
From Distributed Machine Learning to Federated Learning: A Survey

Ji Liu[‡], Jizhou Huang[‡], Yang Zhou[§], Xuhong Li[‡], Shilei Ji[‡], Haoyi Xiong[‡], Dejing Dou^{*‡¶}

Received: date / Accepted: date

Abstract In recent years, data and computing resources are typically distributed in the devices of end users, various regions or organizations. Because of laws or regulations, the distributed data and computing resources cannot be aggregated or directly shared among different regions or organizations for machine learning tasks. Federated learning emerges as an efficient approach to exploit distributed data and computing resources, so as to collaboratively train machine learning models. At the same time, federated learning obeys the laws and regulations and ensures data security and data privacy. In this paper, we provide a comprehensive survey of existing works for federated learning. First, we propose a functional architecture of federated learning systems and a taxonomy of related techniques. Second, we explain the federated learning systems from four aspects: diverse types of parallelism, aggregation algorithms, data communication, and the security of federated learning systems. Third, we present four widely used federated systems based on the functional architecture. Finally, we summarize the limitations and propose future research directions.

Keywords Federated learning · Distributed system · Parallel computing · Security, Privacy

1 Introduction

With billions of connected Internet of Things (IoT) devices [4], smartphones [132] and large websites around the world, in recent years, we have witnessed huge amounts of data generated and dispersed over various mobile devices of end users, or the data centers of different organizations. As the data contain sensitive information of end users or organizations, such as facial images, location-based services, health information [114], or personal economic status, moving the raw data from

* Corresponding author.

‡ Baidu Inc., Beijing, China.

§ Computer Science and Software Engineering Department, Auburn University, Alabama, United States.

¶ Computer and Information Science Department, University of Oregon, United States.

To appear in KAIS

personal devices or data centers of multiple organizations to a centralized server or data center may pose immediate or potential information leakage. Due to the concerns of data security and data privacy, legal restrictions, such as the Cybersecurity Law of the People’s Republic (CLPR) of China [160], the General Data Protection Regulation (GDPR) [133] in European Union, the Personal Data Protection Act (PDP) [32] in Singapore, the California Consumer Privacy Act (CCPA) [1], and the Consumer Privacy Bill of Rights (CPBR) [47] in the United States, have been introduced and put in practice, which makes data aggregation from distributed devices, multiple regions, or organizations, almost impossible [190]. In addition, computing and storage resources are also typically distributed in multiple regions [106] and organizations [98], which cannot be aggregated in a single data center.

Federated Learning (FL) emerges as an efficient approach to exploit the distributed resources to collaboratively train a machine learning model. FL is a distributed machine learning approach where multiple users collaboratively train a model, while keeping the raw data decentralized without being moved to a single server or data center [79, 190]. FL not only exploits the distributed resources to efficiently carry out the training process of machine learning, but also promises to provide security and privacy for the decentralized raw data. Within FL, the raw data, or the data generated based on the raw data with security processing, serves as the training data. FL only allows the intermediate data to be transferred among the distributed computing resources while avoiding the transfer of training data. The distributed computing resources refer to mobile devices of end users or servers of multiple organizations. FL brings the code to the data, instead of bringing the data to the code, and it addresses the fundamental problems of privacy, ownership, and locality of data [120]. In this way, FL can enable multiple users to train a model while satisfying the legal data restrictions.

Traditional centralized machine learning approaches typically gather the distributed raw data generated on different devices or organizations to a single server or a cluster with shared data storage, which may bring serious data privacy and security concerns [205]. The centralized approaches, in general, are associated with diverse challenges, including computational power and training time, and most importantly, security and privacy with respect to distributed data [129]. FL differs from the centralized approach in three aspects. First, FL does not allow direct raw data communication, while the centralized approach has no restriction. Second, FL exploits the distributed computing resources in multiple regions or organizations, while the centralized approach generally only utilizes a single server or a cluster in a single region, which belongs to a single organization. Third, FL generally takes advantage of encryption or other defense techniques to ensure the data privacy or security, while the centralized approach pays little attention to these security issues [205].

The term “federated learning” was first introduced in 2016 [120], which focuses on the unbalanced and non-Independent and Identically Distributed (non-IID) data in mobile devices. The concept of FL was extended to three data scenarios, i.e., horizontal, vertical, and hybrid [190, 205]. The horizontal FL addresses the decentralized data of the same features, while the identifications are different. The vertical FL handles the decentralized data of the same identifications with different features. The hybrid FL deals with the data of different identifications and different features. Then, FL is formally defined as a machine learning approach

where multiple clients collaborate in solving a machine learning problem while the raw data is stored locally and is neither exchanged nor transferred [79].

An FL system is an efficient tool to carry out FL with decentralized data and resources. Several open-source FL systems, e.g., FATE [182], PaddleFL [11], TensorflowFL [55], and Pysyft [136], are now intensively used by both research communities, e.g., healthcare [20], and computer visions [66,109], and by industrial groups, e.g., WeBank [183]. Although various FL systems exist, the architecture of FL systems has common features: In particular, they share the capability to collaboratively train a machine learning model. Most FL systems are composed of four layers, i.e., presentation, user services, FL training, and infrastructure. These four layers enable FL system users to design, execute, and analyze machine learning models with distributed data.

Although FL differs from the centralized machine learning approaches, it not only utilizes novel techniques designed for FL, but also takes advantage of the techniques designed for distributed machine learning. FL exploits parallelization techniques designed for distributed machine learning. For instance, horizontal FL exploits the data parallelism, which trains multiple instances of the same model on different subsets of the training dataset [170]. Vertical FL utilizes model parallelism to distribute parallel paths of a single model to multiple devices in order to handle the data of different features [170]. Multiple aggregation algorithms [30] are proposed to aggregate the models in distributed computing resources. Data transfer techniques are also utilized in FL, e.g., model compression [22]. As FL promises to provide data security and data privacy, diverse defense techniques, e.g., differential privacy [121], homomorphic encryption [59], and Robustness Aggregation [142], are designed to address the possible attacks [49, 71, 181].

There have been a few surveys of FL. Some works [79,89,190] provide a comprehensive study of FL, from the taxonomy of FL to the techniques, e.g., the efficiency, data privacy, security, and applications of FL. Some surveys [89,114,129] focus on the data privacy and security of FL. Other surveys present the application of FL in a specific area, e.g., healthcare informatics [188], mobile edge networks [99], and neural architecture searches [205], and they personalize global models to work better for individual clients [87]. However, few of them present the architecture of FL or analyze parallelization techniques in FL.

In this paper, we provide a survey of federated learning and the related parallelization techniques. The main contributions of this paper are:

- A four-layer FL system architecture, which is useful for discussing the techniques for FL. This architecture can also be a baseline for other work and can help with the assessment and comparison of FL systems.
- A taxonomy of FL-related techniques, including the parallelization techniques, the aggregation algorithms, and the techniques for data communication and security, with a comparative analysis of the existing solutions.
- A discussion of research issues to improve the efficiency and security of FL systems.

This paper is organized as follows. Section 2 gives an overview of the execution of FL, including the FL system architectures and basic functional architecture of FL systems. Section 3 focuses on the techniques used for distributed training of FL and aggregation methods. Section 4 presents the techniques for distributed execution, data communication and data security of FL. Section 5 demonstrates

the existing FL frameworks. Section 6 discusses the open issues raised for the execution of FL with distributed resources. Section 7 summarizes the main findings of this study.

2 An Overview of Federated Learning

In this section, we introduce the basic concepts of federated learning. Then, we present the life cycle of FL models. Afterwards, we detail the functional architecture and the corresponding functionality of FL systems.

2.1 Basic Concepts

Machine learning is the process to automatically extract the models or patterns from data [53]. The models or patterns are expressed as machine learning models. A machine learning model is an ensemble of a model structure, which is typically expressed as a Directed Acyclic Graph (DAG), data processing units, e.g., activation functions in Deep Neural Networks (DNNs), and the associated parameters or hyper-parameters. The input data can be processed through a machine learning model to generate the output, e.g., the prediction results or the classification results, which is the inference process. The machine learning model is generated based on the training data, which is the training process. During the training process, the parameters or the model structure of the machine learning model [61, 68] is adjusted based on a training algorithm in order to improve the performance, e.g., the accuracy or the generalization capacity. The training algorithm is also denoted by machine learning algorithms. The duration of the training process is training time.

According to whether the training data have labels, the training process of machine learning can be classified into four types [170], i.e., supervised learning, unsupervised learning, semi-supervised learning, and reinforcement learning. Supervised learning represents that a machine learning task exploits the training data composed of input features and the corresponding labels [170]. In this paper, we focus on this type of training data. For instance, each data point in the training dataset contains (x, y) , where x represents the input features and y represents the desired output value. Unsupervised learning represents that a machine learning task exploits the training data, which only consists of input features without output values; i.e., each data point only contains x and does not have y . Semi-supervised learning represents that one (generally small) part of the training data contains output values, while the other (generally small) part of the training does not. Reinforcement learning represents that each iteration in the training process considers its observation of the environment from the last iteration.

While the training data become huge, e.g., on the order of terabyte [23], or when the training data is inherently distributed or too big to store on single machines [170], the training process is carried out using distributed resources, which is distributed machine learning. One of the important features of the distributed machine learning is that it can significantly accelerate the training speed so as to reduce the training time. Diverse parallelization techniques are used in distributed machine learning. For instance, Graphics Processing Units (GPUs) using Single

Instruction Multiple Data (SIMD) [45] and Tensor Processing Units (TPUs) using Multiple Instructions Multiple Data (MIMD) [45] are exploited [170]. In addition, distributed machine learning takes advantage of three types of parallelism to parallelize the training process, i.e., data parallelism [170], model parallelism [170], and pipeline parallelism [64, 74, 105]. With the data parallelism approach, the training data is partitioned as many times as the number of computing resources, and all computing resources subsequently apply the same machine learning algorithm to process different chunks of the data sets [170]. With the model parallelism approach, exact copies of the entirety of the data (the training data or the intermediate data) are processed by each computing resource, each of which exploits different parts of the machine learning model [170]. The pipeline parallelism approach combines the data parallelism and the model parallelism. With this approach, each computing resource processes a part of the training data with a part of the machine learning model, while the processing, e.g., computation or communication, at each node can be parallelized [131].

FL is a distributed machine learning approach where multiple users collaboratively train a model, while keeping the raw data distributed without being moved to a single server or data center [79, 190]. The model used for FL is denoted by FL model. FL is first proposed to handle the unbalanced and non-Independent and Identically Distributed (non-IID) data of the same features in mobile devices [120]. Then, the concept of FL is extended to the distributed data of diverse features in multiple organizations [190] or various regions [80]. FL systems are used within one or multiple phases of the life cycle of FL models. An FL system is a distributed system to manage the distributed training process with distributed resources.

FL is a special type of distributed machine learning, which differs from other distributed machine learning approaches in the following three points. First, FL does not allow direct raw data communication, while other approaches have no restriction. As the raw data are of multiple ownerships, FL approaches with this restriction can meet the requirements defined by the related laws, e.g., CLPR [160], GDPR [133], PDPA [32], CCPA [1], and CPBR [47]. In particular, the consent (GDPR Article 6) and the data minimalization principle (GDPR Article 5) limit data collection and storage to only what is consumer-consented and what is absolutely necessary for processing [99]. Second, FL exploits the distributed computing resources in multiple regions or organizations, while the other approaches generally only utilize a single server or a cluster in a single region, which belongs to a single organization. FL enables the collaboration among multiple organizations. Third, FL generally takes advantage of encryption or other defense techniques to ensure the data privacy or security, while the other approaches pay little attention to this security issue [205]. FL promises to ensure the privacy and security of the raw data, as the leakage of information may incur significant financial [35, 150] and reputational [158] losses.

During the training process of FL, an optimization problem is solved as shown in Formula 1. Given n training dataset $\mathcal{D} = D_1, D_2, \dots, D_n$, where each data point $(x, y) \sim \mathcal{D}$, the problem of FL is to learn a function \hat{F} from all possible hypotheses \mathcal{H} , while minimizing the expectation of loss over the distribution of all the dataset \mathcal{D} .

$$\hat{F} = \underset{F \in \mathcal{H}}{\operatorname{argmin}} \mathbb{E}_{(x, y) \in \mathcal{D}} L(y, F(x)), \quad (1)$$

where $L(y, F(x))$ refers to the loss of $F(x)$ to the label y . During the training process, the Stochastic Gradient Descent (SGD) approach [145, 206] is generally used to minimize the loss function using Formula 2.

$$F_{k+1}(x) \leftarrow F_k(x) - \eta_k \nabla F_k(x), \quad (2)$$

where $F_k(x)$ refers to the learned model in the k^{th} iteration, $\nabla F_k(x)$ refers to the gradient of the model at the k^{th} iteration based on the model already obtained $F_k(x)$ and the training dataset, η_k refers to the learning rate, and $F_{k+1}(x)$ refers to the update model of the k^{th} iteration. Within each iteration, there are two phases, i.e., forward propagation and backward propagation. The forward propagation calculates the output based on the input data x using the model, while the backward propagation calculates the gradients $\nabla F_k(x)$ and updates the model. When the calculation is distributed among multiple computing resources, the gradients or models of each computing resource are aggregated using an aggregation algorithm (see details in Section 3.2), in order to achieve consensus of multiple models and to generate a global model. The learning rate can be dynamically adapted using a local adaptive optimizer, e.g., Adam, and/or cross-round learning rate schedulers [199].

2.2 FL Model Life Cycle

The life cycle of an FL model is a description of the state transitions of an FL model from creation to completion [79, 105]. Lo *et al.* [111] propose that the life cycle of an FL model consists of 8 phases: initiated, broadcast, trained, transmitted, aggregated, evaluated, deployed, and monitored. Kairouz *et al.* [79] propose that the life cycle of an FL model includes 6 phases: problem identification, client instrumentation, simulation prototyping, federated model training, model evaluation, and deployment. However, they focus on the FL with distributed data in mobile devices. In this paper, we adopt a combination of workflow life cycle views [79, 111] with a few variations [105, 190], condensed into four phases:

1. The composition phase [79] is for the creation of an FL model, which is used to address a specific machine learning problem, e.g., classification. First, a machine learning model is created to address the problem with certain requirements, e.g., the requirement of accuracy. Then, the machine learning model is adapted to FL scenarios. For instance, if the distributed data is of different features, the machine learning model is partitioned (see details in Section 3.1.2) to process the distributed data.
2. The FL training phase [79, 111, 190] is for the training phase of the FL model. During this phase, a training strategy, which includes parallelism and aggregation algorithms (see details in Section 3), is used to update the parameters, hyper parameters, and even the structure of the network, in order to improve the accuracy and the generalization capacity of the FL model.
3. The FL model evaluation phase [79, 105] is to apply the trained FL models, in order to analyze the performance of the trained FL models on a simulation platform or a real distributed system [65]. As a result, the FL models with the best performance are selected. If the FL models do not meet the requirements, the FL model should be modified, or the training phase should be carried out again.

4. The FL model deployment phase [79] is to deploy the FL model in a real-life scenario to process the data. If the final model can be shared without restriction, there is no difference between the FL model deployment and the model generated from a traditional centralized approach. Otherwise, the deployment of the final model should consider the ownership of the corresponding parts.

2.3 Functional Architecture of FL Systems

The functional architecture of an FL system can be layered as follows [105]: presentation, user services, FL training, and infrastructure. Figure 1 shows this architecture. The higher layers exploit the lower layers to provide their own functionality. A user interacts with an FL system through the presentation layer and realizes independent functionalities at the user services layer. During the training phase of FL models, a Federated Learning Execution Plan (FLEP) is generated, and the corresponding distributed training is carried out at the FL training layer. The FLEP is composed of a type of parallelism, a scheduling strategy, and a fault-tolerance mechanism. The FL system manages the physical resources through the infrastructure layer for the distributed training.

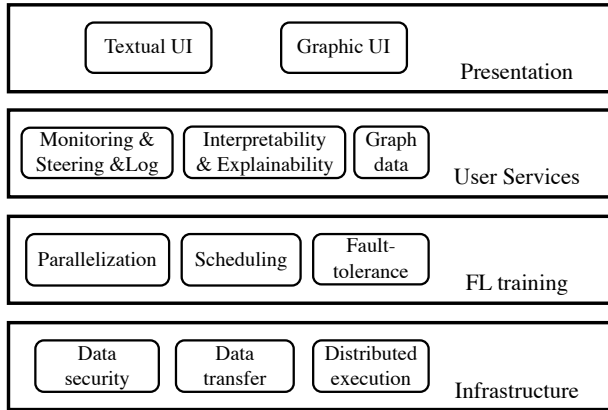


Fig. 1 Functional architecture of an FL system.

2.3.1 Presentation Layer

The presentation layer is a User Interface (UI) for the interaction between Users and FL systems at one or multiple stages of the FL model life cycle. The UI can be textual or graphical. This interface is responsible for designing a new FL model or choosing an existing machine learning model as an FL model. In addition, this layer also supports the modules at the user services layer, e.g., shows the status of the distributed training process. The textual UI is largely used for designing FL models based on the command line or scripts [100]. The models can be directly expressed using Python, with the textual interface in PaddleFL [11],

TensorFlowFL [55], PySyft [148], and FATE [182]. A graphic UI can make the interaction more practical, while the users can drag or drop the data processing element to design an FL model. For instance, FATE [182] provides a Graphic UI (GUI) through a web portal. However, the graphic portal also exploits textual programming languages as inner representations of an FL model.

2.3.2 User Services Layer

The user service layer supports the expected functionalities, i.e., monitoring and steering and log; interpretability and explainability; and graph data. The monitoring enables the users to get the real-time status of the distributed training process. As the training process of FL models can be very long, e.g., from several hours to days [79], it is of much importance to track the execution status, which allows the user to verify whether the training proceeds normally. The log service is generally supported by major FL systems, which can be used to analyze the training process. In addition, the log generated during the training process can be used to debug the system or adjust the FL model. FATE provides a visual monitoring board to users through its GUI. When there are unexpected results or errors during the training process, steering enables users to adjust the training process in order to reduce the time necessary to carry out the distributed training from scratch. Most FL systems can enable the users to stop the training, while the adjustment of parameters is not fully supported by major FL systems. The interpretability of FL is to describe the internals of an FL system in a way that is understandable to humans [52]. The explainability focuses on explaining the representation of data inside an FL model [52]. With interpretability and explainability, the FL system can provide a description of the results of the trained FL model based on the training data and the distributed training process. Shapley values have been used to provide the interpretability [175], while both the interpretability and explainability remain open challenges as each is hard to fully support.

In the real world, graph data widely exist in multiple domains, and a bunch of FL approaches have been proposed to handle the decentralized graph data for community detection [84], financial crime [164], and especially knowledge graph completion [29]. FL is particularly useful in the field of knowledge graph completion, as a knowledge graph could not only contain text but also images or other type of data, i.e., multimodal Knowledge Graphs [202]; and the completion is realized in a collaborative fashion within an FL system [102]. The decentralized graphs can be inter-graph, i.e., the decentralized data belongs to multiple graphs, or intra-graph, i.e., the decentralized data is within one big graph [198], while most of the existing works focus on the intra-graph situation. Horizontal FL techniques can be exploited on the Graph Neural Networks (GNN) [125, 185] with encryption techniques [78] (see details in Section 2.3.4) while the performance of FL may be much worse than that of centralized GNNs [62]. The aggregation algorithms (see details in Section 3.2.1) are also proposed based on the FedAvg [202] or optimal transportation [102]. In addition, decentralized aggregation algorithms (see details in Section 3.2.3) are also proposed to deal with the decentralized graph data for social network [67] and drug discovery [63]. While the fine-tuning of the FL system is time-consuming, Bayesian optimization [203] and evolutionary optimization strategies [174] are utilized to automatically tune the hyper-parameters and the network structure, respectively. Graph data can be vertically distributed where

the features of the nodes are distributed across multiple data owners, and a data owner may or may not have the edges. Vertical FL exploits embeddings [29, 204] or autoencoders [202] to represent the nodes, which can be transferred to train a GNN. In addition, the differential privacy (see details in Section 2.3.4) is combined with the embeddings to protect the data privacy [140] of knowledge graphs.

2.3.3 FL Training Layer

The FL training layer carries out the distributed training process with distributed data and computing resources. This layer consists of three modules: parallelization, scheduling, and fault-tolerance. FL parallelization exploits diverse types of parallelism, e.g., data parallelism, model parallelism, and pipeline parallelism, to generate executable tasks. Through the FL scheduling module, an FL system produces a Scheduling Plan (SP) of executable tasks, which aims at fully exploiting distributed computing resources and preventing training stalling. During the training process, the SP is generally defined by a training algorithm, which aggregates the updates, i.e., gradients or models, from each computing resource in order to generate a final machine learning model. The FL fault-tolerance mechanism handles the failures or errors of task execution and the connection of distributed resources. Reactive approaches are generally exploited, e.g., using check-points, restart, and task replication [18]. A reactive approach reduces the effect of failures after perceiving failures [48]. An FLEP, which captures the execution directives, typically the result of compiling and optimizing the training process of FL models, is generated at this layer.

2.3.4 Infrastructure Layer

The infrastructure layer provides the interaction between an FL system and the distributed resources, including the computing resources, storage resources, network resources, and data resources. This layer contains three modules: a data security module, a data transfer module, and a distributed execution module. The data security module generally exploits Differential Privacy (DP) [3] and encryption techniques, e.g., homomorphic [5], to protect the raw data used during the training process. Although the raw data cannot be directly transferred, intermediate data, e.g., the gradients or models, can be communicated among distributed computing resources. The data transfer module exploits data compression techniques [161] to improve the data transfer efficiency. At this layer, the FLEP generated at the FL training layer is carried out within the distributed execution module; i.e., concrete tasks are executed in distributed computing resources.

3 Distributed Training

In this section, we present the distributed training process for FL. First, we present three types of parallelism in distributed training and the application within FL. The parallelism approaches are generally implemented in the parallelization module. Then, we discuss existing aggregation algorithms for the distributed training, which is implemented in the scheduling module.

3.1 Parallelism & FL Types

Three types of parallelism exist for distributed machine learning: data parallelism, model parallelism, and pipeline parallelism [105,170]. FL can be classified to three types, i.e., horizontal, vertical, and hybrid [190,205]. The horizontal FL generally exploits data parallelism, and the vertical FL typically takes advantage of model parallelism. However, the hybrid FL relies on transfer learning [139], which is not a parallelism approach and is out of the scope of this paper.




Fig. 2 An example of a neural network.

In this section, we take an example of a neural network as shown in Figure 2 to explain the parallelism. In the example, we assume that the model contains three layers and seven data processing nodes (neurons), i.e., $A_1, A_2, B_1, B_2, B_3, C_1, C_2, D$. The arrows represent the data flow among different data processing nodes. The execution of the data processing nodes at each layer can be carried out in parallel, while the execution of different layers should be performed sequentially. The input data contains 4 data points. We assume two/three computing resources owned by two/three users. Each has a part of the input data.

3.1.1 Data Parallelism

Data parallelism is realized by having the data processing performed in parallel at different computing resources, with the same model, on different data points. As shown in Figure 3, data parallelism is exploited when the ensemble of data points is distributed among different computing resources. During the training process of FL, the training data is not transferred among different computing resources, while the intermediate data, e.g., the models or the gradients $\nabla F_k(x)$ in Formula 2, are transferred. The data in each computing resource can be Independent and Identically Distributed Data (IID) or non-IID. FL focuses on the non-IID [120], while other distributed machine learning approaches mainly focus on IID data. With the data parallelism, the FL is horizontal [190], i.e., the data and the calculation are horizontally distributed among multiple computing resources. In addition, this parallelism generally corresponds to the cross-device FL [79], where a large number of devices (mobiles or edge devices) collaboratively participate in training a single global model to have good accuracy. When the number of devices is small, e.g., 2-100, and the computing resources are from diverse organizations, this




Fig. 3 Data parallelism. The forward and backward process of I^1 and I^2 is performed in computing resource 1, while that of I^3 and I^4 is performed in computing resource 2 at the same time. Then the model or gradient is transferred to calculate an average model or gradient to be sent to each computing resource for the following training.

parallelism also corresponds to cross-silo FL [79]. In addition to the general data-parallel schemes for federated learning, some specific privacy-preserved distributed statistical tricks have been invented for federated sparse models [16, 17].

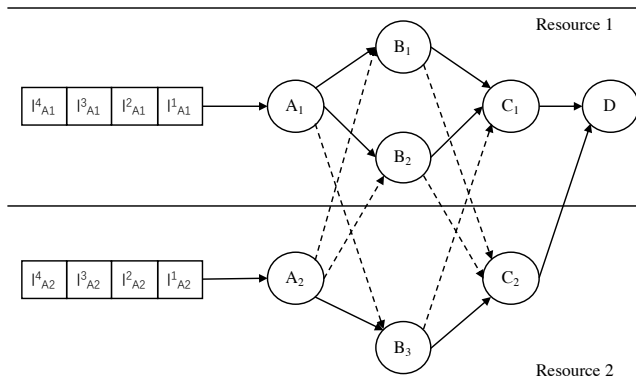


Fig. 4 Model parallelism. The dashed arrows represent inter-computing resource communication. For each input data point I , different parts, i.e., I_{A_1} and I_{A_2} , are distributed at different computing resources.

3.1.2 Model Parallelism

Model parallelism is realized by having independent data processing nodes distributed at different computing resources, so as to process the data points of specific features. Two data processing nodes can be either independent, i.e., the execution of any node does not depend on the output of the other one; or dependent, i.e., there is a data dependency between them [105]. As shown in Figure 4, model parallelism is achieved when different parts of each data point are distributed at different computing resources. For instance, the data process on Node A_1 and that of A_2 can be carried out in parallel. With the model parallelism, vertical FL, where the data points and calculation are vertically distributed among multiple computing resources [59, 190], is realized. In this case, the original model needs to be partitioned to be distributed at different computing resources. Two organizations generally apply this type of FL when each organization owns parts of the features of users and they would like to collaboratively train a model using the data of all the features, which corresponds to cross-silo FL [79]. Most studies of vertical federated learning only support two parties (with or without a central coordinator) [205]. For instance, SecureGBM [43] is proposed to train a tree-based Gradient Boosting Machine (GBM). In order to support multiple parties, the idea of multi-view learning [187] is exploited in a multi-participant, multi-class vertical federated learning framework [42].

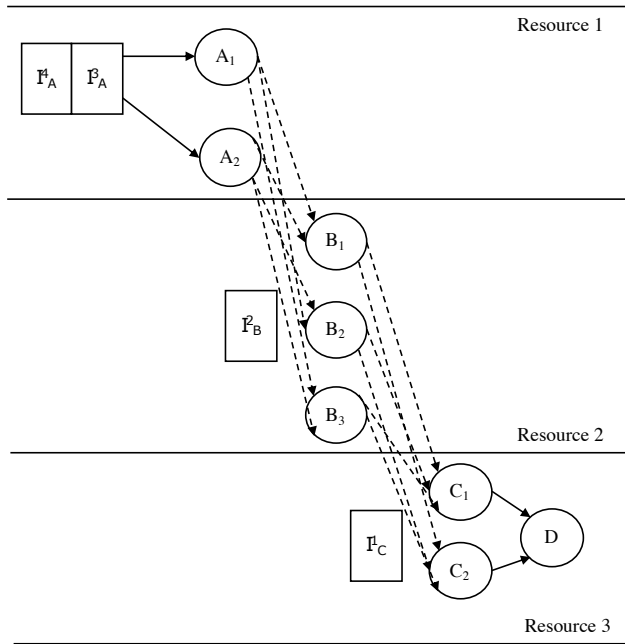


Fig. 5 Pipeline parallelism. The dashed arrows represent inter computing resource communication.

3.1.3 Pipeline Parallelism

Pipeline parallelism is realized by having dependent data processing nodes distributed at different computing resources [74, 131]. As shown in Figure 5, the data processing nodes are distributed at multiple computing resources. While data point I_A^3 is processed in computing resource 1, the outputs of A_1 and A_2 are processed in computing resource 2, and the outputs of B_1 , B_2 , and B_3 are processed in computing resource 3. With this type of parallelism, the dependent data processing nodes can process the data in parallel. As this parallelism may incur many inter-computing resource data transfers, it is not widely used for FL.

3.2 Aggregation Algorithms

With the horizontal FL and data parallelism, aggregation algorithms are used to aggregate the models or gradients generated from the forward and backward propagation in each computing resource. The aggregation algorithms can be either centralized, or hierarchical, and decentralized. The centralized aggregation algorithms generally rely on a centralized server, i.e., a parameter server, to synchronize or schedule the execution of distributed computing resources, while hierarchical aggregation algorithms rely on multiple parameter servers for the model aggregation. The decentralized aggregation algorithms make each computing resource equally perform the calculation based on a predefined protocol, without relying on a centralized server. Please refer to [177] for the details of federated optimization. The characteristics are summarized in Table 1, which can be used to choose appropriate algorithms in a specific situation.

Table 1 Comparison among diverse types of aggregation algorithms. “Complexity” represents the complexity to implement the algorithms (“H” represents high complexity, “M” represents medium complexity, and “L” represents low complexity). “Trust” represents whether the aggregation algorithms require that the data owners trust the centralized server. “Imbalance” represents whether the algorithms can address the unbalanced data. “High-latency” represents whether the algorithms can support the high-latency model or gradient data transfer. “Y” represents that the algorithms support the functionality while “N” represents that the algorithms do not support the functionality.

Type	Complexity	Trust	Imbalance	High-latency
Centralized	L	Y	N	N
Hierarchical	M	Y	Y	Y
Decentralized	H	N	N	Y

3.2.1 Centralized Aggregation

As shown in Figure 6, a single parameter server is used to calculate the average models or gradients sent from multiple computing resources (mobiles). The weights of the model (model) or the gradients are calculated in each computing resource, which are transferred to a parameter server. The parameter server calculates global gradients or global models according to a centralized aggregation algorithm. The global gradients or global models are transferred to each computing resource for




Fig. 6 The architecture of centralized aggregation. “CR” represents computer resource. ω represents the local parameters or the weights of the model calculated in each computing resource. g represents the local gradients in backward propagation in each computing resource. $\bar{\omega}$ represents the global model calculated in the parameter server. \bar{g} represents the global gradients calculated in the parameter server.

the following computation. The update of the model is based on the SGD defined in Formula 2 in both computing resources, or on the parameter server.

A bunch of centralized aggregation algorithms have been proposed. Federated Averaging (FedAvg) [120] algorithm is introduced as the aggregation method in Google’s implementation of an FL system. A centralized server aggregates the machine learning models from selected users. Then, a global model is generated using a weighted sum of each aggregated machine learning model. Afterward, the global model is shared with selected users, and the training process is continued in the computing resource of selected users. However, the trained model of FedAvg may be biased towards computing resources with favorable network conditions [91]. While FedAvg is a straightforward approach, some other methods are proposed to address additional problems. A Federated Stochastic Block Coordinate Descent (FedBCD) [110] algorithm is proposed to reduce the number of communication rounds by enabling multiple local updates before the model communication between a user and the server. In addition, FedBCD also considers the regularization during the training process. The training problem with regularization can be formulated as:

$$\hat{F} = \operatorname{argmin}_{F_{\theta} \in \mathcal{H}} \mathbb{E}_{(x,y) \sim \mathcal{D}} L(y, F_{\theta}(x)) + \lambda \cdot \gamma(\theta), \quad (3)$$

where F , D , \mathcal{H} are the same as those in Formula 1, while $\gamma(\cdot)$ denotes the regularizer and λ is the hyper-parameter. The regularization is exploited to improve the generalization capacity of the trained machine learning model. As the fairness among multiple users is important for an FL system, the Stochastic Agnostic Federated Learning (SAFL) [128] algorithm and the FedMGDA+ [73] algorithm are proposed to achieve fairness during the training process of FL. The fairness represents that the data distribution among multiple users can be equally considered without the influence of unrelated factors. Fairness may also refer to two other concepts: (1) A user gets a final model according to the contribution [115]; and/or (2) Uniform accuracy distribution among all the distributed computing resources [93], which are out of the scope of this paper. In addition, while the computing resources may be heterogeneous, FedProx [92] is proposed to tackle the heterogeneity in an FL system. FedProx enables multiple iterations in each computing resource, while

minimizing a cost function based on the local loss function and the global model. Furthermore, in order to address permutation of data processing nodes during the training process, Federated Matched Averaging (FedMA) [176] is proposed. FedMA exploits an existing approach, i.e., BBP-MAP [193], to generate a matrix, in order to align the data processing nodes of the models from computing resources and the server. SCAFFOLD [81] is proposed to reduce the communication rounds, using stateful variables in the distributed computing resources. Attention-augmented mechanism is exploited in Attentive Federated Aggregation (FedAttOpt) [77] to aggregate the knowledge generated from each computing resource (client), based on the contribution of the model from each client. When the data distribution is heterogeneous among users, personalization remains an open problem. In order to address this problem, the model can be split into local layers and global layers, which has been proposed in adaptive personalized federated learning (APFL) [37], FedPer [6], and pFedMe [38]. The local layers are trained with the decentralized data in each computing resource of users, while the global layers are trained in the computing resources of users and the server. However, it is difficult to choose a dataset and its partition among clients to measure the personalization brought by APFL or FedPer, so as to prove the improvement compared with FedAvg. The attention-augmented mechanism helps reduce the communication rounds. In addition, knowledge distillation can also be exploited to aggregate the models, while requiring that there is data in the centralized server [60]. All these algorithms can handle non-IID data. A comparison among the aforementioned algorithms is proposed in Table 2.

Table 2 Comparison among aggregation algorithms. “Reg” represents regularization. Heterogeneity represents that the computing resources are heterogeneous. “Fairness” represents that the data distribution among multiple users can be equally considered without the influence of unrelated factors. “Permutation” refers to the permutation of data processing nodes during the training process. “C-E” represents communication efficient. “S” represents that the algorithm supports the functionality, while “N” represents that the algorithm does not have support.

Algorithm	Reg	Fairness	Heterogeneity	Permutation	C-E
FedAvg	N	N	N	N	N
FedBCD	S	N	N	N	S
SAFL	N	S	N	N	N
FedMGDA+	N	S	N	N	S
FedProx	N	N	S	N	S
FedMA	N	N	N	S	S
SCAFFOLD	N	N	N	N	S
FedAttOpt	N	N	N	N	S

3.2.2 Hierarchical Aggregation

As shown in Figure 7, a hierarchical architecture is also exploited using multiple parameter servers. A two-layer hierarchical architecture is proposed to reduce the time to transfer models between a parameter server and computing resources [2]. The hierarchical architecture uses a global parameter server (GPS) and multiple region parameter servers. Each region parameter server (RPS) is implemented in a cell base station where the computing resources (mobiles) can be connected,

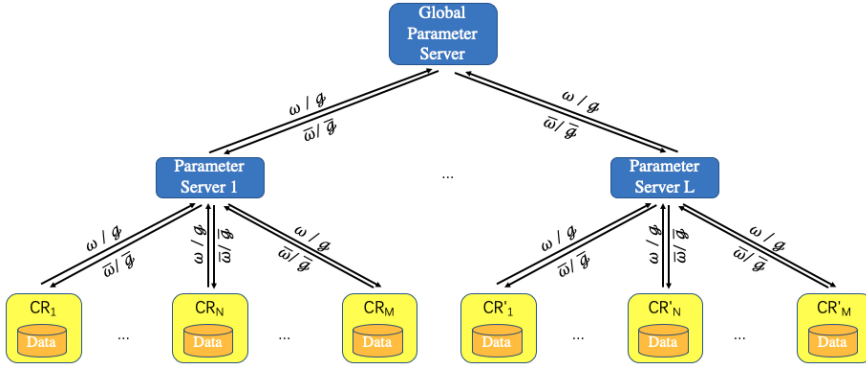


Fig. 7 The architecture of hierarchical aggregation. “CR” represents computer resource. ω represents the local parameters or the weights of the model calculated in each computing resource. g represents the local gradients in backward propagation in each computing resource. $\bar{\omega}$ represents the region or global model calculated in each parameter server. \bar{g} represents the region or global gradients calculated in each parameter server. The region model or gradients are calculated by a region parameter server, while the global model or gradients are calculated by a global parameter server.

with low latency. A hierarchical algorithm, i.e., Hierarchical Federated Learning (HFL) is deployed to realize the model aggregation. Within each iteration of HFL, each RPS calculates an average model using the models of the computing resources within its cluster. It sends the averaged model to the GPS, and it receives a global averaged model at every certain iteration. Afterward, it broadcasts the averaged model to all its computing resources. Some other algorithms, e.g., HierFAVG [107], HFEL [112], and LanFL [192], are similar to HFL, while the SPS is an edge or Local-Area Network (LAN) parameter server and the MPS is a parameter server implemented on the cloud or a Wide-Area Network (WAN). These algorithms take advantage of hierarchical architecture to reduce high-latency model or gradient data transfer, so as to accelerate the training process. In addition, by well-clustering the computing resources to groups, the hierarchical architecture is also exploited to address unbalanced data distributed among multiple computing resources [19, 126], or to address data privacy [172].

3.2.3 Decentralized Aggregation

While collaboratively training a machine learning model with a decentralized aggregation algorithm, the computing resources can be organized with a *connected* topology and can communicate with a peer-to-peer manner, as shown in Figure 8. The degree and connectivity of the topology affects the communication efficiency and the convergence rate of the aggregation algorithm. For a given topology, we define $w_{i,j}$, the weight to scale information flowing from node j to node i , as follows

$$w_{ij} \begin{cases} > 0 & \text{if node } j \text{ is connected to node } i, \text{ or } i = j; \\ = 0 & \text{otherwise.} \end{cases} \quad (4)$$

We further define the **topology matrix** $W = [w_{ij}]_{i,j=0}^{n-1} \in \mathbb{R}^{n \times n}$ as the matrix to represent the topology. In the remainder of this paper, we assume that W satisfies $W\mathbf{1} = \mathbf{1}$ and $\mathbf{1}^T W = \mathbf{1}^T$, i.e., both the row sum and the column sum of W are

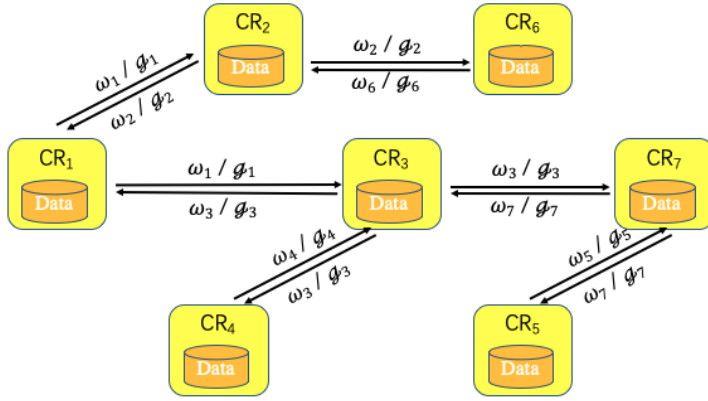


Fig. 8 The architecture of decentralized aggregation. “CR” represents computer resource. ω represents the local parameters, or the weights of the model calculated in each computing resource. g represents the local gradients in backward propagation in each computing resource. When two computing resources are neighbors, they can communicate with each other.

equal to 1, so as to guarantee that the neighborhood averaging will asymptotically approach the global averaging [26, 151, 154]. When a computing resource j is directly connected to computing resource i , i.e., $w_{i,j} \neq 0$, computing resource j is the neighbor of computing resource i . Please note that the weight $w_{i,j}$ denotes the confidence node i has in the information it receives from node j [88], which is different from the bandwidth or data transfer capacity in a network. The centralized aggregation algorithm is a special type of decentralized aggregation with a star topology while only the centralized server communicates with its neighbors. A well designed topology, e.g., an exponential graph [7], can improve the convergence rate, which accelerates the training speed.

With the decentralized SGD (D-SGD), each computing resource maintains a local copy of the global model parameters, and it updates the local copy using the models of its neighbors. According to the order to conduct neighborhood averaging and gradient descent, D-SGD has two common types of realizations: Average-With-Communication (AWC) [88, 96] and Average-Before-Communication (ABC) [28, 178]. AWC can overlap communication and gradient computation, while ABC needs to sequentially calculate and communicate the gradient or model. However, ABC is robust [152], and it converges fast in terms of iterations by exploiting its large learning rate.

In addition, the decentralized aggregation algorithms can be classified to Full Communication (FC) [96] and Partial Communication (PC) [168, 178] according to the number of neighbors. Within the iterations of FC, each computing resource calculates an averaged model or gradient, based on all the models or gradients of the last version from all its neighbors. However, within the iterations of PC, each computing resource calculates an averaged model or gradient based on one or multiple chosen neighbors. With PC, the selection of the neighbors can be based on a gossip algorithm [72]. For instance, a random neighbor can be selected [168]; the neighbors that provide benign models are selected to avoid attack [130].

4 Data Manipulation

At the infrastructure layer of an FL system, there are three types of data manipulation: data security mechanisms, data transfer, and distributed data processing within the distributed execution module. We first present the techniques for the distributed execution in an FL system. Then, we present the techniques for data transfer during the training process of an FL system. Finally, as data security is of much importance to an FL system [129], we present the techniques to protect the data security.

4.1 Distributed Data Processing

While the bandwidth within a single data center is high, e.g., InfiniBand, the High Performance Computing (HPC) libraries, e.g., Message Passing Interface (MPI) [56] or NVIDIA Collective Communications Library (NCCL) [9], are widely exploited for distributed data processing [171]. With MPI or NCCL, the gradients or models in each computing resource can be easily calculated using ring-AllReduce algorithm [51]. However, one of the drawbacks of the HPC libraries is that they lack support for fault-tolerance, as the HPC libraries are designed for high performance servers with high quality networks. When any computing resource within the network becomes unavailable, the distributed training process may be broken.

However, as an FL system is generally implemented for the collaboration of large amounts of mobile device users or different organizations, the network connection among computing resources is of moderate quality, i.e., the bandwidth is not as good as that within a single data center, and the latency is high. For instance, the Internet upload speed is typically much slower than the download speed [86]. Also, some users with unstable wireless communication channels may consequently drop out due to disconnection from the Internet [99]. In this environment, the connection between computing resources and parameter servers has a high possibility of becoming disabled. Thus, Remote Procedure Call (RPC) frameworks are widely exploited, as this kind of framework can ignore the disconnected computing resources and continue the distributed training of an FL system [13], e.g., PaddleFL [11], PySyft [136], or TensorflowFL [55].

4.2 Data Transfer

As the network connection is of moderate quality, the data transfer module mainly focuses on data compression to transfer intermediate data, e.g., gradients or models. Sketched updates are proposed for gradient compression to accelerate the data transfer during the distributed training within a single data center [75, 76, 82, 159]. With the data parallelism and centralized aggregation algorithm, before sending the intermediate data, the intermediate data can be sketched with subsampling [86], quantization [57, 86, 156, 157, 162, 189], sparsification [117, 162], or projection to lower dimensional spaces [147], in each computing resource, in order to reduce the cost to transfer data. Subsampling refers to transferring only a random subset of the intermediate data [86]. Quantization methods encode each value using a fixed number of bits, so as to reduce the length of gradients or models [86].

With the sparsification approach, only selected parts of the intermediate data are transferred, while the selection is based on a threshold, e.g., the gradients larger than a threshold are selected [162]. Then, when the intermediate data is received in the server, they are decompressed to be aggregated according to the aggregation algorithms presented in Section 3.2.1. The convergence of the quantization approach is analyzed in [57], which shows that this approach can also provide good convergence rates [57]. In addition, irrelevant intermediate data can be precluded to be transferred to the server, in order to substantially reduce the communication overhead [180].

4.3 Data Security

Data security is of much importance for data processing. The problem of data security is related to significant financial [35, 150] and reputational [158] losses. For instance, Uber had to pay \$148,000,000 to settle the investigation incurred by a breach of 600,000 drivers' personal information in 2016 [34]. Data security mainly includes two aspects, i.e., data privacy and model security. Data privacy refers to the protection of raw data to avoid raw data information leakage during or after the distributed training of FL systems. Model security refers to the protection of the security of trained models, in order to avoid wrong output based on the trained models incurred by malicious attacks. In this section, we first present the techniques to protect data privacy. Then we present the defense methods for model security.

4.3.1 Data Privacy

The techniques to protect data privacy consist of three types: Trusted Execution Environment (TEE), encryption, Differential Privacy (DP), and anti-Generative Adversarial Network (GAN) methods. These techniques can be combined in FL systems, e.g., the combination of DP and TEE in [58], the combination of encryption and DP [194], and the combination of DP and anti-GAN [166].

A TEE is an environment where the execution is secured and no information can be leaked to unauthorized users. Intel SGX technique [119] has been first proposed as a secure environment while providing a set of security-related instruction codes built within Intel Central Processing Units (CPUs). Then, the implementation of machine learning models has been carried out in the TEE, i.e., Intel SGX, in order to enable collaborative data analysis based on machine learning algorithms while providing a security guarantee [134]. Afterwards, the TEE has been exploited in FL systems, in order to protect the privacy of data in two ways. The first way is to put the entire training process in the TEE of each distributed computing resource to protect the data privacy during the distributed training [31, 127]. The second way is to use TEE to check a small part of the distributed training, while exploiting insecure computing resources, e.g., GPUs, to reduce the training time [200].

As an encryption technique, homomorphic encryption has been used to ensure the data privacy for FL systems [43, 58]. Homomorphic encryption [191] allows specific types of computations to be carried out on encrypted input data, and

to generate an encrypted result, which matches the result of the same computations on the decrypted input data. Two main branches of homomorphic encryption exist, i.e., fully homomorphic encryption and partially homomorphic encryption. The fully homomorphic encryption supports both addition and multiplication on ciphertext, while partially homomorphic encryption only supports either an addition or a multiplication operation on ciphertext, which corresponds to less computational flexibility and better runtime efficiency. Both the fully and partially homomorphic encryptions can be exploited with the horizontal and vertical federated learning. As sharing gradients also leaks the information of training data in horizontal federated learning [49, 95, 201], it is of much importance to protect the privacy of the intermediate data. Thus, the intermediate data can be encrypted using a homomorphic encryption algorithm before being sent to a parameter server [115, 118]. In this way, the intermediate data remain encrypted during the aggregation process while only the computing resource can decrypt the encrypted intermediate data. Even if the transferred encrypted intermediate data is leaked, the information of gradients or models remains safe, and the privacy of the training data is ensured. In addition, partial homomorphic encryption, e.g., Paillier [138], is exploited in vertical federated learning [24, 25]. However, the homomorphic encryption incurs significant costs in computation and communication during distributed training [194]. In order to reduce the overhead of homomorphic encryption, a set of quantized gradients are encrypted [194].

Differential Privacy (DP) protects the data privacy by adding artificial noise to a small part of raw data, while ensuring that the modification does not substantially affect the performance of the machine learning models [3, 39, 50, 184]. DP is widely used in FL systems as the first step to process the raw data, and the output is the training data to be used for the distributed training [83, 97, 108, 141, 149, 184]. With more added noise, the privacy is better protected, i.e., there is less possibility to leak raw data information, while it takes more time to converge for the machine learning models [184]. A trade-off between the privacy protection and the convergence performance can be made by selecting a certain number of distributed resources [153, 167, 184]. However, DP may not be able to ensure the data privacy under certain attacks, e.g., Generative Adversarial Network (GAN) attacks [71].

A well-trained machine learning model can leak information about the training data based on the intermediate data, e.g., gradients [8, 71, 124]. GANs can be used to generate data similar to the training data based on a well-trained machine learning model [54] in either a parameter server [181] or a distributed computing resource [71]. The adversary can reconstruct other participating clients' private data, even if it has no knowledge of the label information using the GANs. Thus, during the distributed training process of FL systems, a malicious user can exploit GANs to infer the training data of other users. DP can be used to prevent the GAN-based attack [71, 166]. In addition, fake training data can be generated based on a GAN and original raw data, which is then used during the distributed training process to prevent the GAN-based attack [113].

4.3.2 Model Security

We mainly focus on poisoning attacks in this section. The objective of poisoning attacks is to reduce the accuracy of machine learning models using artificially designed data, i.e., data poisoning, or models, i.e., model poisoning, in one or

several distributed computing resources during the model aggregation process (see details in Section 3.2.1). There are two ways to carry out poisoning attacks, i.e., data poisoning and model poisoning.

Data poisoning can be realized by modifying the features [46] or the labels [165] of the input data. For instance, malicious users can modify the data points of a certain class C to other classes, and they can then use the modified data points to participate in the distributed training. The modification of the labels is denoted by the label flipping attack. As a result, the accuracy of the trained model has low accuracy in terms of Class C [165]. Model poisoning refers to the attacks in which the updated intermediate data, e.g., gradients or models, are poisoned before being sent to a parameter server in order to reduce the accuracy of the trained model [27, 163]. The goal of the model poisoning is to reduce the performance of the trained model on targeted tasks or classes, while the performance of the model remains unchanged in terms of other tasks or classes [163]. Data poisoning eventually realizes the model poisoning, as it enables some computing resources to update poisoned intermediate data based on the calculation of poisoned training data [46]. However, model poisoning can be more powerful than data poisoning, as model poisoning directly influences the weights of the models and trains in a way that benefits the attack [10]. Both the data poisoning and the model poisoning rely on the backdoor attacks to modify the training data or the intermediate data [27, 46, 163]. Backdoor attacks are performed by embedding the hidden instructions into machine learning models, so that the infected model performs well on benign testing samples when the backdoor is not activated, while its prediction will be changed to the attacker-specified target label when the backdoor is activated by the attacker [94].

In order to defend against these data attacks or model attacks, the malicious users should be identified by analyzing the updated intermediate data using dimensionality reduction methods, e.g., Principal Component Analysis (PCA) [165], anomaly detection [90, 101], or interpretability techniques [14]. In addition, the model poisoning can be incurred by Byzantine failures of certain distributed computing resources [41]. With Byzantine failures, some computing resources (bad users) are manipulated by attackers during the distributed training process, which significantly degrades the performance of the global model in terms of test error [41]. In order to make the training process robust against the Byzantine failures, the bad users can be identified by analyzing the updated intermediate data using a hidden Markov model [40, 130] or via secure aggregation protocols [70].

5 Federated Learning Frameworks

FL systems are widely applied in diverse domains, e.g., mobile service, healthcare [188], and finance [89]. An FL system generally exploits an FL framework, which is deployed on distributed resources. In this section, we present four widely used FL frameworks: PaddleFL [11], TensorFlowFederated [55], FATE [182], and PySyft [136].

5.1 PaddleFL

PaddleFL is an open source federated learning framework based on PaddlePaddle [116], which is supported by Baidu. At the presentation layer, PaddleFL provides a textual UI for the interaction between users and the FL system. At the User Services layer, PaddleFL provides the log and monitoring supports, and it can leverage the interpretability module [12] of PaddlePaddle in the future. At the FL training layer, PaddleFL can realize data parallelism (horizontal FL) and model parallelism (vertical FL). It supports multiple aggregation algorithms, e.g., FedAvg, and fault-tolerance. At the infrastructure layer, PaddleFL exploits RPC for the distributed execution. PaddleFL exploits DP to protect the data security. PaddleFL is widely used in multiple domains, e.g., Natural Language Processing (NLP), Computing Vision (CV) [109], and recommendation.

5.2 TensorFlowFederated

TensorFlow Federated (TFF) [55] is an open-source framework for federated learning on decentralized data, which is supported by Google. TFF also provides a textual UI through Python. TFF supports the monitoring and log functionality at the user service layer. TFF supports data parallelism (horizontal FL), multiple aggregation algorithms, and fault-tolerance of mobile devices. TFF exploits RPC for the distributed execution and DP for the protection of data privacy. TFF enables Android mobile users to predict the next word while using the keyboard on their mobile phones [120, 122].

5.3 FATE

FATE [182] is an open-source FL framework supported by WeBank. FATE provides both a graphical and textual UI. FATE can support the monitoring of distributed training through a web portal. FATE takes advantage of database management systems (DBMS) to track the execution status. FATE can enable horizontal (data parallelism), vertical (model parallelism), and hybrid federated learning. FATE exploits both the DP and HE to protect the data privacy. In addition, FATE exploits RPC to perform the distributed execution.

5.4 PySyft

PySyft [148] is an open-source FL framework based on the PyTorch framework [144]. PySyft is written in Python and provides a textual UI based on Python. PySyft mainly supports the data parallelism and model parallelism based on an aggregator or orchestrating server. The aggregator or orchestrating server sends a part of the model to participating clients to process local data and gets results for federated averaging. PySyft exploits DP and encryption techniques to protect the data security. PySyft exploits multiple communication protocols for distributed execution, e.g., RPC, websocket [44] etc.

5.5 Concluding Remarks

Diverse FL frameworks exist while each has its advantage. We summarize the characteristics of each framework in Table 3, so as to help select a proper framework for use. From the table, we can see that all the frameworks implement the centralized aggregation algorithms, while employing DP and HE for the data security. PaddleFL can exploit Paddle to realize data, model, and pipeline parallelism. FATE and TFF are based on Tensorflow as the engine, while FATE can provide Web portal UI, which is convenient for novices. PySyft is compatible with PyTorch, which can easily handle the PyTorch-based tasks, while PaddleFL is compatible with Paddle, which can easily deal with rich pre-trained models published in PaddleHub [137].

Table 3 Comparison among diverse frameworks. “Aggregation” represents the type of aggregation algorithms. “Textual” represents the textual UI, while “Web” represents Web portal.

Framework	Engine	Aggregation	UI	Parallelism	Security
PaddleFL	Paddle	Centralized	textual	Data/Model/Pipeline	DP/HE
TFF	TensorFlow	Centralized	textual	Data/Model	DP/HE
FATE	TensorFlow	Centralized	Web	Data/Model	DP/HE
PySyft	PyTorch	Centralized	textual	Data	DP/HE

Table 4 represents the support of diverse types of FL in terms of data distribution. All the frameworks support horizontal FL, while vertical FL is supported by three frameworks except TFF. PySyft cannot directly support the vertical FL, while PyVertical [146], which is built upon PySyft, can be used to support vertical FL with the compatibility of PyTorch models. The hybrid FL is only supported by Paddle and FATE. In addition, all the frameworks support the execution with GPU. In practice, although PaddleFL may correspond to slightly longer time, the accuracy of the trained model can be higher than that of TFF and FATE, while PySyft may generate “out of memory” errors [85].

Table 4 Comparison among diverse frameworks for the support of diverse FL types, e.g., horizontal FL, vertical FL, and hybrid FL, and GPU. $\|\checkmark\|$ represents that the support is not realized by itself but a close one.

		PaddleFL	TFF	FATE	PySyft
Types	Horizontal	\checkmark	\checkmark	\checkmark	\checkmark
	Vertical	\checkmark	\times	\checkmark	$\ \checkmark\ $
	Hybrid	\checkmark	\times	\checkmark	\times
GPU		\checkmark	\checkmark	\checkmark	\checkmark

6 Research Directions

Although much work has been done on the FL systems, there remain some limitations, e.g., interpretability of FL, decentralized aggregation, FL on graphs, benchmarks of FL systems, and applications to distributed intelligent systems. This

section discusses the limitations of the existing frameworks and proposes new research directions.

6.1 Benchmarks

Several datasets exist for experiments on FL systems. For instance, Federated Extended MNIST (FEMNIST) [21] is built by partitioning the data in Extended MNIST [33] based on each writer. Shakespeare [120] is built from The Complete Works of William Shakespeare [155] based on each speaking role. Both of these datasets can be used for horizontal FL. However, no public datasets exist for vertical FL or transfer FL. In addition, no open decentralized IID or non-IID distribution of popular datasets, e.g., ImageNet [36], exist for FL systems.

6.2 Interpretability

Deep neural networks have excellent performance in various areas, while it is often difficult to understand the results of deep neural network models, especially within FL systems. Shapley values have been used to provide the interpretability [175], while it focuses on vertical FL. When multiple users collaboratively train an FL model, it remains an open problem to evaluate the contributions of each user, which helps provide evidence for the incentive of each user. The primary incentive for clients to participate in federated learning is obtaining better models [87], while the benefit of participating in federated learning for clients who have sufficient private data to train accurate local models is disputable. Interpretability can help understand the contributions of each user and provide an objective opinion on the incentive strategy within an FL system. In addition, the interpretability helps domain experts to understand the relationship between data and the final trained model in critical domains, e.g., healthcare and finance. However, the interpretability within FL systems remains an open problem.

6.3 Decentralized Aggregation

Current aggregation algorithms of FL systems focus on the full connection or star connection topology, while other topologies, e.g., dynamic exponential-2 graph, may help accelerate the distributed training with FL systems [96]. In addition, well-known graph algorithms, e.g., graph partitioning algorithms, and ad-hoc policies can be exploited to help better distribute computing resources with the topology defined in Section 3.2.3 in order to improve the efficiency of FL systems. While the peer-to-peer communication enables the FL with an arbitrary topology matrix, the data security under diverse attacks, e.g., data or model poisoning, GAN-based attacks, remain open problems and deserve further investigation.

6.4 Federated Learning on Graphs

Graphs or graph neural networks (GNN) [169] have gained increasing popularity in multiple domains, e.g., social network, knowledge graph, and recommender system.

FL frameworks for graphs, i.e., GraphFL [173], and GNN, i.e., SGNN [123], have been proposed to train a model with decentralized graphs. However, the data security of FL on graphs remains an open problem. In addition, while a multimodal knowledge graph could not only contain text but also images or other type of data [202], it is worth further exploration to efficiently support the multimodal knowledge graph construction within an FL system [102].

6.5 Imbalanced Data

Although FL focuses on the non-IID data, the real-world decentralized data usually exhibit an imbalanced distribution [69, 186]. While the imbalanced data exist in multiple areas, such as computer vision [135], bioinformatics, and biomedicine [195], learning from such data requires special attention upon data sampling [195, 197], data augmentation [135], and loss function designs [179]. The imbalanced data is related to diverse tasks, e.g., two-class or multi-class classification [15, 196]. However, an optimized approach can be proposed to address the imbalanced data within FL systems.

6.6 Applications to Distributed Intelligent Systems

Machine learning algorithms have been widely used to boost the performance of intelligent systems, while FL systems could further enhance intelligent systems [103] in distributed computing environments [104, 143] with privacy and security ensured. An intelligent system is a group of machines that has the capacity to gather data, analyze the data, and respond to other systems or the world around. With FL systems, the distributed data can be exploited to generate models of high performance so as to produce smart responses.

7 Conclusion

In this paper, we discussed the current state of the art of FL systems, including the functional architecture of FL systems, distributed training, and data manipulation.

First, we presented an overview of FL systems. In particular, we introduced the life cycle of FL models, including four phases. Then, we presented the four-layer functional architecture of FL systems, including presentation, user services, FL training, and infrastructure, and we presented each layer in detail.

Second, we detailed the distributed training with two parts, i.e., parallelism and aggregation algorithms. We presented three types of parallelism, including data parallelism, model parallelism, and pipeline parallelism. We associate each parallelism to a corresponding type of FL. For instance, data parallelism is associated with the horizontal FL, which corresponds to cross-device or cross-silo FL. Model parallelism is related to vertical FL and cross-silo FL. We presented the features of different aggregation algorithms in three types: centralized aggregation, hierarchical aggregation, and decentralized aggregation.

Third, we presented the techniques for data manipulation within FL systems. We showed that FL systems prefer RPC for the distributed execution, to handle

the fault-tolerance because of moderate network connection. Intermediate data are sketched in order to compress the data, so as to reduce the data communication time. In addition, we presented the data privacy and model security attacks and corresponding defense techniques, e.g., DP, HE, TEE, and the analysis of updated intermediate data for malicious user identification.

We mainly introduced four FL systems: PaddleFL, TensorFlowFederated, FATE, and PySyft. The current solutions primarily focus on the horizontal FL. And we identified five research directions that deserve further investigation: benchmarks, interpretability, decentralized aggregation, FL on graphs, imbalanced data, and the applications of FL systems to distributed intelligent systems.

References

1. California consumer privacy act home page. <https://www.caprivacy.org/>. Online; accessed 14/02/2021.
2. M Salehi Heydar Abad, Emre Ozfatura, Deniz Gunduz, and Ozgur Ercetin. Hierarchical federated learning across heterogeneous cellular networks. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8866–8870, 2020.
3. Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *ACM SIGSAC conf. on computer and communications security*, pages 308–318, 2016.
4. Zakaria Abou El Houda, Abdelhakim Hafid, and Lyes Khoukhi. Co-iot: a collaborative ddos mitigation scheme in iot environment based on blockchain using sdn. In *IEEE Global Communications Conference (GLOBECOM)*, pages 1–6, 2019.
5. Yoshinori Aono, Takuya Hayashi, Lihua Wang, Shiho Moriai, et al. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5):1333–1345, 2017.
6. Manoj Ghuhana Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv preprint arXiv:1912.00818*, 2019.
7. Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Mike Rabbat. Stochastic gradient push for distributed deep learning. In *Int. Conf. on Machine Learning (ICML)*, volume 97, pages 344–353, 2019.
8. Giuseppe Ateniese, Luigi V Mancini, Angelo Spognardi, Antonio Villani, Domenico Vitali, and Giovanni Felici. Hacking smart machines with smarter ones: How to extract meaningful data from machine learning classifiers. *Int. Journal of Security and Networks*, 10(3):137–150, 2015.
9. Ammar Ahmad Awan, Ching-Hsiang Chu, Hari Subramoni, and Dhabaleswar K Panda. Optimized broadcast for deep learning workloads on dense-gpu infiniband clusters: Mpi or nccl? In *European MPI Users’ Group Meeting*, pages 1–9, 2018.
10. Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, pages 2938–2948, 2020.
11. Baidu. Federated deep learning in paddlepaddle. <https://github.com/PaddlePaddle/PaddleFL>. Online; accessed 16/02/2021.
12. Baidu. Paddlepaddle interpretability. <https://github.com/PaddlePaddle/InterpretDL>. Online; accessed 13/03/2021.
13. Daniel J Beutel, Taner Topal, Akhil Mathur, Xinchu Qiu, Titouan Parcollet, and Nicholas D Lane. Flower: A friendly federated learning research framework. *arXiv preprint arXiv:2007.14390*, 2020.
14. Arjun Nitin Bhagoji, Supriyo Chakraborty, Prateek Mittal, and Seraphin Calo. Analyzing federated learning through an adversarial lens. In *Int. Conf. on Machine Learning (ICML)*, pages 634–643, 2019.
15. Jingjun Bi and Chongsheng Zhang. An empirical comparison on state-of-the-art multi-class imbalance learning algorithms and a new diversified ensemble learning scheme. *Knowledge-Based Systems*, 158:81–93, 2018.

16. Jiang Bian, Haoyi Xiong, Wei Cheng, Wenqing Hu, Zhishan Guo, and Yanjie Fu. Multi-party sparse discriminant learning. In *2017 IEEE International Conference on Data Mining (ICDM)*, pages 745–750. IEEE, 2017.
17. Jiang Bian, Haoyi Xiong, Yanjie Fu, Jun Huan, and Zhishan Guo. Mp2sda: Multi-party parallelized sparse discriminant learning. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 14(3):1–22, 2020.
18. Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloé Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Van Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. Towards federated learning at scale: System design. In *Machine Learning and Systems (MLSys)*, 2019.
19. Christopher Briggs, Zhong Fan, and Peter Andras. Federated learning with hierarchical clustering of local updates to improve training on non-iid data. In *Int. Joint Conf. on Neural Networks (IJCNN)*, pages 1–9. IEEE, 2020.
20. Theodora S Brisimi, Ruidi Chen, Theofanie Mela, Alex Olshevsky, Ioannis Ch Paschalidis, and Wei Shi. Federated learning of predictive models from federated electronic health records. *Int. Journal of Medical Informatics (IJMI)*, 112:59–67, 2018.
21. 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.
22. Sebastian Caldas, Jakub Konečný, H Brendan McMahan, and Ameet Talwalkar. Expanding the reach of federated learning by reducing client resource requirements. *arXiv preprint arXiv:1812.07210*, 2018.
23. Kevin Canini, Tushar Chandra, Eugene Ie, Jim McFadden, Ken Goldman, Mike Gunter, Jeremiah Harmsen, Kristen LeFevre, Dmitry Lepikhin, Tomas Lloret Llinares, et al. Sibyl: A system for large scale supervised machine learning. *Technical Talk*, 1:113, 2012.
24. Ferhat Özgür Çatak. Secure multi-party computation based privacy preserving extreme learning machine algorithm over vertically distributed data. In *Int. Conf. on Neural Information Processing (ICONIP)*, pages 337–345, 2015.
25. Ferhat Özgür Çatak and Ahmet Fatih Mustacoglu. Cpp-elm: Cryptographically privacy-preserving extreme learning machine for cloud systems. *Int. Journal of Computational Intelligence Systems*, 11(1):33–44, 2018.
26. Sampri Chatterjee and Eugene Seneta. Towards consensus: Some convergence theorems on repeated averaging. *Journal of Applied Probability*, pages 89–97, 1977.
27. Chien-Lun Chen, Leana Golubchik, and Marco Paolieri. Backdoor attacks on federated meta-learning. *arXiv preprint arXiv:2006.07026*, 2020.
28. Jianshu Chen and Ali H Sayed. Diffusion adaptation strategies for distributed optimization and learning over networks. *IEEE Transactions on Signal Processing*, 60(8):4289–4305, 2012.
29. Mingyang Chen, Wen Zhang, Zonggang Yuan, Yantao Jia, and Huajun Chen. Federated embedding knowledge graphs in federated setting. *arXiv preprint arXiv:2010.12882*, 2020.
30. Yang Chen, Xiaoyan Sun, and Yaochu Jin. Communication-efficient federated deep learning with layerwise asynchronous model update and temporally weighted aggregation. *IEEE Transactions on Neural Networks and Learning Systems*, 31(10):4229–4238, 2019.
31. Yu Chen, Fang Luo, Tong Li, Tao Xiang, Zheli Liu, and Jin Li. A training-integrity privacy-preserving federated learning scheme with trusted execution environment. *Information Sciences*, 522:69–79, 2020.
32. Warren B Chik. The singapore personal data protection act and an assessment of future trends in data privacy reform. *Computer Law & Security Review*, 29(5):554–575, 2013.
33. Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre Van Schaik. Emnist: Extending mnist to handwritten letters. In *Int. Joint Conf. on Neural Networks (IJCNN)*, pages 2921–2926, 2017.
34. Kate Conger. Uber settles data breach investigation for \$148 million. <https://www.nytimes.com/2018/09/26/technology/uber-data-breach.html>. Online; accessed 17/02/2021.
35. Kate Conger. Uber settles data breach investigation for \$148 million, 2018. <https://www.nytimes.com/2018/09/26/technology/uber-data-breach.html>. Online; accessed 28/02/2021.
36. Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009.

37. Yuyang Deng, Mohammad Mahdi Kamani, and Mehrdad Mahdavi. Adaptive personalized federated learning. *arXiv preprint arXiv:2003.13461*, 2020.
38. Canh T Dinh, Nguyen H Tran, and Tuan Dung Nguyen. Personalized federated learning with moreau envelopes. *arXiv preprint arXiv:2006.08848*, 2020.
39. Cynthia Dwork. Differential privacy: A survey of results. In *Int. conf. on theory and applications of models of computation*, pages 1–19, 2008.
40. Sean R Eddy. What is a hidden markov model? *Nature biotechnology*, 22(10):1315–1316, 2004.
41. Minghong Fang, Xiaoyu Cao, Jinyuan Jia, and Neil Gong. Local model poisoning attacks to byzantine-robust federated learning. In *USENIX Security Symposium (USENIX Security)*, pages 1605–1622, 2020.
42. Siwei Feng and Han Yu. Multi-participant multi-class vertical federated learning. *arXiv preprint arXiv:2001.11154*, 2020.
43. Zhi Feng, Haoyi Xiong, Chuanyuan Song, Sijia Yang, Baoxin Zhao, Licheng Wang, Zeyu Chen, Shengwen Yang, Liping Liu, and Jun Huan. Securegbm: Secure multi-party gradient boosting. In *IEEE Int. Conf. on Big Data (Big Data)*, pages 1312–1321, 2019.
44. Ian Fette and Alexey Melnikov. The websocket protocol, 2011.
45. Michael J Flynn. Some computer organizations and their effectiveness. *IEEE Transactions on Computers*, 100(9):948–960, 1972.
46. Clement Fung, Chris JM Yoon, and Ivan Beschastnikh. Mitigating sybils in federated learning poisoning. *arXiv preprint arXiv:1808.04866*, 2018.
47. B. M. Gaff, H. E. Sussman, and J. Geetter. Privacy and big data. *Computer*, 47(6):7–9, 2014.
48. K Ganga and S Karthik. A fault tolerant approach in scientific workflow systems based on cloud computing. In *Int. Conf. on Pattern Recognition, Informatics and Mobile Engineering*, pages 387–390, 2013.
49. Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients—how easy is it to break privacy in federated learning? *arXiv preprint arXiv:2003.14053*, 2020.
50. Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. *arXiv preprint arXiv:1712.07557*, 2017.
51. Andrew Gibiansky. Bringing hpc techniques to deep learning. <https://andrew.gibiansky.com/blog/machine-learning/baidu-allreduce/>, 2017. Online; accessed 2020-08-12.
52. Leilani H Gilpin, David Bau, Ben Z Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal. Explaining explanations: An overview of interpretability of machine learning. In *IEEE Int. Conf. on Data Science and Advanced Analytics (DSAA)*, pages 80–89. IEEE, 2018.
53. Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.
54. Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *Int. Conf. on Learning Representations (ICLR)*, 2015.
55. Google. Tensorflow federated: Machine learning on decentralized data. <https://www.tensorflow.org/federated>. Online; accessed 16/02/2021.
56. William Gropp, William D Gropp, Ewing Lusk, Anthony Skjellum, and Argonne Distinguished Fellow Emeritus Ewing Lusk. *Using MPI: portable parallel programming with the message-passing interface*, volume 1. MIT press, 1999.
57. Farzin Haddadpour, Mohammad Mahdi Kamani, Aryan Mokhtari, and Mehrdad Mahdavi. Federated learning with compression: Unified analysis and sharp guarantees. *arXiv preprint arXiv:2007.01154*, 2020.
58. Meng Hao, Hongwei Li, Guowen Xu, Sen Liu, and Haomiao Yang. Towards efficient and privacy-preserving federated deep learning. In *IEEE Int. Conf. on Communications (ICC)*, pages 1–6, 2019.
59. Stephen Hardy, Wilko Henecka, Hamish Ivey-Law, Richard Nock, Giorgio Patrini, Guillaume Smith, and Brian Thorne. Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption. *arXiv preprint arXiv:1711.10677*, 2017.
60. Chaoyang He, Murali Annavaram, and Salman Avestimehr. Group knowledge transfer: Collaborative training of large cnns on the edge. *arXiv preprint arXiv:2007.14513*, 2020.
61. Chaoyang He, Murali Annavaram, and Salman Avestimehr. Towards non-iid and invisible data with fednas: Federated deep learning via neural architecture search. *arXiv preprint arXiv:2004.08546*, 2020.

62. Chaoyang He, Keshav Balasubramanian, Emir Ceyani, Yu Rong, Peilin Zhao, Junzhou Huang, Murali Annavaram, and Salman Avestimehr. Fedgraphnn: A federated learning system and benchmark for graph neural networks. *arXiv preprint arXiv:2104.07145*, 2021.
63. Chaoyang He, Emir Ceyani, Keshav Balasubramanian, Murali Annavaram, and Salman Avestimehr. Spreadgnn: Serverless multi-task federated learning for graph neural networks. *arXiv preprint arXiv:2106.02743*, 2021.
64. Chaoyang He, Shen Li, Mahdi Soltanolkotabi, and Salman Avestimehr. Pipetransformer: Automated elastic pipelining for distributed training of large-scale models. In *Int. Conf. on Machine Learning*, volume 139 of *Machine Learning Research*, pages 4150–4159, 2021.
65. Chaoyang He, Songze Li, Jinhyun So, Xiao Zeng, Mi Zhang, Hongyi Wang, Xiaoyang Wang, Praneeth Vepakomma, Abhishek Singh, Hang Qiu, et al. Fedml: A research library and benchmark for federated machine learning. *arXiv preprint arXiv:2007.13518*, 2020.
66. Chaoyang He, Alay Dilipbhai Shah, Zhenheng Tang, Di Fan, Adarshan Naiynar Sivashunmugam, Keerti Bhogaraju, Mita Shimpi, Li Shen, Xiaowen Chu, Mahdi Soltanolkotabi, et al. Fedcv: A federated learning framework for diverse computer vision tasks.
67. Chaoyang He, Conghui Tan, Hanlin Tang, Shuang Qiu, and Ji Liu. Central server free federated learning over single-sided trust social networks. *arXiv preprint arXiv:1910.04956*, 2019.
68. Chaoyang He, Haishan Ye, Li Shen, and Tong Zhang. Milenas: Efficient neural architecture search via mixed-level reformulation. In *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2020.
69. Haibo He and Edward A. Garcia. Learning from imbalanced data. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 21(9):1263–1284, 2009.
70. Lie He, Sai Praneeth Karimireddy, and Martin Jaggi. Secure byzantine-robust machine learning. *arXiv preprint arXiv:2006.04747*, 2020.
71. Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. Deep models under the gan: information leakage from collaborative deep learning. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 603–618, 2017.
72. Chenghao Hu, Jingyan Jiang, and Zhi Wang. Decentralized federated learning: A segmented gossip approach. *arXiv preprint arXiv:1908.07782*, 2019.
73. Zeou Hu, Kiarash Shaloudegi, Guojun Zhang, and Yaoliang Yu. Fedmgda+: Federated learning meets multi-objective optimization. *arXiv preprint arXiv:2006.11489*, 2020.
74. Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Mia Xu Chen, Dehao Chen, HyoukJoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, et al. Gpipe: Efficient training of giant neural networks using pipeline parallelism. *arXiv preprint arXiv:1811.06965*, 2018.
75. Nikita Ivkin, Daniel Rothchild, Enayat Ullah, Vladimir Braverman, Ion Stoica, and Raman Arora. Communication-efficient distributed sgd with sketching. *arXiv preprint arXiv:1903.04488*, 2019.
76. Jiawei Jiang, Fangcheng Fu, Tong Yang, and Bin Cui. Sketchml: Accelerating distributed machine learning with data sketches. In *Int. Conf. on Management of Data*, pages 1269–1284, 2018.
77. Jing Jiang, Shaoxiong Ji, and Guodong Long. Decentralized knowledge acquisition for mobile internet applications. *World Wide Web*, pages 1–17, 2020.
78. Meng Jiang, Taeho Jung, Ryan Karl, and Tong Zhao. Federated dynamic gnn with secure aggregation. *arXiv preprint arXiv:2009.07351*, 2020.
79. Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
80. Peter Kairouz, H. Brendan McMahan, Aurélien Bellet, Brendan Avent, Arjun Nitin Bhagoji, Mehdi Bennis, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G.L. D’Oliveira, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Phillip B. Gibbons, Badih Ghazi, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konečný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrède Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang,

- Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1), 2021.
81. Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *Int. Conf. on Machine Learning (ICML)*, pages 5132–5143, 2020.
 82. Sai Praneeth Karimireddy, Quentin Rebjock, Sebastian Stich, and Martin Jaggi. Error feedback fixes signsgd and other gradient compression schemes. In *Int. Conf. on Machine Learning (ICML)*, pages 3252–3261, 2019.
 83. Kleomenis Katevas, Eugene Bagdasaryan, Jason Waterman, Mohamad Mounir Safadieh, Eleanor Birrell, Hamed Haddadi, and Deborah Estrin. Policy-based federated learning. *arXiv e-prints*, pages arXiv–2003, 2020.
 84. Chuyang Ke and Jean Honorio. Federated myopic community detection with one-shot communication. *arXiv preprint arXiv:2106.07255*, 2021.
 85. Ivan Kholod, Evgeny Yanaki, Dmitry Fomichev, Evgeniy Shalugin, Evgenia Novikova, Evgeny Filippov, and Mats Nordlund. Open-source federated learning frameworks for iot: A comparative review and analysis. *Sensors*, 21(1):167, 2021.
 86. Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
 87. Viraj Kulkarni, Milind Kulkarni, and Aniruddha Pant. Survey of personalization techniques for federated learning. In *World Conf. on Smart Trends in Systems, Security and Sustainability (WorldS4)*, pages 794–797, 2020.
 88. Anusha Lalitha, Osman Cihan Kilinc, Tara Javidi, and Farinaz Koushanfar. Peer-to-peer federated learning on graphs. *arXiv preprint arXiv:1901.11173*, 2019.
 89. Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, and Bingsheng He. A survey on federated learning systems: vision, hype and reality for data privacy and protection. *arXiv preprint arXiv:1907.09693*, 2019.
 90. Suyi Li, Yong Cheng, Yang Liu, Wei Wang, and Tianjian Chen. Abnormal client behavior detection in federated learning. *arXiv preprint arXiv:1910.09933*, 2019.
 91. Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
 92. Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In *Machine Learning and Systems*, volume 2, pages 429–450, 2020.
 93. Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*, 2019.
 94. Yiming Li, Baoyuan Wu, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. Backdoor learning: A survey. *arXiv preprint arXiv:2007.08745*, 2020.
 95. Zhaorui Li, Zhicong Huang, Chaochao Chen, and Cheng Hong. Quantification of the leakage in federated learning. *arXiv preprint arXiv:1910.05467*, 2019.
 96. Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5330–5340, 2017.
 97. Zhicong Liang, Bao Wang, Quanquan Gu, Stanley Osher, and Yuan Yao. Exploring private federated learning with laplacian smoothing. *arXiv preprint arXiv:2005.00218*, 2020.
 98. Misbah Liaquat, Victor Chang, Abdullah Gani, Siti Hafizah Ab Hamid, Muhammad Toseef, Umar Shoaib, and Rana Liaquat Ali. Federated cloud resource management: Review and discussion. *Journal of Network and Computer Applications*, 77:87–105, 2017.
 99. Wei Yang Bryan Lim, Nguyen Cong Luong, Dinh Thai Hoang, Yutao Jiao, Ying-Chang Liang, Qiang Yang, Dusit Niyato, and Chunyan Miao. Federated learning in mobile edge networks: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 22(3):2031–2063, 2020.
 100. Bill Yuchen Lin, Chaoyang He, Zihang Zeng, Hulin Wang, Yufen Huang, Mahdi Soltanolkotabi, Xiang Ren, and Salman Avestimehr. Fednlp: A research platform for federated learning in natural language processing. *arXiv preprint arXiv:2104.08815*, 2021.

101. Jierui Lin, Min Du, and Jian Liu. Free-riders in federated learning: Attacks and defenses. *arXiv preprint arXiv:1911.12560*, 2019.
102. Yilun Lin, Chaochao Chen, Cen Chen, and Li Wang. Improving federated relational data modeling via basis alignment and weight penalty. *arXiv preprint arXiv:2011.11369*, 2020.
103. Ji Liu, Carlyna Bondiombouy, Lei Mo, and Patrick Valduriez. Two-phase scheduling for efficient vehicle sharing. *IEEE Transactions on Intelligent Transportation Systems (TITS)*, 2020.
104. Ji Liu, Esther Pacitti, Patrick Valduriez, Daniel De Oliveira, and Marta Mattoso. Multi-objective scheduling of scientific workflows in multisite clouds. *Future Generation Computer Systems*, 63:76–95, 2016.
105. Ji Liu, Esther Pacitti, Patrick Valduriez, and Marta Mattoso. A survey of data-intensive scientific workflow management. *Journal of Grid Computing*, 13(4):457–493, 2015.
106. Ji Liu, Luis Pineda, Esther Pacitti, Alexandru Costan, Patrick Valduriez, Gabriel Antoniu, and Marta Mattoso. Efficient scheduling of scientific workflows using hot metadata in a multisite cloud. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 31(10):1940–1953, 2018.
107. Lumin Liu, Jun Zhang, SH Song, and Khaled B Letaief. Client-edge-cloud hierarchical federated learning. In *IEEE Int. Conf. on Communications (ICC)*, pages 1–6, 2020.
108. Ruixuan Liu, Yang Cao, Masatoshi Yoshikawa, and Hong Chen. FedSel: Federated sgd under local differential privacy with top-k dimension selection. In *Int. Conf. on Database Systems for Advanced Applications*, pages 485–501, 2020.
109. Yang Liu, Anbu Huang, Yun Luo, He Huang, Youzhi Liu, Yuanyuan Chen, Lican Feng, Tianjian Chen, Han Yu, and Qiang Yang. Fedvision: An online visual object detection platform powered by federated learning. In *AAAI Conf. on Artificial Intelligence*, volume 34, pages 13172–13179, 2020.
110. Yang Liu, Yan Kang, Xinwei Zhang, Liping Li, Yong Cheng, Tianjian Chen, Mingyi Hong, and Qiang Yang. A communication efficient collaborative learning framework for distributed features. *arXiv preprint arXiv:1912.11187*, 2019.
111. Sin Kit Lo, Qinghua Lu, Liming Zhu, Hye-young Paik, Xiwei Xu, and Chen Wang. Architectural patterns for the design of federated learning systems. *arXiv preprint arXiv:2101.02373*, 2021.
112. Siqi Luo, Xu Chen, Qiong Wu, Zhi Zhou, and Shuai Yu. Hfel: Joint edge association and resource allocation for cost-efficient hierarchical federated edge learning. *IEEE Transactions on Wireless Communications*, 19(10):6535–6548, 2020.
113. Xinjian Luo and Xiangqi Zhu. Exploiting defenses against gan-based feature inference attacks in federated learning. *arXiv preprint arXiv:2004.12571*, 2020.
114. Lingjuan Lyu, Han Yu, and Qiang Yang. Threats to federated learning: A survey. *arXiv preprint arXiv:2003.02133*, 2020.
115. Lingjuan Lyu, Jiangshan Yu, Karthik Nandakumar, Yitong Li, Xingjun Ma, Jiong Jin, Han Yu, and Kee Siong Ng. Towards fair and privacy-preserving federated deep models. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 31(11):2524–2541, 2020.
116. Yanjun Ma, Dianhai Yu and Tian Wu, and Haifeng Wang. Paddlepaddle: An open-source deep learning platform from industrial practice. *Frontiers of Data and Computing*, 1(1):105, 2019.
117. Amirhossein Malekijoo, Mohammad Javad Fadaeieslam, Hanieh Malekijou, Morteza Homayounfar, Farshid Alizadeh-Shabdiz, and Reza Rawassizadeh. FEDZIP: A compression framework for communication-efficient federated learning. *arXiv preprint arXiv:2102.01593*, 2021.
118. Kalikinkar Mandal and Guang Gong. PrivFL: Practical privacy-preserving federated regressions on high-dimensional data over mobile networks. In *ACM SIGSAC Conf. on Cloud Computing Security Workshop*, pages 57–68, 2019.
119. Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V. Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R. Savagaonkar. Innovative instructions and software model for isolated execution. In *Int. Workshop on Hardware and Architectural Support for Security and Privacy*, 2013.
120. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, pages 1273–1282, 2017.

121. H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. *arXiv preprint arXiv:1710.06963*, 2017.
122. H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *Int. Conf. on Learning Representations (ICLR)*, 2018.
123. Guangxu Mei, Ziyu Guo, Shijun Liu, and Li Pan. Sgmn: A graph neural network based federated learning approach by hiding structure. In *IEEE Int. Conf. on Big Data (Big Data)*, pages 2560–2568, 2019.
124. Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting unintended feature leakage in collaborative learning. In *IEEE Symposium on Security and Privacy (SP)*, pages 691–706, 2019.
125. Chuizheng Meng, Sirisha Rambhatla, and Yan Liu. Cross-node federated graph neural network for spatio-temporal data modeling. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 2021. To appear.
126. Naram Mhaisen, Alaa Awad, Amr Mohamed, Aiman Erbad, and Mohsen Guizani. Optimal user-edge assignment in hierarchical federated learning based on statistical properties and network topology constraints. *IEEE Transactions on Network Science and Engineering*, 2021.
127. Fan Mo and Hamed Haddadi. Efficient and private federated learning using tee. In *EuroSys*, 2019.
128. Mehryar Mohri, Gary Sivek, and Ananda Theertha Suresh. Agnostic federated learning. In *Int. Conf. on Machine Learning (ICML)*, pages 4615–4625, 2019.
129. Virraji Mothukuri, Reza M Parizi, Seyedamin Pouriyeh, Yan Huang, Ali Dehghantanha, and Gautam Srivastava. A survey on security and privacy of federated learning. *Future Generation Computer Systems*, 115:619–640, 2021.
130. Luis Muñoz-González, Kenneth T Co, and Emil C Lupu. Byzantine-robust federated machine learning through adaptive model averaging. *arXiv preprint arXiv:1909.05125*, 2019.
131. Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R Devanur, Gregory R Ganger, Phillip B Gibbons, and Matei Zaharia. Pipedream: generalized pipeline parallelism for dnn training. In *ACM Symposium on Operating Systems Principles*, pages 1–15, 2019.
132. Keiichi Ochiai, Kohei Senkawa, Naoki Yamamoto, Yuya Tanaka, and Yusuke Fukazawa. Real-time on-device troubleshooting recommendation for smartphones. In *ACM SIGKDD Int. Conf. on Knowledge Discovery & Data Mining*, pages 2783–2791, 2019.
133. Official Journal of the European Union. General data protection regulation. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>. Online; accessed 12/02/2021.
134. Olga Ohrimenko, Felix Schuster, Cédric Fournet, Aastha Mehta, Sebastian Nowozin, Kapil Vaswani, and Manuel Costa. Oblivious multi-party machine learning on trusted processors. In *{USENIX} Security Symposium ({USENIX} Security)*, pages 619–636, 2016.
135. Kemal Oksuz, Baris Can Cam, Sinan Kalkan, and Emre Akbas. Imbalance problems in object detection: A review. *IEEE transactions on pattern analysis and machine intelligence*, 2020.
136. OpenMined. Pysyft. <https://github.com/OpenMined/PySyft>. Online; accessed 22/02/2021.
137. PaddlePaddle, Baidu. Paddlehub. <https://github.com/PaddlePaddle/PaddleHub>. Online; accessed 01/10/2021.
138. Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *Int. Conf. on the theory and applications of cryptographic techniques*, pages 223–238, 1999.
139. Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering (TKDE)*, 22(10):1345–1359, 2009.
140. Hao Peng, Haoran Li, Yangqiu Song, Vincent Zheng, and Jianxin Li. Federated knowledge graphs embedding. In *ACM Int. Conf. on Information and Knowledge Management (CIKM)*, pages 1–10, 2021.
141. Hai Phan, My T Thai, Han Hu, Ruoming Jin, Tong Sun, and Dejing Dou. Scalable differential privacy with certified robustness in adversarial learning. In *Int. Conf. on Machine Learning (ICML)*, pages 7683–7694, 2020.

142. Krishna Pillutla, Sham M Kakade, and Zaid Harchaoui. Robust aggregation for federated learning. *arXiv preprint arXiv:1912.13445*, 2019.
143. Luis Pineda-Morales, Ji Liu, Alexandru Costan, Esther Pacitti, Gabriel Antoniu, Patrick Valduriez, and Marta Mattoso. Managing hot metadata for scientific workflows on multi-site clouds. In *IEEE Int. Conf. on Big Data (Big Data)*, pages 390–397, 2016.
144. Pytorch. Pytorch. <https://pytorch.org/>. Online; accessed 13/03/2021.
145. Herbert Robbins and Sutton Monro. A stochastic approximation method. *The annals of mathematical statistics*, pages 400–407, 1951.
146. Daniele Romanini, Adam James Hall, Pavlos Papadopoulos, Tom Titcombe, Abbas Ismail, Tudor Cebere, Robert Sandmann, Robin Roehm, and Michael A. Hoeh. Pyvertical: A vertical federated learning framework for multi-headed splitnn. *arXiv preprint arXiv:2104.00489*, 2021.
147. Daniel Rothchild, Ashwinee Panda, Enayat Ullah, Nikita Ivkin, Ion Stoica, Vladimir Braverman, Joseph Gonzalez, and Raman Arora. Fetchsgd: Communication-efficient federated learning with sketching. In *Int. Conf. on Machine Learning (ICML)*, pages 8253–8265, 2020.
148. Theo Ryffel, Andrew Trask, Morten Dahl, Bobby Wagner, Jason Mancuso, Daniel Rueckert, and Jonathan Passerat-Palmbach. A generic framework for privacy preserving deep learning. *arXiv preprint arXiv:1811.04017*, 2018.
149. César Sabater, Aurélien Bellet, and Jan Ramon. Distributed differentially private averaging with improved utility and robustness to malicious parties. *arXiv preprint arXiv:2006.07218*, 2020.
150. Adam Satariano. Google is fined \$57 million under europe’s data privacy law. <https://www.nytimes.com/2019/01/21/technology/google-europe-gdpr-fine.html>. Online; accessed 28/02/2021.
151. Ali H Sayed. Adaptation, learning, and optimization over networks. *Foundations and Trends in Machine Learning*, 7(ARTICLE):311–801, 2014.
152. Ali H Sayed, Sheng-Yuan Tu, Jianshu Chen, Xiaochuan Zhao, and Zaid J Towfic. Diffusion strategies for adaptation and learning over networks: an examination of distributed strategies and network behavior. *IEEE Signal Processing Magazine*, 30(3):155–171, 2013.
153. Mohamed Seif, Ravi Tandon, and Ming Li. Wireless federated learning with local differential privacy. In *IEEE Int. Symposium on Information Theory (ISIT)*, pages 2604–2609, 2020.
154. Eugene Seneta. *Non-negative matrices and Markov chains*. Springer Science & Business Media, 2006.
155. William Shakespeare. *The complete works of William Shakespeare*. Wordsworth Editions, 2007.
156. Nir Shlezinger, Mingzhe Chen, Yonina C Eldar, H Vincent Poor, and Shuguang Cui. Federated learning with quantization constraints. In *IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8851–8855, 2020.
157. Nir Shlezinger, Mingzhe Chen, Yonina C Eldar, H Vincent Poor, and Shuguang Cui. Uveqfed: Universal vector quantization for federated learning. *IEEE Transactions on Signal Processing*, 2020.
158. JASON SILVERSTEIN. Hundreds of millions of facebook user records were exposed on amazon cloud server. <https://www.cbsnews.com/news/millions-facebook-user-records-exposed-amazon-cloud-server/>. Online; accessed 28/02/2021.
159. Ryan Spring, Anastasios Kyrillidis, Vijai Mohan, and Anshumali Shrivastava. Compressing gradient optimizers via count-sketches. In *Int. Conf. on Machine Learning (ICML)*, pages 5946–5955, 2019.
160. Standing Committee of the National People’s Congress. Cybersecurity law of the people’s republic of china. <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/translation-cybersecurity-law-peoples-republic-china/>. Online; accessed 22/02/2021.
161. Sebastian U Stich, Jean-Baptiste Cordonnier, and Martin Jaggi. Sparsified SGD with memory. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 31, 2018.
162. Haijian Sun, Xiang Ma, and Rose Qingyang Hu. Adaptive federated learning with gradient compression in uplink noma. *IEEE Transactions on Vehicular Technology*, 2020.
163. Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? *arXiv preprint arXiv:1911.07963*, 2019.

164. Toyotaro Suzumura, Yi Zhou, Nathalie Barcardo, Guangnan Ye, Keith Houck, Ryo Kawahara, Ali Anwar, Lucia Larise Stavarache, Daniel Klyashtorny, Heiko Ludwig, and Kumar Bhaskaran. Towards federated graph learning for collaborative financial crimes detection. *arXiv preprint arXiv:1909.12946*, 2019.
165. Vale Tolpegin, Stacey Truex, Mehmet Emre Gursoy, and Ling Liu. Data poisoning attacks against federated learning systems. In *European Symposium on Research in Computer Security*, pages 480–501. Springer, 2020.
166. Aleksei Triastcyn and Boi Faltings. Federated generative privacy. *IEEE Intelligent Systems*, 35(4):50–57, 2020.
167. Stacey Truex, Nathalie Baracaldo, Ali Anwar, Thomas Steinke, Heiko Ludwig, Rui Zhang, and Yi Zhou. A hybrid approach to privacy-preserving federated learning. In *ACM Workshop on Artificial Intelligence and Security*, pages 1–11, 2019.
168. Paul Vanhaesebrouck, Aurélien Bellet, and Marc Tommasi. Decentralized collaborative learning of personalized models over networks. In *Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*, pages 509–517, 2017.
169. Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *Int. Conf. on Learning Representations (ICLR)*, 2018.
170. Joost Verbraeken, Matthijs Wolting, Jonathan Katzy, Jeroen Kloppenburg, Tim Verbeelen, and Jan S Rellermeyer. A survey on distributed machine learning. *ACM Computing Surveys (CSUR)*, 53(2):1–33, 2020.
171. Abhinav Vishnu, Charles Siegel, and Jeffrey Daily. Distributed tensorflow with mpi. *arXiv preprint arXiv:1603.02339*, 2016.
172. Aidmar Wainakh, Alejandro Sanchez Guinea, Tim Grube, and Max Mühlhäuser. Enhancing privacy via hierarchical federated learning. In *IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 344–347, 2020.
173. Binghui Wang, Ang Li, Hai Li, and Yiran Chen. Graphfl: A federated learning framework for semi-supervised node classification on graphs. *arXiv preprint arXiv:2012.04187*, 2020.
174. Chunnan Wang, Bozhou Chen, Geng Li, and Hongzhi Wang. FL-AGCNS: federated learning framework for automatic graph convolutional network search. *arXiv preprint arXiv:2104.04141*, 2021.
175. Guan Wang. Interpret federated learning with shapley values. *arXiv preprint arXiv:1905.04519*, 2019.
176. Hongyi Wang, Mikhail Yurochkin, Yuekai Sun, Dimitris Papailiopoulos, and Yasaman Khazaeni. Federated learning with matched averaging. In *Int. Conf. on Learning Representations (ICLR)*, 2020.
177. Jianyu Wang, Zachary Charles, Zheng Xu, Gauri Joshi, H. Brendan McMahan, Blaise Agüera y Arcas, Maruan Al-Shedivat, Galen Andrew, Salman Avestimehr, Katharine Daly, Deepesh Data, Suhas N. Diggavi, Hubert Eichner, Advait Gadhihar, Zachary Garrett, Antonious M. Girgis, Filip Hanzely, Andrew Hard, Chaoyang He, Samuel Horvath, Zhouyuan Huo, Alex Ingerman, Martin Jaggi, Tara Javidi, Peter Kairouz, Satyen Kale, Sai Praneeth Karimireddy, Jakub Konečný, Sanmi Koyejo, Tian Li, Luyang Liu, Mehryar Mohri, Hang Qi, Sashank J. Reddi, Peter Richtárik, Karan Singhal, Virginia Smith, Mahdi Soltanolkotabi, Weikang Song, Ananda Theertha Suresh, Sebastian U. Stich, Ameet Talwalkar, Hongyi Wang, Blake E. Woodworth, Shanshan Wu, Felix X. Yu, Honglin Yuan, Manzil Zaheer, Mi Zhang, Tong Zhang, Chunxiang Zheng, Chen Zhu, and Wennan Zhu. A field guide to federated optimization. *arXiv preprint arXiv:2107.06917*, 2021.
178. Jianyu Wang, Anit Kumar Sahu, Zhouyi Yang, Gauri Joshi, and Soumya Kar. Matcha: Speeding up decentralized SGD via matching decomposition sampling. In *Indian Control Conference (ICC)*, pages 299–300, 2019.
179. Lixu Wang, Shichao Xu, Xiao Wang, and Qi Zhu. Addressing class imbalance in federated learning. In *AAAI Conf. on Artificial Intelligence*, volume 35, pages 10165–10173, 2021.
180. Luping WANG, Wei WANG, and LI Bo. CMFL: Mitigating communication overhead for federated learning. In *IEEE Int. Conf. on Distributed Computing Systems (ICDCS)*, pages 954–964, 2019.
181. Zhibo Wang, Mengkai Song, Zhifei Zhang, Yang Song, Qian Wang, and Hairong Qi. Beyond inferring class representatives: User-level privacy leakage from federated learning. In *IEEE Conf. on Computer Communications (INFOCOM)*, pages 2512–2520, 2019.
182. WeBank. Federated ai technology enabler (FATE). <https://github.com/FederatedAI/FATE>. Online; accessed 16/02/2021.

183. WeBank. Federated learning white paper v2.0. https://aisp-1251170195.cos.ap-hongkong.myqcloud.com/wp-content/uploads/pdf/%E8%81%94%E9%82%A6%E5%AD%A6%E4%B9%A0%E7%99%BD%E7%9A%AE%E4%B9%A6_v2.0.pdf. Online; accessed 14/02/2021.
184. Kang Wei, Jun Li, Ming Ding, Chuan Ma, Howard H Yang, Farhad Farokhi, Shi Jin, Tony QS Quek, and H Vincent Poor. Federated learning with differential privacy: Algorithms and performance analysis. *IEEE Transactions on Information Forensics and Security*, 15:3454–3469, 2020.
185. Chuhan Wu, Fangzhao Wu, Yang Cao, Yongfeng Huang, and Xing Xie. Fedgnn: Federated graph neural network for privacy-preserving recommendation. *arXiv preprint arXiv:2102.04925*, 2021.
186. Tong Wu, Ziwei Liu, Qingqiu Huang, Yu Wang, and Dahua Lin. Adversarial robustness under long-tailed distribution. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8659–8668, 2021.
187. Chang Xu, Dacheng Tao, and Chao Xu. A survey on multi-view learning. *arXiv preprint arXiv:1304.5634*, 2013.
188. Jie Xu, Benjamin S Glicksberg, Chang Su, Peter Walker, Jiang Bian, and Fei Wang. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, pages 1–19, 2020.
189. Jinjin Xu, Wenli Du, Yaochu Jin, Wangli He, and Ran Cheng. Ternary compression for communication-efficient federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
190. Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
191. Xun Yi, Russell Paulet, and Elisa Bertino. Homomorphic encryption. In *Homomorphic Encryption and Applications*, pages 27–46. Springer, 2014.
192. Jinliang Yuan, Mengwei Xu, Xiao Ma, Ao Zhou, Xuanzhe Liu, and Shangguang Wang. Hierarchical federated learning through lan-wan orchestration. *arXiv preprint arXiv:2010.11612*, 2020.
193. Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Nghia Hoang, and Yasaman Khazaeni. Bayesian nonparametric federated learning of neural networks. In *Int. Conf. on Machine Learning (ICML)*, pages 7252–7261, 2019.
194. Chengliang Zhang, Suyi Li, Junzhe Xia, Wei Wang, Feng Yan, and Yang Liu. Batchcrypt: Efficient homomorphic encryption for cross-silo federated learning. In *USENIX Annual Technical Conference (USENIX ATC)*, pages 493–506, 2020.
195. Chongsheng Zhang, Jingjun Bi, and Paolo Soda. Feature selection and resampling in class imbalance learning: Which comes first? an empirical study in the biological domain. In *Int. Conf. on Bioinformatics and Biomedicine (BIBM)*, pages 933–938, 2017.
196. Chongsheng Zhang, Jingjun Bi, Shixin Xu, Enislay Ramentol, Gaojuan Fan, Baojun Qiao, and Hamido Fujita. Multi-imbalance: An open-source software for multi-class imbalance learning. *Knowledge-Based Systems*, 174:137–143, 2019.
197. Chongsheng Zhang, Paolo Soda, Jingjun Bi, Gaojuan Fan, George Almpandis, and Salvador Garcia. An empirical study on the joint impact of feature selection and data resampling on imbalance classification. *arXiv preprint arXiv:2109.00201*, 2021.
198. Huanding Zhang, Tao Shen, Fei Wu, Mingyang Yin, Hongxia Yang, and Chao Wu. Federated graph learning - A position paper. *arXiv preprint arXiv:2105.11099*, 2021.
199. Tuo Zhang, Chaoyang He, Tianhao Ma, Mark Ma, and Salman Avestimehr. Federated learning for internet of things: A federated learning framework for on-device anomaly data detection. *arXiv preprint arXiv:2106.07976*, 2021.
200. Xiaoli Zhang, Fengting Li, Zeyu Zhang, Qi Li, Cong Wang, and Jianping Wu. Enabling execution assurance of federated learning at untrusted participants. In *IEEE INFOCOM Conf. on Computer Communications*, pages 1877–1886, 2020.
201. Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. idlg: Improved deep leakage from gradients. *arXiv preprint arXiv:2001.02610*, 2020.
202. Yuchen Zhao, Payam Barnaghi, and Hamed Haddadi. Multimodal federated learning. *arXiv preprint arXiv:2109.04833*, 2021.
203. Longfei Zheng, Jun Zhou, Chaochao Chen, Bingzhe Wu, Li Wang, and Benyu Zhang. Asfgnn: Automated separated-federated graph neural network. *Peer-to-Peer Networking and Applications*, 14(3):1692–1704, 2021.
204. Jun Zhou, Chaochao Chen, Longfei Zheng, Huiwen Wu, Jia Wu, Xiaolin Zheng, Bingzhe Wu, Ziqi Liu, and Li Wang. Vertically federated graph neural network for privacy-preserving node classification. *arXiv preprint arXiv:2005.11903*, 2020.

-
205. Hangyu Zhu, Haoyu Zhang, and Yaochu Jin. From federated learning to federated neural architecture search: a survey. *Complex & Intelligent Systems*, 2021.
 206. Martin Zinkevich, Markus Weimer, Alexander J Smola, and Lihong Li. Parallelized stochastic gradient descent. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 4, page 4. Citeseer, 2010.