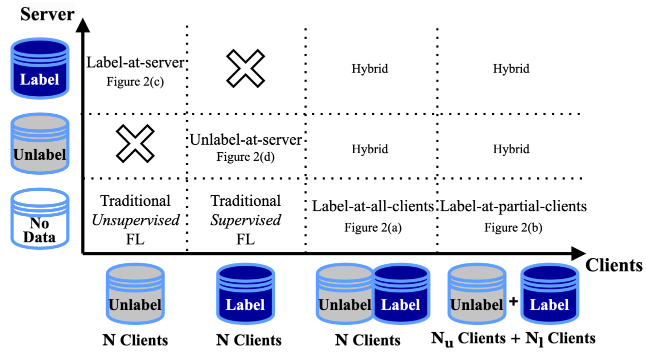

Figure 1: All possible FSSL settings.

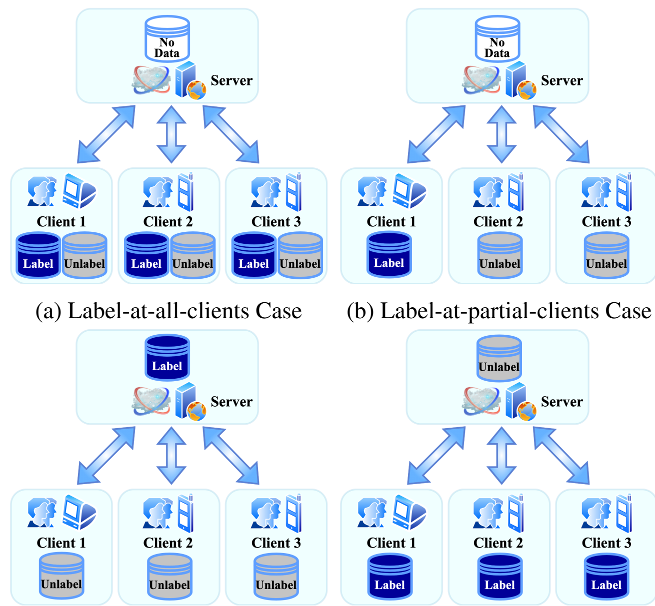

(a) Label-at-all-clients Case (b) Label-at-partial-clients Case

(c) Label-at-server Case (d) Unlabel-at-server Case

Figure 2: Four basic problem settings in FSSL.

3 Problem Settings in FSSL

Incorporating SSL into FL renders the problem more intricate and multifaceted. Figure 1 sketches the various common scenarios that often arise in real-world applications. For simplicity's sake, we make the following assumptions for FSSL.

1. All clients hold data, either labeled, unlabeled, or both.
2. The server can possess data as well.
3. The system as a whole must contain a mix of labeled and unlabeled data, excluding solely one data type.

For the first assumption, clients without data do not actively contribute to the learning process. For the second, the server may hold labeled or unlabeled data, which can improve model performance. We simplify by not considering the server holding both data types, treating it as a simple combination of the two scenarios. The third assumption aligns with the semi-supervised framework. We concentrate on four basic FSSL scenarios based on the locations of labeled and unlabeled data (Figure 2), treating other situations as combinations of these. We consider a system with N clients.

- *Label-at-all-client* All the data are located on the client side $\mathcal{D} = \{\mathcal{D}_i\}_{i=1}^N$ and each client has both labeled and unlabeled data $\mathcal{D}_i = \mathcal{L}_i \cup \mathcal{U}_i$. Here, $\mathcal{L}_i = \{(x_i^j, y_i^j)\}_{j=1}^{n_i}$ with $n_i = |\mathcal{L}_i|$ and $\mathcal{U}_i = \{x_i^k\}_{k=1}^{m_i}$ with $m_i = |\mathcal{U}_i|$.
- *Label-at-partial-client* All the data are located on the client side, but some clients only have labeled data while the remaining ones have unlabeled data $\mathcal{D} = \{\mathcal{D}_i^l\}_{i=1}^{N_l} \cup \{\mathcal{D}_i^u\}_{i=1}^{N_u}$ with $N_u + N_l = N$. Here, $\mathcal{D}_i^l = \{(x_i^j, y_i^j)\}_{j=1}^{n_i}$ with $n_i = |\mathcal{D}_i^l|$ and $\mathcal{D}_i^u = \{x_i^k\}_{k=1}^{m_i}$ with $m_i = |\mathcal{D}_i^u|$.
- *Label-at-server* All the data are located on both the client and the server sides $\mathcal{D} = \mathcal{D}_s \cup \{\mathcal{D}_i\}_{i=1}^N$. The server has the labeled data only $\mathcal{D}_s = \{(x_s^j, y_s^j)\}_{j=1}^{n_s}$ with $n_s = |\mathcal{D}_s|$ while each client has the unlabeled data only $\mathcal{D}_i = \{x_i^k\}_{k=1}^{m_i}$ with $m_i = |\mathcal{D}_i|$.
- *Unlabel-at-server* All the data are located on both the client and the server sides $\mathcal{D} = \mathcal{D}_s \cup \{\mathcal{D}_i\}_{i=1}^N$. The server has the unlabeled data only $\mathcal{D}_s = \{x_s^k\}_{k=1}^{m_s}$ with $m_s = |\mathcal{D}_s|$ while each client has the labeled data only $\mathcal{D}_i = \{(x_i^j, y_i^j)\}_{j=1}^{n_i}$ with $n_i = |\mathcal{D}_i|$.

4 Approaches to FSSL

4.1 Label-at-all-client Case

Unique Challenges and Generic Solution Framework

This case is the most prevalent scenario, and it incorporates the semi-supervised learning setting into the traditional federated learning system in a most straightforward manner. Label-at-all-client case is widely applicable in real-world situations. For example, consider a company aiming to develop an object detection FL model using smartphone-captured images. The company has no access to the local data of users, and only users can annotate their images. However, it is common for users to refrain from labeling every image, resulting in a scenario where all clients possess partially labeled data, epitomizing the "label-at-all-client" setting.

A naive solution to this case involves applying existing off-the-shelf semi-supervised learning techniques at each client, coupled with the use of federated learning algorithms to aggregate the locally trained weights back to the server. The detailed training procedure is shown in Algorithm 1 and the training objective can be formulated as

$$\min_\theta \sum_{i=1}^N w_i \underbrace{\left(\sum_{(x_i,y_i)\in\mathcal{L}_i} \ell_s(x_i, y_i; \theta) + \lambda \sum_{x_i \in \mathcal{D}_i} \ell_u(x_i; \theta) \right)}_{\ell_i(\mathcal{D}_i;\theta)}.$$

However, this approach does not fully exploit the knowledge of the multiple models trained on heterogeneous data

Algorithm 1 Generic label-at-all-client FSSL framework

Input: Initialized model parameters $\theta = \theta^1$
1: **for** communication round $t = 1$ **to** R **do**
2: Randomly select a subset of clients $\mathcal{C}^t \subseteq \{1, \ldots, N\}$.
3: **for** each selected client $i \in \mathcal{C}^t$ in parallel **do**
4: Train a local model θ_i^t with $\mathcal{D}_i$ and θ^t using **SSL**.
5: **end for**
6: Server aggregates weights $\{\theta_i^t\}_{i\in\mathcal{C}^t}$ and updates θ^{t+1}.
7: **end for**

distributions. The client heterogeneity makes the simple combination of SSL and FL algorithms suffer severely from slow convergence and performance degradation. This heterogeneity manifests in two primary forms: statistical and system. The former refers to the prevalence of non-IID data, which undermines algorithmic convergence as many FL algorithms (e.g. FedAvg) are inefficient in dealing with this issue. As for the latter, the clients may have varied system capabilities and computational resources, which can lead to inconsistent client participation and exacerbate the straggler problem, where slower clients delay overall progress, thereby significantly reducing the speed of the learning algorithm.

Representative Methods

FedMatch [Jeong *et al.*, 2021] is the first to improve upon naive combinations of FL and SSL. It incorporates an inter-client consistency loss to encourage the same predictions across clients. Additionally, FedMatch divides the model parameters into two sets: one dedicated to supervised learning and the other to unsupervised learning. This disjoint learning reduces the communication overhead. FedMatch is even designed with flexibility, allowing easy adaptation to the following label-at-server case.

FedAvg-DS [Nandury *et al.*, 2021] builds upon the foundational FedAvg algorithm by introducing a sophisticated aggregation mechanism that accounts for the diversity in updates from clients.This enhancement is particularly effective when it is combined with existing SSL methods, addressing the issue of client heterogeneity.

FedTriNet [Che *et al.*, 2021] first pretrains the model on the labeled data using FedAvg and then utilizes a dynamic quality control mechanism to generate high-quality pseudo labels of the unlabeled data for retraining [Lin *et al.*, 2021].

FedFAPL [Wei and Huang, 2022] derives a fairer federated model across all clients. It strategically balances the involvement of active unlabeled samples (AUS) in the local model training. By setting global numerical restrictions on AUS participation and then allocating these restrictions into personalized local constraints for each client, FedFAPL facilitates more effective local pseudo-labeling.

FedGAN [Zhao *et al.*, 2022] adapts the Triple GAN architecture [Li *et al.*, 2017], originally designed for conventional SSL, to the FL context. It consists of three generators and one discriminator to learn the relationship between labeled and unlabeled data effectively. Importantly, FedGAN dynamically adjusts aggregation weights of local models based on optimization difficulty posed by non-IID data distributions.

F2CMT [Wen *et al.*, 2022] improves over the mean teacher model [Tarvainen and Valpola, 2017], a popular consistency

Representative Methods	SSL Techniques	Data Heterogeneity	Data Privacy	Communication Efficiency	Advantages	Disadvantages
FedMatch [Jeong *et al.*, 2021]	Consistency regularization	✗	✗	Parameter decomposition	First work	Information leakage
FedAvg-DS [Nandury *et al.*, 2021]	✗	Diversity Scaling	✗	✗	FedAvg extension	Not specific for SSL
FedTriNet [Che *et al.*, 2021]	Pseudo-label	✗	✗	✗	High-quality pseudo-label	Extra pre-training overhead
FedFAPL [Wei and Huang, 2022]	Pseudo-label	Balancing #AUS	✗	✗	Fair accuracy parity	Extra communication overhead
FedGAN [Zhao *et al.*, 2022]	SSL GAN	Dynamic aggregation	✗	✗	Robust	Unstable training
F2CMT [Wen *et al.*, 2022]	Consistency regularization	✗	✗	✗	Reliable targets generation	Extra local training overhead
FedCPSL [Wang *et al.*, 2023b]	Pseudo-label	Client variance reduction	✗	Local momentum	Personalization	Information leakage
DS-FL [Itahara *et al.*, 2023]	Consistency regularization	Entropy reduction average	✗	Model-output exchange	Communication-efficient	Extra local training overhead
FedLoKe [Zhang *et al.*, 2023a]	Pseudo-label	Local knowledge enhancement	✗	✗	Overfitting prevention	Extra local training overhead

Table 1: Comparison of different FSSL methods for the **label-at-all-client** case. ✗ denotes that the method does not focus on this issue.

regularization method, with a cross-clients ensemble module and a model-wise self-ensembling module. F2CMT can yield reliable teacher targets to guide the learning process of student models on local clients.

FedCPSL [Wang *et al.*, 2023b; Wang *et al.*, 2023c] focuses on model personalization within the FSSL framework. It introduces novel strategies such as adaptive client variance reduction, local momentum, and normalized global aggregation to combat the challenge of device heterogeneity. These strategies are designed to enhance algorithm convergence, theoretically achieving a sublinear convergence rate.

DS-FL [Itahara *et al.*, 2023] exchanges model outputs instead of model parameters, significantly enhancing communication efficiency, especially as model sizes increase. It also presents a new model-output aggregation method designed to be robust against the heterogeneity in model outputs caused by non-IID data distributions.

FedLoKe [Zhang *et al.*, 2023a] generates more accurate pseudo-labels with local knowledge to boost the performance of the global model in the non-IID scenarios. The local model is also trained on the pseudo-label dataset generated by the global model, transferring general knowledge to the local model and preventing it from overfitting.

4.2 Label-at-partial-client Case

Unique Challenges and Generic Solution Framework

This case can be viewed as an extreme case of the previous label-at-all-client case. In the conventional label-at-all-client setting, each client is homogeneous, containing a mix of both labeled and unlabeled data. In contrast, the label-at-partial-client case now features heterogeneous clients: some possess exclusively labeled data, while others have only unlabeled data. Here, the "semi-supervised" aspect applies at the intra-client level, with each client holding data that is either entirely labeled or completely unlabeled.

The strategies effective in the label-at-all-client situation often falter in the label-at-partial-client case due to the extreme non-IID nature of client data. The division of labeled and unlabeled data across multiple clients prevents the application of standard semi-supervised learning techniques within one single client for local training. However, it remains feasible to engage in supervised learning with labeled clients and unsupervised learning with unlabeled clients. The local models' weights can then be aggregated into the global model, taking into account the specific data type each client contributes. The detailed training procedure is shown in Algorithm 2 and the training objective can be formulated as

Algorithm 2 Generic label-at-partial-client FSSL framework

Input: Initialized model parameters $\theta = \theta^1$
for communication round $t = 1$ **to** R **do**
  Randomly select a subset of clients $\mathcal{C}^t \subseteq \{1, \ldots, N\}$.
  for each selected client $i \in \mathcal{C}^t$ in parallel **do**
    if client i is a labeled client **then**
      Train a local model θ_i^t with $\mathcal{D}_i^l$ and the initialized θ^t.
    end if
    if client i is an unlabeled client **then**
      Train a local model θ_i^t with $\mathcal{D}_i^u$ and the initialized θ^t.
    end if
  end for
  Server aggregates $\{\theta_i^t\}_{i \in \mathcal{C}^t}$ based on client type, gets θ^{t+1}.
end for

$$\min_{\theta} \sum_{i=1}^{N_l} w_i^l \sum_{(x_i, y_i) \in \mathcal{D}_i^l} \ell_s(x_i, y_i; \theta) + \sum_{i=1}^{N_u} w_i^u \sum_{x_i \in \mathcal{D}_i^u} \ell_u(x_i; \theta).$$

Representative Methods

FedIRM [Liu *et al.*, 2021a] enhances the consistency regularization framework by introducing an inter-client relation matching strategy. The proposed learning scheme fosters coherence in learning between labeled and unlabeled clients. It ensures that the unlabeled clients mirror the class relationships observed in labeled clients and preserves the discriminative task knowledge.

Fed-Consist [Yang *et al.*, 2021] also introduces a consistency-based method, in which different augmentations were applied to unlabeled images with their predictions similarity maximized. However, its performance significantly decreases when the proportion of unlabeled clients increases.

Representative Methods	SSL Techniques	Data Heterogeneity	Data Privacy	Communication Efficiency	Advantages	Disadvantages
FedIRM [Liu *et al.*, 2021a]	Consistency regularization	✗	✗	✗	Knowledge preservation	Failure with imbalanced classes
Fed-Consist [Yang *et al.*, 2021]	Consistency regularization	✗	✗	✗	Domain specific for segmentation	Failure with large number of unlabeled clients
RSCFed [Liang *et al.*, 2022]	Consistency regularization	Distance-reweighted model aggregation.	✗	✗	Consideration of uneven reliability in non-IID clients	Extra aggregation overhead
FedSSL-DP [Fan *et al.*, 2022]	SSL GAN	✗	Differential privacy	✗	No data leakage	Unstable training
CBAFed [Li *et al.*, 2023a]	Pseudo-label	Class balanced adaptive thresholds	✗	✗	Effective with imbalanced classes	Extra communication cost
UM-pFSSL [Shi *et al.*, 2023]	Pseudo-label	Uncertainty minimization	✗	Ranking update protocol	Personalization	Extra local training overhead

Table 2: Comparison of different FSSL methods for the **label-at-partial-client** case. ✗ denotes that the method does not focus on this issue.

This is due to the potential for the global model to become overly influenced by the unlabeled data, which can skew learning outcomes and degrade overall model performance.

RSCFed [Liang *et al.*, 2022] addresses the uneven reliability of non-IID local clients by moving beyond straightforward aggregation of local models. It introduces a novel concept of updating the global model through the aggregation of multiple sub-consensus models. This is achieved by randomly sub-sampling clients to form sub-consensus models and employing a distance-reweighted aggregation module to integrate these models during each synchronization round.

FedSSL-DP [Fan *et al.*, 2022] designs a mixed-data generation strategy to utilize both labeled and unlabeled clients by establishing a unified data space without direct data exchange. Furthermore, the differential privacy (DP) scheme can be integrated smoothly into the model, prohibiting excessive access to the labeled data with theoretical guarantees.

CBAFed [Li *et al.*, 2023a] rethinks the standard pseudo-labeling methods used in FSSL. Recognizing the potential imbalance in the training distribution of unlabeled data due to the non-IID issue, CBAFed proposes class-balanced adaptive thresholds. These thresholds are dynamically adjusted based on the empirical distribution of all training data across local clients, as observed in the previous communication round, to accommodate the non-IID data distributions more effectively.

UM-pFSSL [Shi *et al.*, 2023] investigates the personalized FL model under the label-at-partial-client case. Unlabeled clients struggle to obtain competent personalized models due to insufficient knowledge of their local data distributions, while labeled clients may dominate the collaborative training and obtain superior performance. UM-pFSSL enables unlabeled clients to assimilate knowledge from selected "helper" clients, thereby acquiring reliable pseudo-labels guided by an uncertainty metric. To reduce the communication cost, a ranking update protocol is designed to select the suitable helper clients as well.

4.3 Label-at-server Case

Unique Challenges and Generic Solution Framework

Unlike two previous cases where the server acts as the parameter server only, the server now has direct access to its own data. The label-at-server case is quite common in real-world applications. For instance, a healthcare system may involve a central hub ("server") with domain experts and a limited number of labeled data, such as medical records, together with many rural branches ("clients") with large volumes of unlabeled data without expert annotation.

Algorithm 3 Generic label-at-server FSSL framework

Input: Initialized model parameters $\theta = \theta^1$
1: **for** communication round $t = 1$ **to** R **do**
2: Server trains the model θ^t in a supervised manner with $\mathcal{D}_s$.
3: Randomly select a subset of clients $\mathcal{C}^t \subseteq \{1, \ldots, N\}$.
4: **for** each selected client $i \in \mathcal{C}^t$ in parallel **do**
5: Train a local model θ_i^t with $\mathcal{D}_i$ and the initialized θ^t.
6: **end for**
7: Server aggregates $\{\theta_i^t\}_{i\in\mathcal{C}^t}$ and updates θ^{t+1}.
8: **end for**

The main challenge of this case lies in the lack of direct supervision on the client side due to the disjoint distribution of labeled and unlabeled data, and thus, no classic SSL methods can be directly applied to either the client or the server. The clients now must rely on indirect supervision, such as the pseudo-labels generated based on the server's model, to learn from their unlabeled data, making the learning process less efficient and potentially less accurate. Therefore, the model is trained in a supervised manner on the server before it is distributed to the clients, as shown in Algorithm 3. The corresponding training objective is given as

$$\min_{\theta} \sum_{(x_i,y_i)\in\mathcal{D}_s} \ell_s(x_i, y_i; \theta) + \lambda \sum_{i=1}^{N} w_i \sum_{x_i\in\mathcal{D}_i} \ell_u(x_i; \theta).$$

Representative Methods

FedMix [Zhang *et al.*, 2021a] designs parameter decomposition strategies for disjointed learning of labeled and unlabeled data. To alleviate the non-IID issue, FedMix proposes a novel aggregation rule, which dynamically adjusts the weight of each local model based on the client's participation frequency, thus balancing the influence of varying data distributions across clients.

Representative Methods	SSL Techniques	Data Heterogeneity	Data Privacy	Communication Efficiency	Advantages	Disadvantages
FedMix [Zhang *et al.*, 2021a]	Consistency regularization	Frequency-reweighted model aggregation	✗	Parameter decomposition	Robust and stable	Information leakage
GDST [Liu *et al.*, 2021b]	Self-training	✗	✗	✗	Easy to implement	Failure with non-IID setting
CRL-Grouping [Zhang *et al.*, 2021b]	Consistency regularization	Grouping-based model aggregation	✗	✗	Easy to implement	Failure with distribution shift
FedIL [Yang *et al.*, 2022]	Consistency regularization	✗	✗	Group-based asynchronous training	Provable convergence	Not scalable
SSFL [Diao *et al.*, 2022]	Pseudo-label	✗	✗	Alternate training	Communication-efficient	Information leakage
FedFAME [Malaviya *et al.*, 2023]	Contrastive learning	Knowledge distillation	✗	✗	Data augmentation free	Information leakage
pFedKnow [Wang *et al.*, 2023a]	Pseudo-label	Collaborative distillation	✗	Network pruning	Personalization	Extra pre-training overhead

Table 3: Comparison of different FSSL methods for the **label-at-server** case. ✗ denotes that the method does not focus on this issue.

GDST [Liu *et al.*, 2021b] integrates global distillation and self-training into local training within the FL framework, jointly with server-side fine-tuning, to further stabilize and enhance the learning process of the global model.

CRL-Grouping [Zhang *et al.*, 2021b] combines consistency regularization with group normalization (GN), reducing gradient diversity and improving test accuracy. It employs a grouping-based averaging technique to expedite convergence, offering a substantial improvement in speed over FedAvg.

FedIL [Yang *et al.*, 2022] enforces the consistency between the predictions made by clients and the server during the training process, achieving a provable convergence guarantee. FedIL also introduces a group-based asynchronous training algorithm in combination with a time-slot-based task scheduling to allow more clients to participate in training simultaneously.

SSFL [Diao *et al.*, 2022] resorts to an alternate training scheme, which fine-tunes the aggregated global model with labeled data and generates pseudo-labels only once upon receiving the global model from the server. Compared to the existing solutions, which all train and aggregate server and client models in parallel and generate pseudo-labels with the training models for every batch of unlabeled samples, SSFL is much more communication efficient.

FedFAME [Malaviya *et al.*, 2023] offers a versatile framework that eliminates the need for data augmentation in local model training by utilizing contrastive learning. This feature makes FedFAME particularly suitable for domains lacking standard augmentation techniques, like text or graphs, providing a robust solution across diverse data types.

pFedKnow [Wang *et al.*, 2023a] generates lightweight personalized client models via neural network pruning techniques to reduce communication costs. Moreover, it leverages pretrained large models as a form of prior knowledge to guide the aggregation of personalized client models and further enhance the learning efficiency.

4.4 Unlabel-at-server Case

Unique Challenges and Generic Solution Framework

In the unlabel-at-server case, the locations of labeled and unlabeled data are swapped compared to the label-at-server case, with the server now holding unlabeled data while clients possess labeled data. This setup is relevant in scenarios where data privacy is a critical concern. For instance, a research institution ("server") working with local hospitals ("clients") on a new drug development project illustrates this well. The hospitals have access to labeled patient data for individuals diagnosed with rare conditions but are restricted from sharing this sensitive information due to privacy regulations. Conversely, the research institution has a wealth of unlabeled patient data derived from public datasets and published research.

Algorithm 4 Generic unlabel-at-server FSSL framework

Input: Initialized model parameters $\theta = \theta^1$
1: **for** communication round $t = 1$ **to** R **do**
2: Server trains model θ^t on $\mathcal{D}_s$ with pseudo-labels.
3: Randomly select a subset of clients $\mathcal{C}^t \subseteq \{1, \ldots, N\}$.
4: **for** each selected client $i \in \mathcal{C}^t$ in parallel **do**
5: Train a local model θ_i^t with $\mathcal{D}_i$ and the initialized θ^t.
6: **end for**
7: Server aggregates $\{\theta_i^t\}_{i\in\mathcal{C}^t}$, pseudo-labels and updates θ^{t+1}.
8: **end for**

The main challenge is that the server cannot directly evaluate the quality of the local models during the server aggregation since it lacks labeled data. It must employ indirect measures to assess and incorporate client contributions effectively. Hence, we train the model with the help of pseudo-labels from the selected local models. The training procedure is shown in Algorithm 4 with the training objective given as

$$\min_{\theta} \sum_{x_i \in \mathcal{D}_s} \ell_u(x_i; \theta) + \lambda \sum_{i=1}^{N} w_i \sum_{(x_i, y_i) \in \mathcal{D}_i} \ell_s(x_i, y_i; \theta).$$

Representative Methods

PATE-G [Papernot *et al.*, 2017] collects local models from clients to serve as teacher models, which then generate pseudo-labels for the server's unlabeled data based on their collective predictions. These pseudo-labeled datasets are leveraged to refine the global model. To ensure the privacy of this knowledge transfer process, PATE-G utilizes the moments accountant technique, which facilitates the training of student models under stringent privacy constraints, providing meaningful privacy guarantees for the data involved.

Representative Methods	SSL Techniques	Data Heterogeneity	Data Privacy	Communication Efficiency	Advantages	Disadvantages
PATE-G [Papernot *et al.*, 2017]	SSL GAN	✗	Differential privacy	✗	Privacy preservation	Failure with non-IID
Ada-FedSemi [Wang *et al.*, 2022]	Pseudo-label	✗	✗	Online learning	Communication-efficient	Failure with heterogeneous clients
HeteroAda-FedSemi [Xu *et al.*, 2023]	Pseudo-label	Layer-wise global aggregation	✗	Online learning	Effective with heterogeneous clients	Extra local training overhead

Table 4: Comparison of different FSSL methods for the **unlabel-at-server** case. ✗ denotes that the method does not focus on this issue.

Ada-FedSemi [Wang *et al.*, 2022] selects partial local models to produce pseudo-labels for the unlabeled data. A multi-armed bandit (MAB) based online learning algorithm is introduced to adaptively determine the participating fraction and confidence threshold during training. **HeteroAda-FedSemi** [Xu *et al.*, 2023] extends Ada-FedSemi for the heterogeneous clients with diverse computation and communication resources. It customizes local models derived from the same global model to match the capabilities of individual clients by adjusting the model depths.

5 Future Directions

FSSL is an emerging research topic. Although significant progresses have been made for FSSL, there still remain plenty of research directions worthy of future explorations.

- **FSSL for Vertically Partitioned Data**: Investigating FSSL for vertically partitioned data presents an intriguing frontier with unique challenges. This new VFL setting, where different clients hold different features or the labels for the overlapping samples, requires the clients to communicate with the server for each iteration of local training (rather than after several epochs under the horizontal FSSL setting), introducing extremely high communication costs [Sun *et al.*, 2023]. Therefore, the key research question revolves around effectively and efficiently leveraging the limited supervision information across diverse feature spaces while solving extra issues related to data alignment, feature heterogeneity, and privacy preservation across vertical partitions [Kang *et al.*, 2022]. An in-depth exploration of vertical FSSL enables collaborative learning of feature-partitioned partially-labeled data distributed across multiple institutions in industrial applications.
- **FSSL across Diverse Data Modalities**: The current exploration of FSSL has predominantly focused on image data, yet the potential of FSSL extends far beyond to include a variety of data modalities encountered in real-world scenarios, such as text, graphs, and time series. Specifically, the complex non-Euclidean structure of graph data [Ma *et al.*, 2023; Song *et al.*, 2023c; Song *et al.*, 2023a; Song *et al.*, 2023b; Zhang *et al.*, 2023c; Song *et al.*, 2024], necessitates customized designs for effective SSL within an FL framework at both the node level [Yao *et al.*, 2023] and graph level [Tao *et al.*, 2022]. The adaptation of FSSL to these diverse data types or even environments with multi-modal clients is a critical area for future research. The main challenge lies in handling the heterogeneity inherent to different data types.
- **FSSL with Enhanced Model Robustness**: There has been an emerging concern about the robustness of FSSL, especially in some safety-critical healthcare applications like medical diagnosis. The paramount goals of enhancing robustness against noisy labels and adversarial attacks hold significant implications. Noisy labels can severely degrade FSSL model performance by introducing inaccuracies in training data [Kim *et al.*, 2022], and adversarial attacks pose a substantial risk, potentially manipulating FSSL model predictions [Fu *et al.*, 2022; Liu *et al.*, 2022]. Therefore, addressing these open problems is not merely a technical endeavor but a means to safeguard the trustworthiness of FSSL applications in sensitive sectors where the stakes involve human lives.
- **FSSL under Additional Learning Settings**: Recent studies have shifted from basic FSSL investigation to exploring advanced FSSL with additional learning settings' constraints. FedPU [Lin *et al.*, 2022] focuses on positive and unlabeled learning under the FSSL framework, where the labeled data is only of the positive class. FedoSSL [Zhang *et al.*, 2023b] extends open-world SSL to an FL context with unseen classes in test data. These developments push the boundaries of FSSL and highlight the need for practical FSSL models that can operate effectively under extra learning conditions encountered in real-world environments.

6 Conclusion

In conclusion, this survey underscores the burgeoning field of FSSL, highlighting its potential to synergize the privacy-preserving aspects of FL with the data efficiency of SSL. Despite existing progress, considerable obstacles remain, uncovering a wealth of opportunities for crafting diverse FSSL models for practical applications.

Acknowledgments

The work described in this paper was partially supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (RGC GRF 2151185).

References

[Che *et al.*, 2021] Liwei Che, Zewei Long, Jiaqi Wang, Yaqing Wang, Houping Xiao, and Fenglong Ma. Fedtrinet: A pseudo labeling method with three players for federated semi-supervised learning. In *IEEE BigData*, 2021.

[Diao *et al.*, 2022] Enmao Diao, Jie Ding, and Vahid Tarokh. Semifl: Semi-supervised federated learning for unlabeled clients with alternate training. In *NeurIPS*, 2022.

[Fan *et al.*, 2022] Chenyou Fan, Junjie Hu, and Jianwei Huang. Private semi-supervised federated learning. In *IJCAI*, pages 2009–2015. ijcai.org, 2022.

[Fu *et al.*, 2022] Chong Fu, Xuhong Zhang, Shouling Ji, Jinyin Chen, Jingzheng Wu, Shanqing Guo, Jun Zhou, Alex X. Liu, and Ting Wang. Label inference attacks against vertical federated learning. In *USENIX Security Symposium*, pages 1397–1414, 2022.

[Itahara *et al.*, 2023] 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 TMC*, 2023.

[Jeong *et al.*, 2021] Wonyong Jeong, Jaehong Yoon, Eunho Yang, and Sung Ju Hwang. Federated semi-supervised learning with inter-client consistency & disjoint learning. In *ICLR*, 2021.

[Jin *et al.*, 2023] Yilun Jin, Yang Liu, Kai Chen, and Qiang Tang. Federated learning without full labels: A survey. *IEEE Data Eng. Bull.*, 46(1):27–51, 2023.

[Kang *et al.*, 2022] Yan Kang, Yang Liu, and Xinle Liang. Fedcvt: Semi-supervised vertical federated learning with cross-view training. *ACM Trans. Intell. Syst. Technol.*, 13(4):64:1–64:16, 2022.

[Kim *et al.*, 2022] Sangmook Kim, Wonyoung Shin, Soohyuk Jang, Hwanjun Song, and Se-Young Yun. Fedrn: Exploiting k-reliable neighbors towards robust federated learning. In *CIKM*, pages 972–981. ACM, 2022.

[Lee and others, 2013] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, 2013.

[Li *et al.*, 2017] Chongxuan Li, Taufik Xu, Jun Zhu, and Bo Zhang. Triple generative adversarial nets. In *NIPS*, pages 4088–4098, 2017.

[Li *et al.*, 2023a] Ming Li, Qingli Li, and Yan Wang. Class balanced adaptive pseudo labeling for federated semi-supervised learning. In *CVPR*, pages 16292–16301, 2023.

[Li *et al.*, 2023b] Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, Yuan Li, Xu Liu, and Bingsheng He. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Trans. Knowl. Data Eng.*, 35(4):3347–3366, 2023.

[Liang *et al.*, 2022] Xiaoxiao Liang, Yiqun Lin, Huazhu Fu, Lei Zhu, and Xiaomeng Li. Rscfed: Random sampling consensus federated semi-supervised learning. In *CVPR*, pages 10144–10153. IEEE, 2022.

[Lin *et al.*, 2021] Haowen Lin, Jian Lou, Li Xiong, and Cyrus Shahabi. Semifed: Semi-supervised federated learning with consistency and pseudo-labeling. *CoRR*, abs/2108.09412, 2021.

[Lin *et al.*, 2022] Xinyang Lin, Hanting Chen, Yixing Xu, Chao Xu, Xiaolin Gui, Yiping Deng, and Yunhe Wang. Federated learning with positive and unlabeled data. In *ICML*. PMLR, 2022.

[Liu *et al.*, 2021a] Quande Liu, Hongzheng Yang, Qi Dou, and Pheng-Ann Heng. Federated semi-supervised medical image classification via inter-client relation matching. In *MICCAI*, 2021.

[Liu *et al.*, 2021b] Xinyi Liu, Linghui Zhu, Shu-Tao Xia, Yong Jiang, and Xue Yang. GDST: global distillation self-training for semi-supervised federated learning. In *GLOBECOM*, pages 1–6. IEEE, 2021.

[Liu *et al.*, 2022] Yi Liu, Xingliang Yuan, Ruihui Zhao, Cong Wang, Dusit Niyato, and Yefeng Zheng. Poisoning semi-supervised federated learning via unlabeled data: Attacks and defenses. 2022.

[Ma *et al.*, 2023] Yueen Ma, Zixing Song, Xuming Hu, Jingjing Li, Yifei Zhang, and Irwin King. Graph component contrastive learning for concept relatedness estimation. In *AAAI*, pages 13362–13370. AAAI Press, 2023.

[Malaviya *et al.*, 2023] Shubham Malaviya, Manish Shukla, Pratik Korat, and Sachin Lodha. Fedfame: A data augmentation free framework based on model contrastive learning for federated semi-supervised learning. In *SAC*, pages 1114–1121. ACM, 2023.

[McMahan *et al.*, 2017] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, volume 54, 2017.

[Nandury *et al.*, 2021] Kishore Nandury, Anand Mohan, and Frederick Weber. Cross-silo federated training in the cloud with diversity scaling and semi-supervised learning. In *ICASSP*, pages 3085–3089. IEEE, 2021.

[Nguyen *et al.*, 2021] Dinh C. Nguyen, Ming Ding, Pubudu N. Pathirana, Aruna Seneviratne, Jun Li, and H. Vincent Poor. Federated learning for internet of things: A comprehensive survey. *IEEE Commun. Surv. Tutorials*, 23(3):1622–1658, 2021.

[Nguyen *et al.*, 2023] Dinh C. Nguyen, Quoc-Viet Pham, Pubudu N. Pathirana, Ming Ding, Aruna Seneviratne, Zihuai Lin, Octavia A. Dobre, and Won-Joo Hwang. Federated learning for smart healthcare: A survey. *ACM Comput. Surv.*, 55(3):60:1–60:37, 2023.

[Papernot *et al.*, 2017] Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian J. Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. In *ICLR*, 2017.

[Shi *et al.*, 2023] Yanhang Shi, Siguang Chen, and Haijun Zhang. Uncertainty minimization for personalized federated semi-supervised learning. *IEEE Trans. Netw. Sci. Eng.*, 10(2):1060–1073, 2023.

[Song *et al.*, 2023a] Zixing Song, Yifei Zhang, and Irwin King. No change, no gain: Empowering graph neural networks with expected model change maximization for active learning. In *NeurIPS*, 2023.

[Song *et al.*, 2023b] Zixing Song, Yifei Zhang, and Irwin King. Optimal block-wise asymmetric graph construction for graph-based semi-supervised learning. In *NeurIPS*, 2023.

[Song *et al.*, 2023c] Zixing Song, Yuji Zhang, and Irwin King. Towards fair financial services for all: A temporal GNN approach for individual fairness on transaction networks. In *CIKM*, pages 2331–2341. ACM, 2023.

[Song *et al.*, 2024] Zixing Song, Ziqiao Meng, and Irwin King. A diffusion-based pre-training framework for crystal property prediction. In *AAAI*, pages 8993–9001. AAAI Press, 2024.

[Sun *et al.*, 2023] Jingwei Sun, Ziyue Xu, Dong Yang, Vishwesh Nath, Wenqi Li, Can Zhao, Daguang Xu, Yiran Chen, and Holger R. Roth. Communication-efficient vertical federated learning with limited overlapping samples. In *ICCV*, pages 5180–5189. IEEE, 2023.

[Tao *et al.*, 2022] Ye Tao, Ying Li, and Zhonghai Wu. Semigraphfl: Semi-supervised graph federated learning for graph classification. In *PPSN (1)*, volume 13398 of *Lecture Notes in Computer Science*, pages 474–487, 2022.

[Tarvainen and Valpola, 2017] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NIPS*, pages 1195–1204, 2017.

[Wang *et al.*, 2022] Lun Wang, Yang Xu, Hongli Xu, Jianchun Liu, Zhiyuan Wang, and Liusheng Huang. Enhancing federated learning with in-cloud unlabeled data. In *ICDE*, pages 136–149. IEEE, 2022.

[Wang *et al.*, 2023a] Jiaqi Wang, Shenglai Zeng, Zewei Long, Yaqing Wang, Houping Xiao, and Fenglong Ma. Knowledge-enhanced semi-supervised federated learning for aggregating heterogeneous lightweight clients in iot. In *SDM*, pages 496–504. SIAM, 2023.

[Wang *et al.*, 2023b] Shuai Wang, Yanqing Xu, Yanli Yuan, and Tony Q. S. Quek. Towards fast personalized semi-supervised federated learning in edge networks: Algorithm design and theoretical guarantee. *IEEE Transactions on Wireless Communications*, pages 1–1, 2023.

[Wang *et al.*, 2023c] Shuai Wang, Yanqing Xu, Yanli Yuan, Xiuhua Wang, and Tony Q. S. Quek. Boosting semi-supervised federated learning with model personalization and client-variance-reduction. In *ICASSP*, 2023.

[Wei and Huang, 2022] Xiao-Xiang Wei and Hua Huang. Balanced federated semisupervised learning with fairness-aware pseudo-labeling. *IEEE TNNLS*, 2022.

[Wen *et al.*, 2022] Tingjie Wen, Shengjie Zhao, and Rongqing Zhang. Federated semi-supervised learning through a combination of self and cross model ensembling. In *IJCNN*, pages 1–8. IEEE, 2022.

[Xie *et al.*, 2020] Qizhe Xie, Zihang Dai, Eduard H. Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation for consistency training. In *NeurIPS*, 2020.

[Xu *et al.*, 2023] Yang Xu, Lun Wang, Hongli Xu, Jianchun Liu, Zhiyuan Wang, and Liusheng Huang. Enhancing federated learning with server-side unlabeled data by adaptive client and data selection. *IEEE TMC*, pages 1–18, 2023.

[Yang *et al.*, 2021] Dong Yang, Ziyue Xu, Wenqi Li, Andriy Myronenko, Holger R. Roth, Stephanie A. Harmon, Sheng Xu, Baris Turkbey, Evrim Turkbey, Xiaosong Wang, Wentao Zhu, Gianpaolo Carrafiello, Francesca Patella, Maurizio Cariati, Hirofumi Obinata, Hitoshi Mori, Kaku Tamura, Peng An, Bradford J. Wood, and Daguang Xu. Federated semi-supervised learning for COVID region segmentation in chest CT using multi-national data from china, italy, japan. *Medical Image Anal.*, 2021.

[Yang *et al.*, 2022] Nan Yang, Dong Yuan, Yuning Zhang, Yongkun Deng, and Wei Bao. Asynchronous semi-supervised federated learning with provable convergence in edge computing. *IEEE Netw.*, 36(5):136–143, 2022.

[Yang *et al.*, 2023] Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. A survey on deep semi-supervised learning. *IEEE TKDE*, 35(9):8934–8954, 2023.

[Yao *et al.*, 2023] Yuhang Yao, Weizhao Jin, Srivatsan Ravi, and Carlee Joe-Wong. FedGCN: Convergence-communication tradeoffs in federated training of graph convolutional networks. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.

[Yin *et al.*, 2022] Xuefei Yin, Yanming Zhu, and Jiankun Hu. A comprehensive survey of privacy-preserving federated learning: A taxonomy, review, and future directions. *ACM Comput. Surv.*, 54(6):131:1–131:36, 2022.

[Zhang *et al.*, 2021a] Zhe Zhang, Shiyao Ma, Jiangtian Nie, Yi Wu, Qiang Yan, Xiaoke Xu, and Dusit Niyato. Semi-supervised federated learning with non-iid data: Algorithm and system design. In *HPCC/DSS/SmartCity/DependSys*. IEEE, 2021.

[Zhang *et al.*, 2021b] Zhengming Zhang, Yaoqing Yang, Zhewei Yao, Yujun Yan, Joseph E. Gonzalez, Kannan Ramchandran, and Michael W. Mahoney. Improving semi-supervised federated learning by reducing the gradient diversity of models. In *IEEE BigData*, 2021.

[Zhang *et al.*, 2023a] Chao Zhang, Fangzhao Wu, Jingwei Yi, Derong Xu, Yang Yu, Jindong Wang, Yidong Wang, Tong Xu, Xing Xie, and Enhong Chen. Non-iid always bad? semi-supervised heterogeneous federated learning with local knowledge enhancement. In *CIKM*, 2023.

[Zhang *et al.*, 2023b] Jie Zhang, Xiaosong Ma, Song Guo, and Wenchao Xu. Towards unbiased training in federated open-world semi-supervised learning. In *ICML*, 2023.

[Zhang *et al.*, 2023c] Yifei Zhang, Yankai Chen, Zixing Song, and Irwin King. Contrastive cross-scale graph knowledge synergy. In *KDD*, pages 3422–3433. ACM, 2023.

[Zhang *et al.*, 2023d] Yifei Zhang, Dun Zeng, Jinglong Luo, Zenglin Xu, and Irwin King. A survey of trustworthy federated learning with perspectives on security, robustness and privacy. In *WWW (Companion Volume)*, 2023.

[Zhao *et al.*, 2022] Chen Zhao, Zhipeng Gao, Qian Wang, Zijia Mo, and Xinlei Yu. Fedgan: A federated semi-supervised learning from non-iid data. In *WASA (2)*, volume 13472 of *Lecture Notes in Computer Science*, pages 181–192. Springer, 2022.