From Euclidean to Graph-Structured Data: A Survey of Collaborative Learning

Anonymous authors
Paper under double-blind review

Abstract

The conventional approach to machine learning, that is, collecting data, training models, and performing inference in a single location, faces fundamental limitations, including scalability and privacy, that restrict its applicability. To address these challenges, recent research has explored *collaborative learning* approaches, including *federated learning* and *decentralized learning*, where individual agents perform training and inference locally, with limited collaboration. Most collaborative learning research focuses on Euclidean data with regular, grid-like structure (e.g., images, text). However, these approaches fail to capture the relational patterns in many real-world applications, best represented by graphs. Learning on graphs relies on message-passing mechanisms to propagate information between connected nodes, making it conceptually well-suited for collaborative environments where agents must exchange information. Yet, the opportunities and challenges of learning on graph-structured data in collaborative settings remain largely underexplored. This survey provides a comprehensive investigation of collaborative learning from Euclidean to graph-structured data, aiming to consolidate this emerging field. We begin by reviewing its foundational principles for Euclidean data, organizing them along three core dimensions: learning effectiveness, efficiency, and privacy preservation. We then extend the discussion to graph-structured data, introducing a taxonomy of graph distribution scenarios, characterizing associated statistical heterogeneities, and developing standardized problem formulations and algorithmic frameworks. Finally, we systematically identify open challenges and promising research directions. By bridging established techniques for Euclidean data with emerging methods for graph learning, our survey provides researchers and practitioners with a well-structured foundation of collaborative learning, supporting further development across a wide range of scientific and industrial fields [1].

[1] Resources will be available after acceptance on GitHub.

1 Introduction

Over the past decade, Deep Learning has turned into a success story, largely due to the combination of four factors (Goodfellow et al., 2016): (i) neural networks as *universal* function approximators (Rosenblatt, 1958; Rumelhart et al., 1986) (ii) *sample efficient* training algorithms (Bottou, 2010), (iii) *parallelization* schemes that align with hardware architectures (Hooker, 2021); and (iv) the availability of *massive* amounts of high-quality data.

While the first three factors have been consolidated over the years, the fourth requirement, *access to massive datasets*, remains a major obstacle. Indeed, today data ecosystem remains highly fragmented, constrained by privacy legislation (European Union, 2016), data ownership concerns, integration complexity, and the inherent coordination difficulties of multi-agent systems. This access to data is partially addressed for LLMs by scaling on *public* web data (Naveed et al., 2023; Baack, 2024), but the fundamental challenge persists: the absence of mechanisms to collect, align, and utilize data across organizational and technical boundaries.

Collaborative learning, often termed under the umbrella of *federated* or *decentralized learning*, provides a framework to address this challenge by allowing multiple agents to train a shared model directly on their private datasets. In this framework, only model updates or aggregated information are exchanged between agents rather than entire datasets, drastically reducing the amount of data communicated. Since the data remains at its source, the framework enhances privacy through data minimization and anonymization (Bonawitz et al., 2022; Daly et al., 2024), enabling learning from private datasets in compliance with modern data protection frameworks (European Union, 2016).

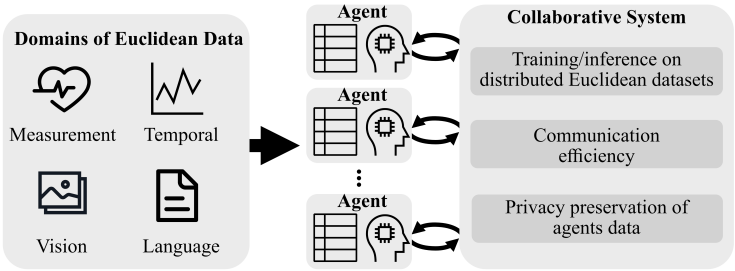

(a) Collaborative Learning on Euclidean Data

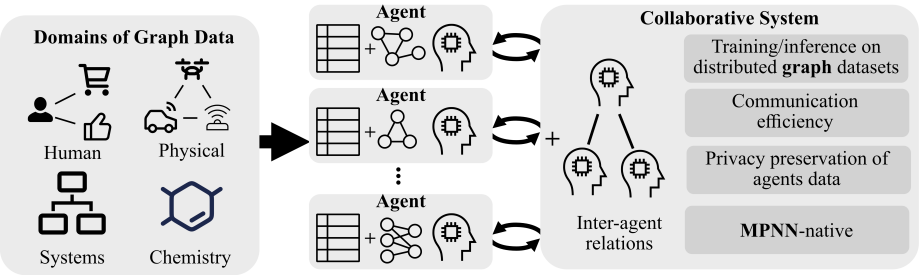

(b) Collaborative Learning on Graph Data

Figure 1: **Collaborative Learning from Euclidean to Graph Data.** Collaborative learning refers to a system of agents that collectively train and infer ML models on their local private data. Collaborative learning has been extensively studied for Euclidean data (a), comprising highly popular domains like vision, language, temporal sequences, and measurement data. Conversely, collaborative learning on graph data (b) is an emerging field. Graphs canonically represent data structured with relations, arising in domains including human behavior, physical systems, chemistry, and information systems. Compared to *centralized learning*, collaborative learning offers three main advantages: enabling training and inference on distributed datasets, communication efficiency, and privacy preservation of agents data. Collaborative learning on graph data aligns naturally with *Message Passing Neural Networks (MPNNs)*.

Collaborative learning solutions can be divided into two categories based on data structure: *Euclidean data*, which exhibits regular, grid-like structure (e.g., images as 2D grids, text as 1D sequences), and *non-Euclidean data*, where relationships are defined by irregular structures such as *graphs* (Bronstein et al., 2021), as illustrated in Figure 1.

For Euclidean data, typical use cases include cross-device scenarios involving numerous agents, each possessing limited data insufficient for generating accurate models independently. These agents may include mobile phones (Xu et al., 2023), autonomous vehicles (Hellström et al., 2022), mobile base stations (Zhang et al., 2022), or Internet of Things (IoT) devices (Zhou et al., 2023). Cross-silo applications involve large datasets held by distinct actors who cannot share data to preserve privacy, prevalent in healthcare and finance (Li et al., 2020). Collaboration typically involves exchanging and aggregating model parameters during training, while inference is generally performed locally by individual agents.

For graph-structured data, two distinct scenarios emerge. The first scenario is learning on a large collection of independent graph instances, typical in computational biology (e.g., predicting molecular properties) (He et al., 2022), involves learning parameters of a graph-based model (e.g., a GNN) using tools conceptually similar to those for Euclidean data, as each graph instance can be treated analogously to a single data sample. The second scenario arises when agents hold a portion of a *shared* global graph, both inference and learning necessitate information exchange among agents, as the graph topology encodes critical relationships spanning agent boundaries. Applications include collaborative autonomous agents in IoT systems (Liu et al., 2025), autonomous vehicles (Blumenkamp et al., 2022; Pan et al., 2023), and privacy-preserving learning over large-scale social network or recommender system graphs (He et al., 2021; Dong et al., 2023; Agrawal et al., 2024; Han et al., 2024).

While these diverse applications demonstrate broad potential, their implementation presents significant challenges that differ fundamentally between Euclidean and non-Euclidean settings. Modern Machine Learning (ML) relies predominantly on iterative optimization algorithms (Bottou et al., 2018), which have long posed difficulties for distributed systems (Bertsekas, 2011; Ram et al., 2009; Low et al., 2010; Xing et al., 2016). These challenges are exacerbated in collaborative learning due to the substantial size of the models that need to be communicated and inherent system discrepancies, driving active research in system architecture and algorithm design (Kairouz et al., 2021; Daly et al., 2024).

For graph-structured data, the challenges are further amplified by the non-Euclidean nature of the data. The graph topology is often captured by *Message Passing Neural Networks* (Gilmer et al., 2017; Hamilton et al., 2017), where the nodes iteratively exchange and aggregate messages with their neighbors. These frameworks are inherently distributed iterative algorithms, which makes them conceptually well-suited for collaborative settings. However, when each agent holds a partition of the shared global graph and cannot freely share data, the requirement for information to propagate through the graph topology creates novel challenges beyond those in Euclidean settings, leading to a rapidly expanding area of research.

1.1 The Objectives of the Survey and Related Work

While existing surveys have extensively covered *federated learning* on Euclidean data, and to a lesser extent on graph data, none have systematically analyzed their intersection. This survey bridges that gap by providing a unified view that maps the design space of collaborative learning from Euclidean to graph-structured data. We outline how different approaches relate to each other and identify research challenges that remain unanswered for the less understood case of graphs. Crucially, while methods for Euclidean data are well-established, graph-structured data introduces unique complexities that draw upon knowledge from multiple domains. Collaborative learning for graph data stands on three foundational pillars, borrowing from distinct research areas as illustrated in Figure 2:

- **Collaborative Learning**, for example *federated learning* (McMahan et al., 2017) allows privacy-preserving training and inference on private distributed datasets through centralized or decentralized approaches.
- **Graph Representation Learning** focuses on learning representations from graph-structured data typically employing deep neural network architectures (Kipf & Welling, 2017).

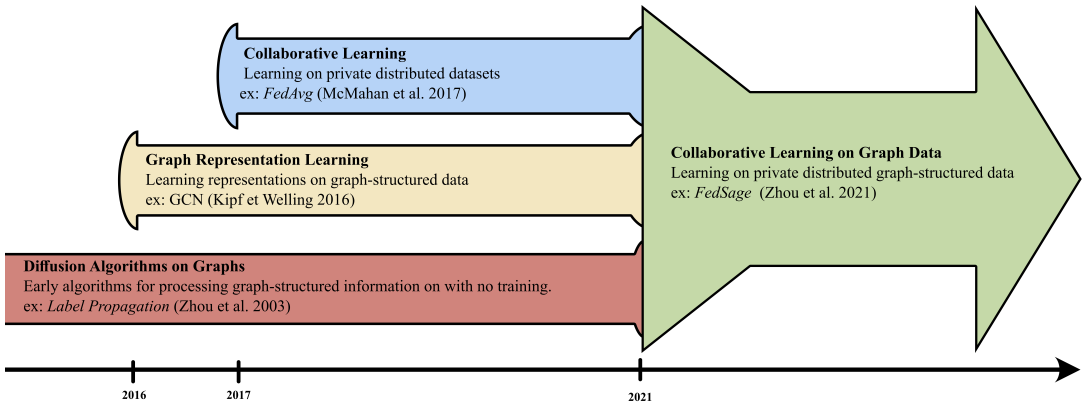

Figure 2: **Collaborative Learning on graph data and its three pillars.** Collaborative learning on graph data relies on three distinct pillars: diffusion algorithm on graphs, graph representation learning, and collaborative learning.

- **Diffusion algorithms on graphs** possess an extensive history, beginning with seminal works on distributed averaging (Tsitsiklis et al., 1986) and other elementary function computations performed among distributed agents communicating via communication graphs or peer-to-peer networks (Olfati-Saber et al., 2007), and further extended for solving semi-supervised tasks on graphs (Zhou et al., 2003).

This survey connects these three domains to examine how established results can accelerate research in emerging fields and novel applications. Table 1 presents representative surveys and foundational papers across these pillars, providing comprehensive details that were necessarily excluded to maintain focus. The table also includes works on parallel processing of big data, an area with partial overlap with federated learning. To limit the scope of the survey, we focus on supervised learning. Interested readers may see the surveys (Gronauer & Diepold, 2022; Chen et al., 2025) on reinforcement learning in distributed environments.

1.2 The Organization of the Paper

In this survey, we provide a structured overview of learning from distributed private datasets, including the basic concepts, the research challenges, and proposed solutions, starting from classical collaborative learning, to learning over distributed private graph datasets. The skeleton of the paper is shown in Figure 3. In the first part of the paper, we discuss collaborative learning over Euclidean data, that is, classical centralized FL and its distributed counterparts. We formulate the learning problem, discuss how the data can be partitioned across the agents, and what heterogeneity aspects need to be considered. Then we discuss proposed solutions (i) for effective learning, that is, to achieve models with high accuracy despite the various forms of heterogeneity, (ii) for efficiency, in terms of the use of communication and computation resources, and finally, (iii) solutions for privacy preservation. Since collaborative learning over Euclidean data is an established research area, we review the key contributions in this part.

In the second part of the paper, we follow the same structure to discuss the emerging research field of collaborative learning on graph-structured data. The ways data (or information) can be distributed are significantly more varied now than in the Euclidean case, so we introduce a taxonomy for data partition. Then, we follow the structure of the first part to find relevant research results and questions that are still open for the design of collaborative learning system on graph data. This part is based on a rigorous survey of all recent papers with keywords *federated graph* and *decentralized graph learning*. The notations and the glossary of key terminology used throughout this survey can be found in Appendix A.

Topic	Representative Works
Collaborative Learning on Euclidean Data	
Federated Learning	(Kairouz et al., 2021; Zhang et al., 2021; Yang et al., 2019; Liu et al., 2022)
Vertical Federated Learning	(Liu et al., 2024)
Data heterogeneity and personalization	(Zhu et al., 2021; Tan et al., 2022; Liu et al., 2024)
Privacy and security	(Lyu et al., 2020; Daly et al., 2024)
Parallel training	(Xing et al., 2016)
Machine Learning in wireless networks	(Hellström et al., 2022)
Diffusion algorithms on graphs	
Distributed averaging and consensus	(Bertsekas & Tsitsiklis, 2015; Olfati-Saber et al., 2007)
Label propagation	(Zhou et al., 2003; Zhu & Ghahramani, 2002)
Graph Representation Learning	
Graph Neural Architectures	(Wu et al., 2020)
Graph Foundation Models	(Wang et al., 2025)
Collaborative Learning on graph data	
Federated Learning with GNNs	(He et al., 2021; Zhang et al., 2021; Liu et al., 2025; Fu et al., 2024)
Applications and use cases	(He et al., 2021; Dong et al., 2023)
GNNs in wireless communications	(He et al., 2021; Lee et al., 2022)
Parallel training for large GNNs	(Shao et al., 2024; Lin et al., 2023; Liu et al., 2022)

Table 1: Key surveys and seminal works on collaborative learning on graph data, including the three pillars of *Collaborative Learning*, *Diffusion algorithms on graph-structued data*, and *Graph Representation Learning*.

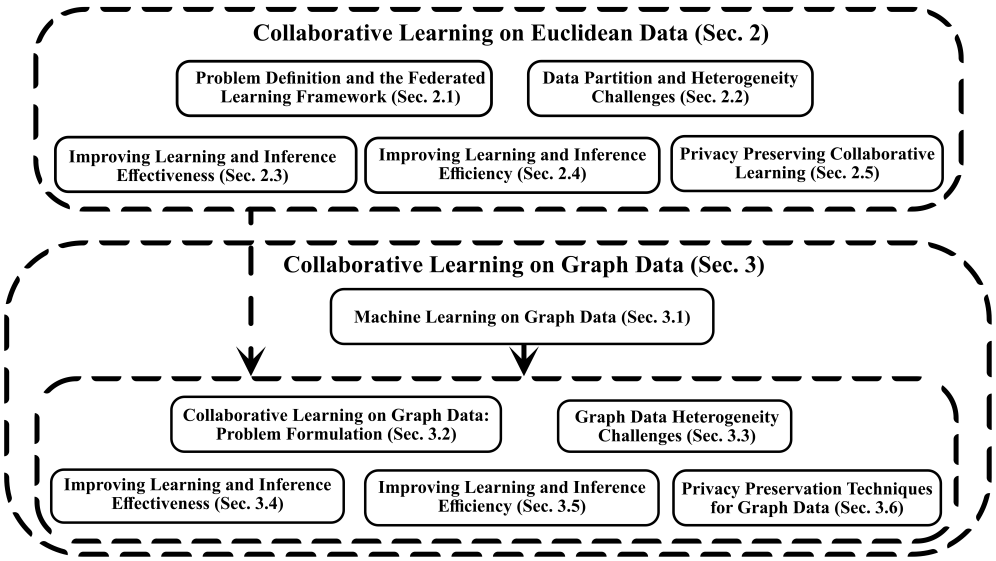

Figure 3: The structure of the paper. Our contributions reside in extending and mapping the concepts found for collaborative learning on Euclidean data to collaborative learning on graph data, thereby consolidating insights across these separate research areas. Our analysis of the design choices focuses on three directions: improving learning and inference *effectiveness*, improving learning and inference *efficiency*, and *privacy*-preservation techniques.

2 Collaborative Learning on Euclidean Data

The concept of *Collaborative Learning* emerges from multiple taxonomies, which are all cast under the umbrella term of *distributed Machine Learning* as outlined in Figure 4. The term *distributed machine*

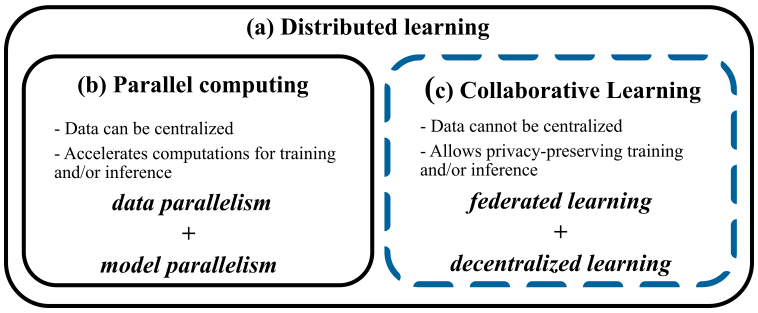

Figure 4: Taxonomies of distributed learning and the **scope of our survey**. *Distributed learning* (a) is often used as an umbrella term encompassing various scenarios and techniques. In particular, *parallel computing* (b) aims to accelerate training and/or inference under the assumption that the data can be centralized, and includes both *data-parallelism* and *model-parallelism* approaches. Our focus, however, is on a different scenario that we term *collaborative learning* (c), which addresses cases where the data remains distributed and private to its owner, covering both *federated learning* and *decentralized learning*.

learning originally served as a broad label for methods designed to train and/or infer ML models on massive datasets by distributing computations across multiple machines, often in a cluster (Xing et al., 2016; Kairouz et al., 2021). Distribution is then largely associated with parallelization, combining *data parallelism*, where datasets are divided among machines, and *model parallelism*, where parts of the model are divided among machines and updated in parallel.

The term *collaborative learning* refers specifically to scenarios where local data remains strictly private to its owner. For clarity, we assume that local data is generated locally at multiple agents, where it persists throughout training and inference phases while remaining exclusively accessible to the data owner, a case commonly termed *federated learning* or *decentralized learning* in the literature. This privacy constraint necessitates precise problem formulation and distinct design methodologies to address real-world implementation challenges, which we develop in this section. Our focus here is on *Euclidean data*, defined as data where each sample is represented as a feature vector in an Euclidean space (in practice some $\mathbb{R}^d$ with the standard inner product).

2.1 Collaborative Learning: Problem Formulation and the Federated Learning Framework

Problem definition. Consider a set of agents, denoted by $\mathcal{K} = \{1, \ldots, K\}$, as illustrated in Figure 5, where the data of each agent is sampled from an unknown distribution $\mathcal{D}_k$ on $\mathcal{X} \times \mathcal{Y}$ where $\mathcal{X}$ is the feature Euclidean domain and $\mathcal{Y}$ the label domain. Given a loss function $\ell : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}^+$, each agent seeks to learn a mapping $F : \mathcal{X} \to \mathcal{Y}$ among a class of function $\mathcal{F}$ that minimizes the local population loss

$$\mathcal{L}_k(F) := \mathbb{E}_{(x,y)\sim\mathcal{D}_k} \ell(F(x), y). \tag{1}$$

If the data distributions $\mathcal{D}_k$ were known, agents could determine an optimal solution $F_k^* \in \mathcal{F}$, by directly minimizing their respective local loss (1) on $\mathcal{F}$. However, in practice, the true distributions $\mathcal{D}_k$ are unknown. Instead, agents rely on empirical samples forming local datasets S_k, drawn i.i.d. from $\mathcal{D}_k$. These datasets allow agents to construct a surrogate $\hat{\mathcal{L}}_k$, which represents an empirical approximation of the true loss $\mathcal{L}_k$ based on their local data only

$$\hat{\mathcal{L}}_k(F) := \frac{1}{|S_k|} \sum_{(x,y)\in S_k} \ell(F(x), y). \tag{2}$$

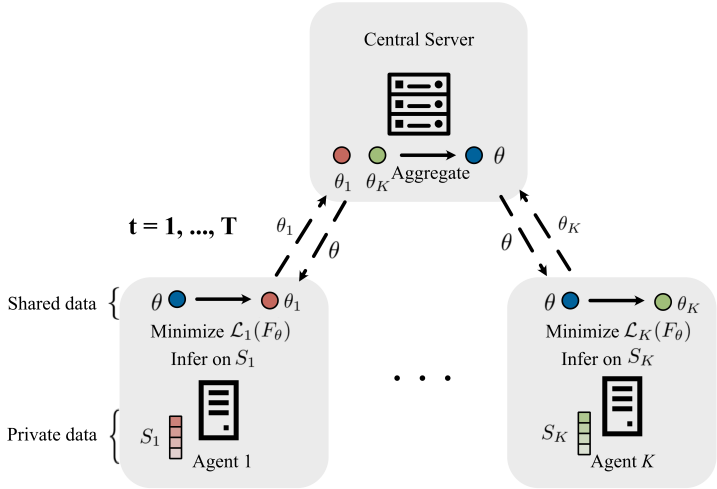

Figure 5: *Federated Learning*, the basic approach to collaborative learning represented here for K agents. Agents possess private local datasets S_k. In each round t they train their local model F_θ on their local loss and communicate the model parameter updates $\theta_k(t)$. The central server aggregates in turn the local models into a global model with parameters $\theta(t)$, which is then shared again with the agents for another round $t+1$.

This approach provides a local estimate $\hat{F}_k^* \in \mathcal{F}$. However, the local datasets S_k are typically too small to accurately approximate the underlying distribution $\mathcal{D}_k$. As a result, the local estimate $\hat{F}_k^*$ poorly approximates the true optimal function F_k^*, leading to high estimation error (Valiant, 1984; Bottou & Bousquet, 2007). To address this issue, a key assumption is that although the distributions $\mathcal{D}_k$ may differ across agents, they share strong similarities. Therefore, collaboration among agents enables them to benefit from the aggregated knowledge contained in the union of local datasets $\{S_k\}_{k\in\mathcal{K}}$.

In the most favorable case, when the local distributions D_k are the same, an effective solution is to collaborate on building **a single model common to all agents**. The problem can then be framed as identifying $F^* \in \mathcal{F}$ that minimizes the average of local population losses among all agents

$$\mathcal{L}(F) := \frac{1}{K}\sum_{k=1}^{K} \mathcal{L}_k(F). \tag{3}$$

Since the population losses $\mathcal{L}_k$ are generally uncomputable, (3) is typically approximated using local empirical losses. Most existing work focuses on minimizing (4), which is computationally feasible, rather than the true loss

$$\hat{\mathcal{L}}(F) := \frac{1}{K}\sum_{k=1}^{K} \hat{\mathcal{L}}_k(F). \tag{4}$$

This line of work typically assumes that $\hat{F}^* \approx F^*$,that is, minimizing (4) provides a good approximation of minimizing (3). This assumption holds only when the generalization gap $|\mathcal{L}(\hat{F}^*) - \hat{\mathcal{L}}(\hat{F}^*)|$ is negligible (Bottou et al., 2018).

To minimize (4), it can be convenient to parametrize F_θ with $\theta \in \Theta$ where Θ is the domain parameter. Typically, F can be a neural network, parametrized by its learnable weights θ. The problem can be framed as identifying θ^*, a solution of

$$\min_{\theta\in\Theta} \hat{\mathcal{L}}(F), \text{ or equivalently, } \quad \min_{\theta\in\Theta} \frac{1}{K}\sum_{k=1}^{K} \hat{\mathcal{L}}_k(F_\theta). \tag{5}$$

Federated learning framework. The basic approach for collaborative learning on distributed datasets is the *Federated Learning (FL)* framework (Konečný et al., 2016; McMahan et al., 2017; Kairouz et al., 2021).

The key idea of Federated Learning is to *federate* (McMahan et al., 2017) the training of agents through a common round-based protocol coordinated by a central server, as shown on Figure 5. At the end of training, each agent receives the jointly trained model, which can be used to perform inference independently on local data.

Federated Averaging (FedAvg) protocol. The most popular and simplest implementation of FL is *Federated Averaging (FedAvg)* (McMahan et al., 2017) detailed in Algorithm 1. It proceeds in T communication rounds between a leader, referred to as the server, and the agents. At the beginning of each round t, the server randomly selects $K_{\max}$ agents (line 2) and communicates the latest global model parameters $\theta(t-1)$ to them (line 3). Upon reception, each selected agent initializes its temporary local model $\theta_k(t-\frac{1}{2})$ to the received global model $\theta(t-1)$. Then, each agent k performs E local update steps (7) on its local dataset S_k to minimize its empirical loss (2) using the *Stochastic Gradient Descent (SGD)* algorithm (Robbins & Monro, 1951). After completing the local updates, each agent sends its resulting local model $\theta_k(t)$ back to the server (line 10). The server then aggregates all received models to compute the updated global model $\theta(t)$ via averaging, denoted by $\bar{\theta}(t)$ (11). Under specific conditions, this iterative process converges to θ^* a solution of (5) (Yu et al., 2019; Stich, 2019).

Algorithm 1 Federated Averaging (FedAvg) adapted from McMahan et al. (2017)

Require: Set of agents $\mathcal{K} = \{1, \ldots, K\}$, local datasets S_k, total number of communication rounds T, number of local training iterations E, number of sampled clients K_{max}, initial global model $\theta(0)$, learning rate η

1: **for** each round $t = 1, \ldots, T$ **do**

Server:

2: Sample uniformly a random subset $\mathcal{K}(t) \subseteq \mathcal{K}$ of m clients ($|\mathcal{K}(t)| = \mathcal{K}_{max}$)

3: Distribute global model parameters $\theta(t-1)$ to all sampled clients $k \in \mathcal{K}_t$

4: **for** each agent $k \in \mathcal{K}(t)$ **in parallel do**

5: Upon reception of $\theta(t-1)$, initialize local model parameter $\theta_k(t-\frac{1}{2}) = \theta(t-1)$

6: **for** $e = 1, \ldots, E$ **do**

7: Update local model using gradient descent [2]:

$$\theta_k(t-\frac{1}{2}) \leftarrow \theta_k(t-\frac{1}{2}) - \eta \nabla \left(\sum_{(x,y)\in S_k} \ell(F_{\theta_k(t-\frac{1}{2})}(x), y \right)$$

8: **end for**

9: **end for**

10: Communicate local model $\theta_k(t) = \theta_k(t-\frac{1}{2})$ to the server

Server:

11: Upon reception of models $\{\theta_k(t)\}_{k\in\mathcal{K}(t)}$, average local models:

$$\bar{\theta}(t) = \sum_{k\in\mathcal{K}(t)} \frac{n_k}{n} \theta_k(t)$$

12: Set global model: $\theta(t) \leftarrow \bar{\theta}(t)$

13: **end for**

14: **return** Final global model $\theta(T)$

This basic approach for FL builds on several unrealistic assumptions. In the sequel, we present the characteristics of collaborative learning that need to be considered to address realistic use cases.

[2] In practice and as proposed in the original algorithm(McMahan et al., 2017), the dataset S_k is divided into batches of fixed size and the model is updated with the gradient on each batch of data, elided here for simplicity.

2.2 Data Partition and Heterogeneity Challenges

The first deviation from the ideal assumptions outlined in Section 2.1 concerns data-related aspects, notably the partitioning of data across agents and the statistical imbalance.

2.2.1 Data Partition

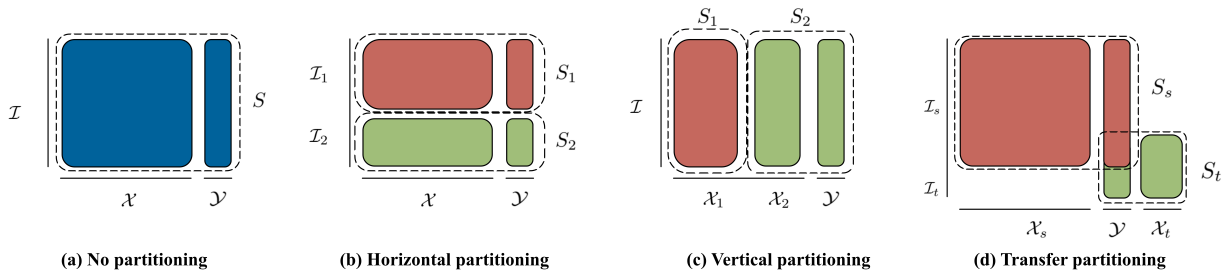

Figure 6: The different partitions of Euclidean data for a two-agent case, where S_k denotes the dataset of agent k. Compared to *no partitioning* (a), *horizontal partitioning* (b) splits the sample space $\mathcal{I}$, *vertical partitioning* (c) splits the feature space $\mathcal{X}$, and *transfer partitioning* (d) splits both, with the source agent s holding most of the data.

FL assumes a significant shared knowledge among the agents through the homogeneity of the local datasets. However, in practice, local datasets $\{S_k\}_{k\in\mathcal{K}}$ often vary significantly across the agents, with heterogeneity arising from the partition of the data or statistical imbalance. An example of these partitions for two agents is shown in Figure 6, with a summary provided in Table 2. In the native scenario (Section 2.1), agents collect similar types of features and labels, represented by a common feature and label spaces $\mathcal{X} \times \mathcal{Y}$. However, inconsistency in the features and labels collected is frequent in data collection practices. While these challenges have long been recognized in centralized ML (Pan & Yang, 2009; Zhu, 2005), they are amplified in collaborative settings where the dataset S, that would ideally be centrally available, is instead partitioned across agents in local datasets $\{S_k\}_{k\in\mathcal{K}}$ which stay private to each agent. To formalize, we assume the existence of a global dataset

$$S = \left\{(x^i, y^i)\right\}_{i\in\mathcal{I}}.$$

where $\mathcal{I}$ is the global sample index, and each agent $k \in \mathcal{K}$ holds a private partition S_k of S. Different partitions exist for S:

- *Horizontal partition.* Under horizontal partition, the index $\mathcal{I}$ of S is partitioned across agents into local indexes $\mathcal{I}_k \subset \mathcal{I}$ such as $\cup_{k\in\mathcal{K}}\mathcal{I}_k = \mathcal{I}$

  $$S_k = \left\{(x^i, y^i)\right\}_{i\in\mathcal{I}_k}.$$

  Therefore, agents possess data about distinct samples indexed by $\mathcal{I}_k$, but with the same feature $\mathcal{X}$ and label space $\mathcal{Y}$. This is the default setting of FL presented in Algorithm 1, and the most studied partition (Kairouz et al., 2021) in the FL literature. This partition further distinguishes *cross-device* scenarios, where millions of agents with limited data (e.g., mobile devices) collaborate, from *cross-silo* scenarios, where a few institutions holding vast amounts of data "silos" (e.g., banks, hospitals, e-commerce) collaborate.

- *Vertical partition.* Under vertical partitions, the domains $\mathcal{X}$ and $\mathcal{Y}$, usually some Euclidean vectorial spaces, are partitioned across agents into local feature $\mathcal{X}_k$ and label $\mathcal{Y}_k$ domains such that $\mathcal{X} = \oplus_{k\in\mathcal{K}}\mathcal{X}_k$ and $\mathcal{Y} = \oplus_{k\in\mathcal{K}}\mathcal{Y}_k$. Consequently, each sample (x^i, y^i) is partitioned into $(x^i_k, y^i_k) \in \mathcal{X}_k \times \mathcal{Y}_k$ to form the local datasets

  $$S_k = \left\{(x^i_k, y^i_k)\right\}_{i\in\mathcal{I}}.$$

Therefore, agents observe the same samples but with distinct views, as represented by different feature and label domains (Liu et al., 2024). This naturally arises in multi-source scenarios where agents collect features and labels specific to their source (e.g., a hospital will diagnose respiratory diseases using lung radiography techniques, while another will use respiratory tests)

A special case of vertical partition is the *transfer partition*. It occurs when one or more *source* agents (e.g., agent $s \in \mathcal{K}$) hold a dataset from one feature space $\mathcal{X}_s$, while another *target* agent (e.g., agent $t \in \mathcal{K}$) holds a dataset from a different feature space $\mathcal{X}_t$. For simplicity, we describe the case involving two agents s and t

$$\forall k \in \{s, t\}, \quad S_k = \left\{(x^i, y^i) \in \mathcal{X}_k \times \mathcal{Y}_k\right\}_{i \in \mathcal{I}_k}.$$

Despite differences in feature domains, the agents may share a subset of common samples

$$\forall k \in \{s, t\}, \quad \left\{(x_k^i, y_k^i)\right\}_{i \in \mathcal{I}_s \cap \mathcal{I}_t} \subset S_k.$$

These act as anchor points for transferring information between agents. Transfer partitions often involve significant data imbalance, where $|\mathcal{I}_s| \gg |\mathcal{I}_t|$, consistent with the idea of "transferring" substantial knowledge from s to t. This setting frequently arises when organizations employ different data acquisition modalities but share overlapping entities. For example, a hospital with a large archive of traditional diagnostic records may support another institution, using a small set of shared patient samples.

2.2.2 Statistical Imbalance

In the native model, local datasets S_k are assumed to be of the same size and sampled from a shared distribution $\mathcal{D}$ on $\mathcal{X} \times \mathcal{Y}$. However, in practice, the statistical properties of the local datasets can vary significantly from one agent to another, leading to different modes of imbalance (Kairouz et al., 2021; Paulik et al., 2021).

Quantity imbalance. The most obvious phenomenon comes from the *quantity imbalance* ($|\mathcal{S}_k| \gg |\mathcal{S}_l|$) where the number of samples collected locally varies significantly across the agents. For example, a small clinic might manage thousands of patient records, while a multinational institution handles millions. This quantity imbalance is a natural phenomenon observed in many fields (Nisonger, 2008) where most of the world's data is held by a minority of data holders. The significant imbalance has given rise to different regimes known as *cross-device*, where millions of agents usually corresponding to users devices (e.g. a phone) collaborate, and *cross-silo*, where a few agents holding vast silos of data (e.g. an hospital) data collaborate.

Distributional shift. In addition, statistical imbalance in local datasets can arise from the distribution itself, where the datasets are drawn. Factors like population bias, dynamic environments, multiple modalities, and varying labeling standards can influence a shift in the marginal distribution $\mathcal{D}_k(x, y)$ from one agent to another, known as *distributional shift* (Moreno-Torres et al., 2012; Quiñonero-Candela et al., 2022; Kairouz et al., 2021; Zhu et al., 2021). We distinguish:

- *Covariate shift* occurs when the feature $\mathcal{D}_k(x)$ or label distribution $\mathcal{D}_k(y)$ observed locally, differ across agents but the concept $\mathcal{D}(x \mid y)$ or $\mathcal{D}(y \mid x)$ remains common. Covariate shift is a ubiquitous phenomenon in statistics dating back to early statistical studies (Snow, 1855) and is the most studied case for FL (Kairouz et al., 2021; Zhu et al., 2021; McMahan et al., 2017; Hsieh et al., 2020; Li et al., 2022). This often arises from population selection bias or non-stationary environments (Moreno-Torres et al., 2012). This distributional shift can be observed from two different perspectives. Firstly, differences in observed label distribution $\mathcal{D}_k(y)$, known as *label shift*, lead to to differences in observed features distribution $\mathcal{D}_k(x) = \mathcal{D}(x \mid y).\mathcal{D}_k(y)$. For example, a hospital located in a polluted area records a higher frequency of respiratory disease diagnoses compared to others in a clean environment. Similarly, a difference in observed feature distributions $D_k(x)$, known as *feature shift*, leads to a difference in label distribution $\mathcal{D}_k(x) = \mathcal{D}(y \mid x).\mathcal{D}_k(x)$. For example, one hospital

focuses on diagnosing a young population while another an older population, leading to differences in features collected, yet the diagnostic knowledge remains the same.

- *Concept shift* occurs when the when the conditional distribution $\mathcal{D}_k(y \mid x)$ (or $\mathcal{D}_k(x \mid y)$) varies across agents. This corresponds to cases where the understanding of the concept varies from one agent to another. For example, the standards used for diagnosing hypertension from blood pressure differ between hospitals. This scenario is rather rare (Moreno-Torres et al., 2012), and challenging for most applications that assume shared mapping $F \in \mathcal{F}$ from features to labels (Huang et al., 2021).

2.2.3 System Heterogeneity

The native forms of FL, such as *FedAvg* in Algorithm 1, work best if the capabilities and tasks of the agents are similar. In practice, real-world systems include a variety of devices, networks, and tasks, and these system heterogeneities affect the learning and inference performance. We briefly characterize the main types of system heterogeneity that impact collaborative training and inference.

Device and network constraints. Agent computations are performed by various physical machines with varying capabilities. These machines range from IoT sensors and mobile devices to telecom base stations, small containers with cloud access, powerful servers, and even supercomputers. This diversity introduces differences in computational and communication capacities (Xie et al., 2020; Even et al., 2024).

The available computing resources define the time needed to calculate the local model updates. As the computing capability of the devices can differ in order of magnitudes, the computation times may also differ significantly. The computing capabilities may also vary in time, due to parallel computations, or varying energy resources.

The agents may use different network technologies (fiber, wifi, 5G, low-power IoT). This defines the *communication topology*, typically a star with a central base station, or a mesh, with direct communication between pairs of agents. The network technology, in addition, defines the *achievable bitrate*, and consequently the transmission delays. When agents utilize wireless links, the quality of these links varies due to interference and noise, which introduces temporal variations and affects the reliability of the communication. Finally, the topology, the reliability, and the transmission rate of the network may be evolving dynamically, due to agents joining, leaving, or simply moving (Hellström et al., 2022; Wang et al., 2020).

This heterogeneity in the *computation* and *communication* capabilities affects both the agents and the system as a whole. For individual agents, these variations affect resource consumption costs, such as energy, time, and communication overhead, when participating in the collaborative learning process (Wang et al., 2022; Meng et al., 2025). On the system level, slower agents, so-called *stragglers*, can delay the entire learning process, creating bottlenecks or biases that influence all participants (Li et al., 2022).

Task heterogeneity. The native model assumes that all agents share the same task, represented as a common mapping from the feature domain $\mathcal{X}$ to the label domain $\mathcal{Y}$. In practice, however, agents may use their data to perform different tasks, thereby optimizing distinct mappings (Tan et al., 2022). For example, different hospitals may use the same radiography images to diagnose different diseases, resulting in heterogeneous local losses $\mathcal{L}_k$. Another case of task heterogenity emerges when models are shared among applications with compound ML models, as exemplified by certain Android implementations (Huang et al., 2021). Furtheremore, agents may employ different model architectures $\mathcal{F}$ due to limitations in memory or computational resources (Park & Joe-Wong, 2024).

The Collaborative Training Trilemma

Maintaining effective and efficient learning under the privacy requirements of collaborative learning (Section 2.1), while simultaneously accommodating for data and system heterogeneities (Sections 2.2 and 2.2.3), poses significant challenges. These challenges can be captured by the *trilemma* illustrated in Figure 7, which serves as the foundation for analyzing the design choices discussed in the following subsections.

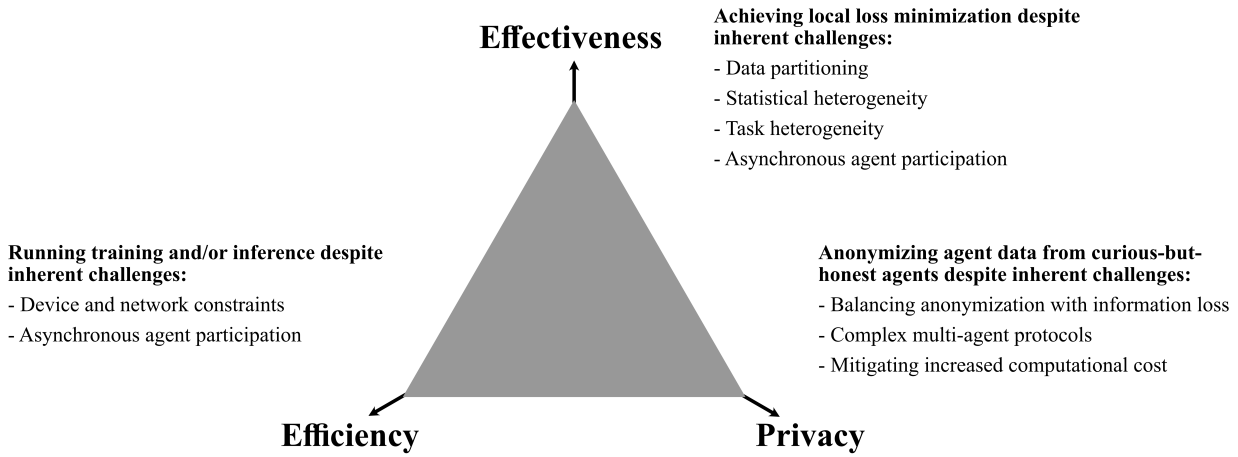

Figure 7: The trilemma of collaborative learning: balancing effectiveness (Section 2.3), efficiency (Section 2.3) and privacy (Section 2.5).

2.3 Improving Learning Effectiveness

Learning algorithms are effective if the generated models manage to minimize the learning loss of the agents, according to (1). The native proposition of collaborative learning (Algorithm 1) offers strong theoretical guarantees (Yu et al., 2019; Stich, 2019) to convergence to the solution of (5), particularly when the local datasets $\mathcal{D}_{k\in\mathcal{K}}$ are i.i.d. However, real-world deployments often diverge from this ideal case, as outlined in Section 2.1.

In such circumstances, the commonly used *FedAvg* algorithm becomes suboptimal to achieve adequate minimization of local objectives. This motivates the exploration of alternative strategies aimed at **improving the effectiveness of the models used by agents to minimize their respective local objectives**, as discussed in the following sections and outlined in Figure 8.

2.3.1 Data-based Techniques

Federated learning employs a **shared global model**, computed by the server as an average of the models over some participants in (see step (11)), to compute an optimal solution θ^* of the optimization problem (5). This approach performs effectively in its native application of horizontal cross-device scenarios with small, balanced, and similar datasets (Yang et al., 2018; Hard et al., 2018). However, in horizontal settings, under statistical data heterogeneity (seen in Section 2.2), or unbalanced agent participation (seen in Section 2.2.3), the returned local models $\theta_k(t)$ show high variance, which famously leads the average model $\bar{\theta}(t)$ to drift away from θ^*, the solution of (5), a phenomenon well-known in FL as *client drift* (Karimireddy et al., 2020; Li et al., 2022). To mitigate this client drift, a variety of methods have been developed that retain a shared global model paradigm but modify the *FedAvg* Algorithm 1 to achieve tighter convergence to an optimal solution θ^*. These methods, collectively known as *data-based solutions*, range from regularizing the local models to influencing the underlying data distribution through data augmentation and client sampling.

Local regularization. Regularization approaches aim to minimize the discrepancy between the local models $\{\theta_k(t)\}_{k\in\mathcal{K}_t}$ returned to the server by modifying the local updates (7). Bounding the variance of local gradient updates in (7) across agents has shown theoretical guarantees (Li et al., 2020) as well as empirical evidence (Li et al., 2020) for *FedAvg* to converge to the optimal solution θ^*, also in non-i.i.d. settings. Local regularization methods can be broadly categorized into three approaches (Kairouz et al., 2021):

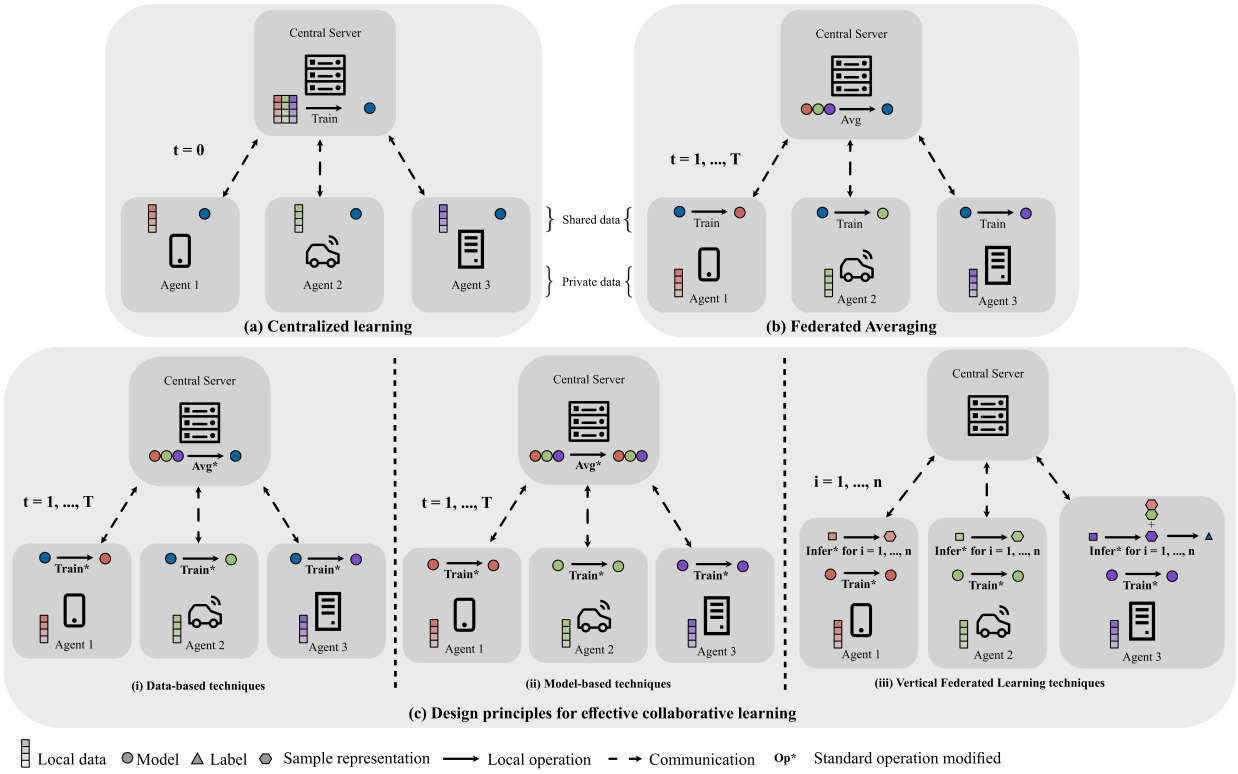

Figure 8: Variations of federated learning that aim at improving the learning effectiveness. In contrast to *centralized learning* (a), where data is collected from agents and training occurs on a central server, (b) *Federated Averaging* represents the baseline approach for collaborative learning systems. Methods to improve effectiveness (c) focus on *data-based* (i) and *model-based* (ii) techniques, typically applied to horizontal partitions. An alternative to the baseline is *Vertical Federated Learning (VFL) (iii)* where the inference becomes *collaborative*.

- *Gradient control methods* reduce the gradient variance across agents and within local iterations by incorporating variance correction terms into the local updates (7). These terms are usually computed from the previous iterations of the local $\theta_k(t-1)$ and global $\theta(t-1)$ models. The main approaches rely on adaptive algorithms (Karimireddy et al., 2020; Khanduri et al., 2021; Kim et al., 2024) or dual-based optimization (Fan et al., 2022; Gong et al., 2022; Zhou & Li, 2023).
- *Selective update strategies* reduce noise by refining which model components are updated in (7), either by restricting updates to the least noisy layers (Park & Joe-Wong, 2025) or by using local optimizers that ensure smoother gradients (Wang et al., 2020).
- *Objective modification approaches* modify the local empirical objective (2) by penalizing divergence from the global model $\theta(t-1)$ (Li et al., 2020; Yao & Sun, 2020; Li et al., 2021) or, conversely, encouraging deviation from the previous local model $\theta_k(t-1)$ (Li et al., 2021). Alternatively, second-order extensions to SGD provide smoother local updates (Fallah et al., 2020; T Dinh et al., 2020; Luo et al., 2025).

Data augmentation. Data augmentation addresses distributional shifts, including quantity, feature, and label shifts, as well as missing labels, by modifying the local datasets $\mathcal{S}_k$ used for local training (see (7) in Algorithm 1). Augmentation strategies in federated settings can be categorized as follows:

- *Server-instructed augmentation.* A central server may augment local datasets by redistributing samples between agents to balance distributional shifts (Zhao et al., 2018), by instructing selec-

tive down-sampling of overrepresented classes in agents local datasets (Duan et al., 2020), or by generating synthetic samples and communicating them to the agents (Jeong et al., 2018).

- *Independent local augmentation.* Agents may augment their local datasets independently. Techniques include traditional transformations (Zhang et al., 2018) or learned generative models (Wu et al., 2020). When labels are missing, contrastive losses can be used to train on the unlabeled local dataset, augmented with negative samples (Li et al., 2022; He et al., 2025).

- *Collaborative local data augmentation.* Agents may augment their local datasets by collaborating with other agents. This is particularly salient in vertical partitioning, where common samples have to be aligned between agents (Yang et al., 2022; Kang et al., 2022). When not all local samples of the local datasets can be aligned, the pseudo features or labels can be estimated with generative models (Kang et al., 2022), or by matching features across agents (Yang et al., 2022).

Client selection. Client selection techniques tackle the statistical imbalance seen in horizontal partition settings by modifying the server strategy for sampling agents in line 3 of Algorithm 1. There are two approaches:

- *Total availability.* Some works assume total availability of agents. Strategies are designed to select underperforming clients (Li et al., 2021; Cho et al., 2020), group clients with similar performance in the same round (Chai et al., 2020; Yang et al., 2021), or prioritize clients that contribute to improving the global accuracy (Wang et al., 2020). These strategies can be further optimized using reinforcement learning techniques on the server side, such as Q-learning (Wang et al., 2020) or multi-armed bandits (Yang et al., 2021).

- *Partial availability.* Other works assume that not all agents are available at the same time. Several participation patterns are studied, ranging from cyclical patterns (Cho et al., 2023) to temporal participation modeled by stationary Markov chains (Wang & Ji, 2022; Rodio et al., 2023), or a combination of both using R-separated Markov chains (Sun et al., 2025; Xiang et al., 2024). Some works look at spatio-temporal correlation using Markov chains (Rodio et al., 2023) or assuming independent agent unavailability (Wang & Ji, 2024) or handle agent turnover with departing and newly arriving agents (Ruan et al., 2021). The techniques involve maintaining estimate of the participation statistics to modify the sampling strategy (line 2 in Algorithm 1) (Cho et al., 2020; Wang & Ji, 2022; Ruan et al., 2021; Ribero et al., 2022), the aggregation (line 11 in Algorithm 1) (Ruan et al., 2021; Ribero et al., 2022; Wang & Ji, 2024), and the local update (line 19 in Algorithm 1) (Ruan et al., 2021) accordingly to guarantee a biased convergence to θ^*. Theoretical results on correlated participation (Rodio et al., 2023; Cho et al., 2023) establish the trade-off of fast and tight convergence to θ^*.

2.3.2 Model-based Techniques

In the first place, FL aims to minimize the local objectives of (1) among all agents. However, the optimal solutions F^* of the different local objectives in (1) might be different between agents due to factors like distributional shifts (Section 2.2) or task heterogeneity (Section 2.2.3). This implies that the convenient approach of looking for a model solution $F_\theta^* \in \mathcal{F}$, common to all agents, is suboptimal. Therefore, some efforts have been devoted to allow model solutions that are personalized to each agent, while retaining the main steps of the collaborative learning framework depicted in Section 2.1. With personalization, the optimization problem becomes

$$\min_{\theta_1 \in \Theta_1, \ldots, \theta_K \in \Theta_K} \frac{1}{K} \sum_{k=1}^{K} \mathcal{L}_k(F_{\theta_k}). \tag{6}$$

Personalization is a well-established direction in federated learning (Tan et al., 2022; Sabah et al., 2024). It can be achieved using two approaches: keeping a common architecture while personalizing the model parameters $\theta_k \in \Theta$, or personalizing the model architectures $\{\Theta_k\}_{k \in \mathcal{K}}$ themselves.

Parameter personalization. $\forall l \neq k, \theta_k \neq \theta_l; \Theta_k = \Theta_l$ Parameter personalization maintains a common domain Θ for the personalized models but allows models to converge to distinct parameters. These techniques apply exclusively to horizontal partitions, where all agents maintain similar model architectures. We distinguish:

- *Local adaptation.* Local adaptation extends Algorithm 1 by allowing each agent to modify the received global model parameters $\theta(t)$ into a personalized version $\theta_k(t+\frac{1}{2})$ through additional steps in the local training round. This adaptation can occur after the federated learning process, as in *meta-learning*, where $\theta(T)$ is rapidly fine-tuned on local data (Finn et al., 2017; Jiang et al., 2019). Alternatively, adaptation can be performed in every round t by modifying the local model initialization step (line 5 in Algorithm 1), e.g., by interpolating between $\theta(t)$ and $\theta_k(t-1)$ (Deng et al., 2020), using learnable local masks (Li et al., 2021; Dai et al., 2022; Huang et al., 2022), or applying a low-rank decomposition to $\theta(t)$ (Dadras et al., 2024).

- *Hypernetwork-based.* A hypernetwork (Klein et al., 2015), that is, a network that generates the weights of a target model, can be used to produce personalized parameters for each agent. Given an agent representation vector c_k, the server-side hypernetwork $\mathcal{H}_\alpha$ with trainable parameters α outputs the personalized model $\theta_k(t) = \mathcal{H}_\alpha(c_k)$. Instead of maintaining a shared model $F_{\theta(t)}$, all agents share a *common* hypernetwork with learnable weights α jointly trained across agents by minimizing (Shamsian et al., 2021)

$$\frac{1}{K}\sum_{k=1}^{K}\hat{\mathcal{L}}_k\big(F_{\mathcal{H}_\alpha(c_k)}\big). \tag{7}$$

  Training is performed on the server, where agents transmit $\nabla_\alpha \hat{\mathcal{L}}_k$ and the server updates α via gradient steps. The representation c_k can be fixed (e.g., encoding an agent type) or updated jointly with α.

- *Cluster-based.* Assignments can be static (Duan et al., 2021), where clusters operate independently, or dynamic, allowing clustering and learning to occur simultaneously. A straightforward approach leverages similarities for clustering. These methods typically alternate between clustering and local updates, framing the problem as an optimization task (Long et al., 2023). Clustering optimization aims to minimize inter-cluster similarities (Briggs et al., 2020; Werner et al., 2024) or maximize intra-cluster similarities (Sattler et al., 2020; Briggs et al., 2020; Long et al., 2023; Werner et al., 2024). Similarities are computed between clients (Briggs et al., 2020), clients and clusters (Ghosh et al., 2020; Mansour et al., 2020; Werner et al., 2024), or clients and the global model (Sattler et al., 2020), using losses (Ghosh et al., 2020; Sattler et al., 2020) or gradients (Werner et al., 2024). Techniques include greedy assignment (Ghosh et al., 2020; Mansour et al., 2020), recursive bi-partitioning (Sattler et al., 2020), recursive fusion (Briggs et al., 2020), and threshold-based clustering robust to Byzantine settings (Werner et al., 2024). Assignments can be performed server-side (Sattler et al., 2020; Long et al., 2023) or client-side, although this requires sending all cluster models (Ghosh et al., 2020). Additionally, some studies also address the assignment of new entrants to clusters (Duan et al., 2021).

- *Similarities between agents* can be used to favor aggregation between the most similar agents. A similarity metric $\Omega = [\omega_{k,l}] \in \mathbb{R}^{K^2}$, measuring the proximity between local objectives, is maintained by the server to build personalized models. These similarities can be computed explicitly using metrics on local models $\theta_k(t)_{k\in\mathcal{K}}$ (Huang et al., 2021; Chen & Zhang, 2022) or by evaluating local empirical objectives on models of other agents $\hat{f}_k(\theta_l)_{(k,l)\in\mathcal{K}^2}$ (Onoszko et al., 2021; Zec et al., 2022). Given Ω, models can be personalized by averaging the most similar ones (Onoszko et al., 2021) or using weighted averages (Zec et al., 2022; Huang et al., 2021). Alternatively, Ω can be learned as a problem parameter by adding a regularization term $\mathcal{R}(\Omega, [\theta_{k,l}]_{(k,l)\in\mathcal{K}^2})$ to (6): (Smith et al., 2017; Huang et al., 2021; Shoham et al., 2019; Chen & Zhang, 2022; Chen et al., 2022). This requires alternating optimization, updating Ω and model parameters in separate steps (Smith et al., 2017; Huang et al., 2021; Chen et al., 2022). While learning Ω introduces significant communication

overhead, decentralized protocols (Zec et al., 2022) and smart sampling (Zec et al., 2024) help mitigate scaling issues.

Architecture personalization. $\forall l \neq k, \theta_k \neq \theta_l; \Theta_k \neq \Theta_l$ Agents can utilize different model architectures, for example, neural networks of different sizes. The need for architecture personalization is driven both by concept statistical imbalance (seen in Section 2.2.2) but also by system heterogeneity constraints like device heterogeneity and application heterogeneity (seen in Section 2.2.3). We distinguish:

- *Model decoupling.* Local models, θ_k, are divided into a shared component, $\theta_k^c = \bar{\theta}$, trained using the *FedAvg* algorithm, and a private component, θ_k^p. In neural networks, this split is typically *layer-wise* (depth) (Wang et al., 2016; 2023; Liang et al., 2020; Arivazhagan et al., 2019; Bui et al., 2019), but may also occur *width-wise* (Diao et al., 2021) or in a hybrid manner (Kang et al., 2025). The structure of the split directly influences communication costs. Some strategies reduce these costs by replicating shared layers to private components, especially in CNNs (Park & Joe-Wong, 2024), or by focusing on the most impactful layers (Park & Joe-Wong, 2025). The level of decoupling impacts model performance – greater device heterogeneity benefits from more private neurons, while noisy local updates favor a larger shared component (Liang et al., 2020; Collins et al., 2021). The optimal split can be determined through *adaptive* (Vahidian et al., 2021) or even *learnable* (Deng et al., 2022) strategies. However, incomplete aggregation schemes introduce security risks (Park & Joe-Wong, 2024).
- *Multi-modal personalization.* The server maintains M common models, $\theta = \cup_{m=1,\ldots,M} \theta_m^c$, where model θ_m^c corresponds to a data modality m. Agents selectively combine these models based on their local data, addressing feature shifts. Identifying the modality can rely on prior knowledge (Chen & Zhang, 2022) or statistical estimation (Huang et al., 2019). Model combinations vary: agents may assign distinct models per sample (Huang et al., 2019) or aggregate outputs from multiple models when tasks require multi-modal fusion (Chen & Zhang, 2022). Training often leverages *parameter personalization*, using clustering (Huang et al., 2019) or agent similarity (Smith et al., 2017; Chen & Zhang, 2022) to train these models jointly.
- *Knowledge distillation.* (Hinton et al., 2014) is a technique used for transferring the knowledge of a model to another model of a different architecture. In practice, it enriches the dataset of agent k (the student), $\mathcal{S}_k$ (student), with a dataset $\mathcal{S}_c$ generated by another participant (the teacher). The teacher model is trained to align its predictions on $\mathcal{S}_c$. This is often used to provide rapid training of a model with different architectures (for example, for model compression) on a similar task by leveraging the knowledge of a larger pre-trained model. In client-server architectures, knowledge transfer can occur server-to-client (Li & Wang, 2019), peer-to-peer (Zhu et al., 2021; Bistritz et al., 2020), or bidirectionally (He et al., 2020; Lin et al., 2020). This is particularly useful in settings with limited local labeled data (Bistritz et al., 2020).

2.3.3 Training under Vertical Partitioning

The techniques presented in Sections 2.3.1 and 2.3.2 consider *horizontal partitioning* of the data (Section 2.2.1) However, in practice, agents might obtain a more effective model by collaborating under the vertical partitioning described in Section 2.2.1. Training under vertical partitioning and its variants can be achieved under *collaborative* inference or *isolated* inference.

Training under collaborative inference. Models can be trained under vertical partitions by collaboratively optimizing a collaborative inference loss. In previous Sections 2.3.2 and 2.3.1, each agent trains a local model θ_k used to perform inference *in isolation* on its local data. Collaboration occurs solely during training when agents periodically share their local model iterate $\theta_k(t)$. However, when agents know about the same sample, as seen in the example of *vertical data partition* (Section 2.2.1), agents might benefit from collaborating during inference as well, for building a common prediction. This collaborative inference imposes a reevaluation of the problem definition of standard FL in (5), and a reevaluation of the training protocol. Under vertical partition, input features of a sample (x^i, y^i) are split across the agents $k \in \mathcal{K}$. Each agent k holds

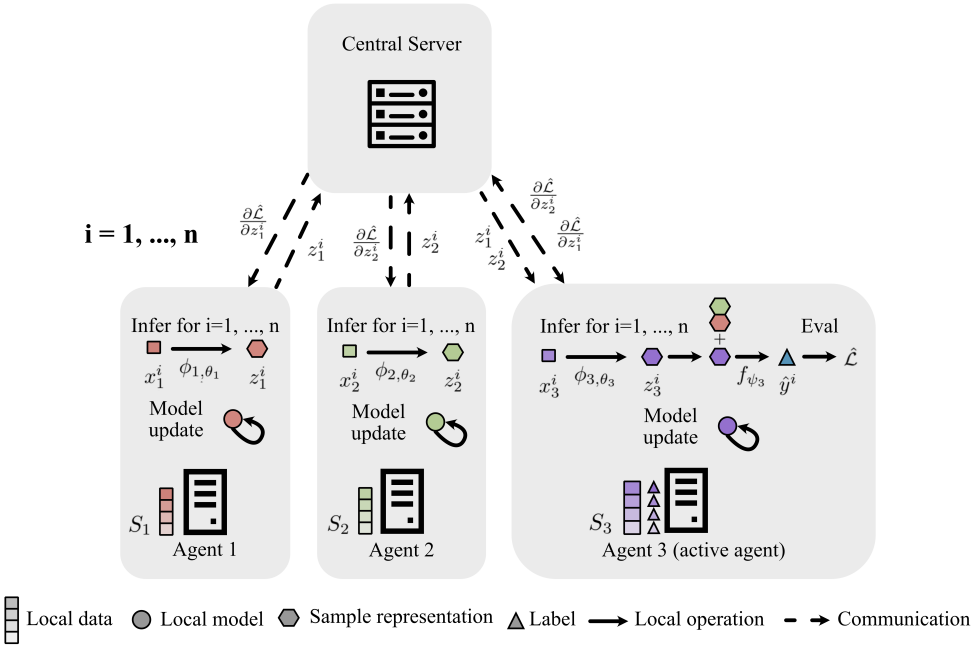

Figure 9: Training under vertical partitioning with three agents. During inference, the computations of the global model $F_{\theta,\psi}$ are distributed across agents: to produce a prediction $\hat{y}^i$ for input x^i, each agent k communicates its partial representation $z_i^k = \phi_{k,\theta_k}(x_k^i)$ of the sample. Only one or a subset of the agents (here, agent 3) possesses the true label, which can evaluate the training loss $\hat{\mathcal{L}}$. During training, the gradients with respect to the learnable parameters $\theta_1, \theta_2, \theta_3, \psi_3$ must be communicated across agents in the form of $\frac{\partial \hat{\mathcal{L}}}{\partial z_k^i}$.

x_k^i, a subset of the features x^i such as $x^i = \cup_{k\in\mathcal{K}} x_k^i$, forming the local dataset $\mathcal{S}_k = \{x_k^i \in \mathcal{X}\}_{i\in\mathcal{I}}$. For simplicity, we assume that a single agent, agent k^* has access to the labels, hence $\mathcal{S}_{k^*} = \{(x_{k^*}^i, y^i) \in \mathcal{X} \times \mathcal{Y}\}_{i\in\mathcal{I}}$. Since the features of each sample are split across agents, the global model $F \in \mathcal{F}$, used to infer on the centrally available sample (x^i, y^i), can be naturally decomposed to process the features locally [3]. Consider that each agent k holds a local model component ϕ_{k,θ_k} mapping from $\mathcal{X}_k$ to latent space $\mathcal{Z}_k$, parameterized by $\theta_k \in \Theta_k$ while a final predictor $f_{\psi_{k^*}}$, parameterized by $\psi_{k^*} \in \Psi_{k^*}$, is held by agent k^*. Letting F's model parameters be $\theta = (\theta_1, \ldots, \theta_K, \psi_{k^*})$ with parameter space $\Theta = \Theta_1 \times \cdots \times \Theta_K \times \Psi_{k^*}$, the inference of F_θ on a sample indexed by i can be expressed as

$$\hat{y}^i = F_\theta(x^i) = F_\theta((x_1^i, \ldots, x_K^i)) = f_{\psi_{k^*}}\left(\phi_{1,\theta_1}(x_1^i), \ldots, \phi_{K,\theta_K}(x_K^i))\right). \tag{8}$$

Therefore, the partition of the global model F results in a *collaborative inference* described in Figure 9: for performing a prediction on x^i, each agent k communicates to contribute its partial representation $\phi_{k,\theta_k}(x_k^i)$ of the sample.

The population loss can be defined for all agents as

$$\mathcal{L}_{VFL}(F_{\theta_1,\ldots,\theta_K,\psi_{k^*}}) := \mathbb{E}_{((x_1,\ldots,x_K),y)\sim\mathcal{D}}\ell\left(F_{\theta_1,\ldots,\theta_K,\psi_{k^*}}\left((x_1,\ldots,x_K)\right), y\right). \tag{9}$$

Since the data distribution $\mathcal{D}$ is unknown, the loss (9) is approximated by an empirical counterpart defined over S_k

$$\hat{\mathcal{L}}_{VFL}(F_{\theta_1,\ldots,\theta_K,\psi_{k^*}}) := \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}} \ell\left(F_{\theta_1,\ldots,\theta_K,\psi_{k^*}}\left(x_1^i,\ldots,x_K^i\right), y^i\right). \tag{10}$$

Notably, both (9) and (10) are common to all agents, allowing the global loss (3) to be equivalently expressed in this collaborative setting as

$$\frac{1}{K}\sum_{k=1}^{K}\mathcal{L}_{VFL}(F) = \mathcal{L}_{VFL}(F). \tag{11}$$

However, unlike in (3), where local losses can be computed independently by each agent, the structure of (9) requires collaborative inference to evaluate the loss. As a result, computing the objective requires inter-agent communication, and we refer (9) and (10) to as *collaborative inference loss*. Despite this difference, the learning goal remains the same: to approximate a minimizer of the population loss (9) by minimizing its empirical surrogate (10) over a parametrized hypothesis class. This leads to the following optimization problem of *Vertical Federated Learning*, in direct analogy with (5):

$$\min_{\theta_1\in\Theta_1,\ldots,\theta_K\in\Theta_K,\psi_{k^*}\in\Psi_{k^*}} \hat{\mathcal{L}}_{VFL}(F_{\theta_1,\ldots,\theta_K,\psi_{k^*}}). \tag{12}$$

The typical training pipeline follows two steps (Liu et al., 2024):

- *Privacy-preserving sample alignment.* The formulation in (10) assumes a shared understanding of sample indices $\mathcal{I}$, their features (x^i), and labels (y^i) across agents. However, in practice, each agent only sees its local view, making the computation of (10) infeasible without preliminary coordination. To address this problem, the first step is to identify the common index set $\mathcal{I}$ across the agents. It is usually achieved through the *Private Set Intersection (PSI)* protocols (Liang & Chawathe, 2004; Pinkas et al., 2018). PSI is extended to multiparty setups in Zhou et al. (2021); Lu & Ding (2020),
- *Collaborative training.* The optimization problem in (12) is solved collaboratively between agents $\mathcal{K} = \{1, \ldots, K\}$, as summarized in Algorithm 2. It proceeds in T rounds, where each round consists of two phases:
  - *Collaborative Inference (lines 3–9):* Each agent k computes intermediate representations $z_k^i = \phi_{k,\theta_k}(x_k^i)$ (line 3) for each sample and sends them to the active agent k^* (line 4), which aggregates them and computes the empirical loss $\hat{\mathcal{L}}(F_\theta)$ on the mini-batch (line 7), updates ψ_{k^*} (line 8), and sends partial gradients $\frac{\partial\hat{\mathcal{L}}}{\partial z_k^i}$ to each agent (line 9).

[3] A distinct case, known as *split learning* (Gupta & Raskar, 2018; Vepakomma et al., 2019), arises when the inference model is partitioned across multiple agents, with only one holding the actual data. This setup falls outside the scope of our study, where our fundamental assumption is that data is distributed and private to its owner.

– *Collaborative Model Update (lines 11–12):* Each agent applies the chain rule to compute $\nabla_{\theta_k}\hat{\mathcal{L}}$ (line 11) and updates its local parameters (line 12).

Algorithm 2 Vertical Federated Learning algorithm (VFL) adapted from Liu et al. (2024)

Require: Set of agents $\mathcal{K} = \{1, \ldots, K\}$, shared sample index set $\mathcal{I}$, learning rates η_1, η_2, local parameters $\theta(0) = \{\theta_k(0)\}_{k=1}^{K} \cup \{\psi_{k^*}(0)\}$, number of communication rounds T, mini-batch size B

1: **for** each round $t = 1, \ldots, T$ **do**

Collaborative Inference:

2: **for** each agent $k \in \mathcal{K}$ **in parallel do**

3: Compute local representations: $z_k^i = \phi_{k,\theta_k(t-1)}(x_k^i)$ for all [4] $i \in \mathcal{I}$

4: Send $\{z_k^i\}_{i\in\mathcal{I}}$ to the active agent k^*

5: **end for**

Active agent k^*:

6: Aggregate $\{z_k^i\}$ and compute empirical predictions for each $i \in \mathcal{I}$: $\{\hat{y}^i\}_{i\in\mathcal{I}} = \{f_{\psi_{k^*}}(z_1^i, \ldots, z_K^i\}_{k\in\mathcal{K}})\}_{i\in\mathcal{I}}$:

7: Compute partial empirical loss on samples of index $\mathcal{I}$, $\hat{\mathcal{L}}_{VFL}\left(F_{\psi_{k^*},\theta_1,\ldots,\theta_K}\right) = \frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}} \ell\left(\hat{y}^i, y^i\right)$

8: Update predictor model: $\psi_{k^*}(t) \leftarrow \psi_{k^*}(t-1) - \eta_1 \nabla_{\psi_{k^*}} \hat{\mathcal{L}}_{VFL}$

9: Compute and send partial derivatives to each agent k: $\left\{\frac{\partial \hat{\mathcal{L}}}{\partial z_k^i}\right\}_{i\in\mathcal{B}_t}$

Collaborative Model Update:

10: **for** each agent $k \in \mathcal{K}$ **in parallel do**

11: Compute gradients via chain rule: $\nabla_{\theta_k}\hat{\mathcal{L}} = \sum_{i\in\mathcal{I}} \frac{\partial \hat{\mathcal{L}}_{VFL}}{\partial z_k^i} \cdot \frac{\partial z_k^i}{\partial \theta_k}$

12: Update local model: $\theta_k(t) \leftarrow \theta_k(t-1) - \eta_2 \nabla_{\theta_k} \hat{\mathcal{L}}_{VFL}$

13: **end for**

14: **end for**

15: **return** Final parameters $\theta = (\theta_1(T), \ldots, \theta_K(T), \psi_{k^*}(T))$ and sample predictions $\{\hat{y}^i\}_{i\in\mathcal{I}}$

Algorithm 2 can be extended in several ways to adapt to more realistic settings.

- *Pre-training.* To decrease the level of communication required by the collaboration, agents may pre-train their local models on private data prior to collaborative training (Feng et al., 2022; Shen et al., 2025). This approach leads to scalable VFL, but may introduce bias from local data distributions (Shen et al., 2025).
- *Model averaging.* The VFL protocol in Algorithm 2 assumes that each agent $k \in \mathcal{K}$ maintains a private local model ϕ_{k,θ_k} with domain parameters Θ_k which are individually updated based on the collaborative model update phase (lines 9–12). However, in practice, certain agents might use the same model architecture, leading to the same domain space $\Theta_k = \Theta_l$. This can occur, for example, when agents (e.g., hospitals) collect the same features related to the same samples (e.g., patients). In such cases, it can be beneficial to introduce a periodic averaging of the models, $\theta_k(t) = \theta_l(t) = \frac{\theta_k(t)+\theta_l(t)}{2}$, in the collaborative model update steps (Fu et al., 2024).
- *Unlabeled data.* The formulation in (12) assumes full access to supervision through the labels y^i. However, in practical scenarios, a significant portion of the data may be unlabeled. To handle such cases, the empirical loss in (10) can be extended to include unsupervised objectives, such as contrastive losses over unlabeled samples (Li et al., 2022; He et al., 2025). Additionally, the local inference modules ϕ_{k,θ_k} can be instantiated as autoencoders, enabling self-supervised training via reconstruction losses (Feng, 2022). These extensions apply primarily to the collaborative inference phase (lines 3–4) of Algorithm 2, and they allow agents to learn useful representations even in the absence of labels.

Training for isolated inference. Transfer techniques aim to train models capable of performing isolated inference on vertically partitioned data while achieving effectiveness comparable to collaborative inference.

This addresses the primary limitation of collaborative inference, which is the communication overhead at inference time, while still exploiting the knowledge contained in vertically partitioned datasets. Consider a transfer partition as defined in Section 2.2.1, involving a *source* agent s and a *target* agent t, each aware of their common sample set $\mathcal{I}_s \cap \mathcal{I}_t$, obtained, for example, via *PSI* protocols. Two main approaches are commonly used depending on the structure of the partition (Liu et al., 2024):

- *Knowledge distillation.* When the feature spaces of the agents exhibit significant overlap, knowledge distillation (KD) techniques, also discussed for horizontal partitioning in Section 2.3.2, can be adapted to vertical settings. The objective is to transfer the knowledge of the collaborative inference model $(\phi_{s,\theta_s}, \phi_{t,\theta_t}, f_{\psi_s})$, pre-trained on domains $\mathcal{X}_s$ and $\mathcal{X}_t$ by agents s and t according to (12), to a student model $(\phi_{\theta_t^t}, f_{\psi_t^t})$ for agent t, that performs *isolated inference* on domain $\mathcal{X}_t$. To this end, agent t trains its model on a distillation loss that aligns the representations of $\phi_{\theta_t^t}$, ϕ_{s,θ_s}, and ϕ_{t,θ_t} in the shared latent space $\mathcal{Z}_s \times \mathcal{Z}_t$, using aligned unlabeled samples (Huang et al., 2023; Ren et al., 2022).

- *Transfer learning.* When the number of aligned samples is limited and only the source agent s has sufficient data to train an independent model, transfer learning techniques (Liu et al., 2020) can be employed. The goal is to adapt a model pre-trained by agents s on its dataset S_s and domain $\mathcal{X}_s$ to the target domain $\mathcal{X}_t$, using a shared embedding space $\mathcal{Z}$. The core assumption in this setting is that unaligned samples from s and t still yield semantically similar representations in $\mathcal{Z}$. The process involves two steps. First, agent s trains a local encoder $\phi_{s,\theta_s} : \mathcal{X}_s \to \mathcal{Z}$ and a predictor f_{ψ_s} on S_s, yielding representations $Z_s = \{z_s^i, i \in \mathcal{I}_s\} = \{\phi_{s,\theta_s}(x_s^i), i \in \mathcal{I}_s\}$. Then, agent t initializes its own encoder $\phi_{t,\theta_t} : \mathcal{X}_t \to \mathcal{Z}$ and predictor f_{ψ_t}. The training of parameters ψ_t and θ_t proceeds by aligning $Z_t = \{z_t^i, i \in \mathcal{I}_t\} = \{\phi_{t,\theta_t}(x_t^i), i \in \mathcal{I}_t\}$ with Z_s in the latent space $\mathcal{Z}$ using a transfer loss, while keeping ψ_s, θ_s frozen. Since Z_s and Z_t may differ in size, advanced transfer losses are typically required, such as feature matching (Yang et al., 2022), pseudo-label matching (Feng et al., 2022), or adversarial methods (Kang et al., 2022).

2.4 Improving Learning Efficiency

While Section 2.3 explored techniques to enhance the effectiveness of the learning process, these approaches often overlook the constraints of the communication and computing resources introduced in Section 2.2.3. This section examines strategies for *efficient* learning, that is, ways to achieve the fast convergence of the learning models despite the constraints given by the network topology, the heterogeneity of the devices and network connections, and the often costly and unreliable communication links.

2.4.1 Model Aggregation

The efficiency of collaborative learning depends primarily on the model aggregation pattern, defined by the topology and the timing of the model parameter exchange and aggregation. This section discusses aggregation schemes in three aggregation topologies: centralized, decentralized, and hierarchical, as shown in Figure 10.

Centralized aggregation. Centralized aggregation is the standard strategy for combining local models, as employed in Algorithm 1. In this setting, a central server coordinates training: at each round $t = 1, \ldots, T$, it collects local model parameters $\theta_k(t)$, each obtained after training for E local epochs, aggregates them, and redistributes the result to all agents.

The motivation for centralized aggregation originates in distributed computing, where distributing stochastic gradient computations can theoretically reduce training time by a factor proportional to the number of agents K (Zinkevich et al., 2009; McDonald et al., 2010). However, naive parallelization with aggregation at every epoch ($E = 1$) incurs significant communication overhead, which can offset the computational benefits depending on system cost constraints described in Section 2.2.3.

When communication overhead outweighs the gains from faster convergence, a simple workaround is *one-shot aggregation*, where models are aggregated once after local training completes ($T = 1$)(Mcdonald et al., 2009;

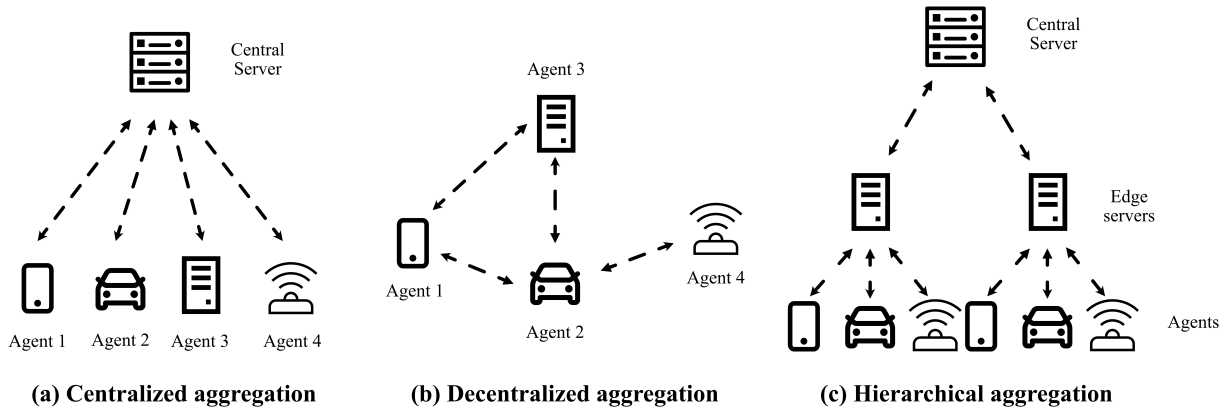

Figure 10: Different aggregation patterns for collaborative learning. *Centralized aggregation* (a) uses a central server that collects and combines model updates from all agents. *Decentralized aggregation* (b) allows agents to exchange and aggregate updates over a network with restricted communication links. *Hierarchical aggregation* (c) relies on a multi-level structure where a central server coordinates with edge servers, which in turn oversee the aggregation of updates of their associated agents.

Zinkevich et al., 2010). While highly communication-efficient (Zhang et al., 2012), this strategy relies on strong assumptions, such as statistical similarity across agents datasets, that become increasingly difficult to satisfy as K increases (Li et al., 2020; Khaled et al., 2020).

A more balanced alternative is to let agents perform multiple local updates ($E > 1$) before periodically aggregating their models. This approach, known as *local SGD*, leverages the robustness of SGD to noise and relaxed synchronization (Dean et al., 2012; Bottou et al., 2018). It has proven successful both in practice and theory, achieving near-linear speedup under mild i.i.d. assumptions (Stich, 2019; Yu et al., 2019; Li et al., 2020).

Further refinements include hybrid schemes combining one-shot and local SGD (Hou et al., 2022), the use of server-side optimizers (Reddi et al., 2021), and dual-based methods such as ADMM (Boyd et al., 2011; Zhou & Li, 2023).

Decentralized aggregation. Decentralized aggregation is the opposite of centralized aggregation: agents aggregate models directly without a central coordinator. It is often proposed as a more robust alternative, mitigating the single point of failure inherent to centralized schemes. In some cases, it is also necessitated by system constraints. Unlike centralized approaches, both training and aggregation occur over a mesh network, introducing additional challenges related to network topology.

Distributed aggregation builds on the early results of *distributed averaging* (DeGroot, 1974; Tsitsiklis et al., 1986; Bertsekas & Tsitsiklis, 2015). For simplicity, consider agents connected according to a mesh topology modeled as a graph, and holding initial value vectors $\pi_i(0)$, with an average of $\bar{\pi}(0)$. At each step t, the updates follow a linear rule: $\forall k \in \mathcal{K}, \pi_k(t+1) = w_{kk}\pi_k(t) + \sum_{l \in \mathcal{N}(k)} w_{kl}(t)\pi_l(t)$, where $\mathcal{N}(k)$ denotes the neighbors of agent k. These updates can be expressed with a Markov chain as $\Pi(t+1) = W(t)\Pi(t)$ where $W(t) = [w_{kl}(t)]_{(k,l)\in\mathcal{K}^2}$ and $\Pi(t) = [\pi_{kl}(t)]$. Convergence to a consensus state occurs when all nodes reach the same value $\forall k \in \mathcal{K}, \lim_{t\to\infty} \pi_k(t) = \pi^*$. Average consensus is achieved if, in addition, $\pi^* = \bar{\pi}(0)$. In the literature, we find two algorithmic approaches for distributed averaging (Denantes et al., 2008):

- *Average consensus* algorithms assume that all nodes communicate and update their values simultaneously. The problem is formulated by minimizing the *Dirichlet energy*, expressed as

$$\frac{1}{2}\,\Pi(t)^\top L \Pi(t) \;=\; \sum_{i\in\mathcal{K}} \sum_{j\in\mathcal{N}(i)} \big(\pi_i(t) - \pi_j(t)\big)^2, \tag{13}$$

which quantifies the quadratic disagreement between neighboring nodes, where L is the graph Laplacian (Mohar et al., 1991) and $\mathcal{N}(i)$ denotes the neighbors of node i in the mesh network. In *continuous time*, updates towards a minimum of (13) can be expressed as a diffusion equation as $\dot{\pi}(t) = -L\pi(t)$ in (Cortés, 2006; Jiang & Wang, 2009; Olfati-Saber et al., 2007). In *discrete time*, the update takes a Markov chain form $\Pi(t+1) = W(t)\Pi(t)$ in (Blondel et al., 2005; DeGroot, 1974; Nedic et al., 2009; Olshevsky & Tsitsiklis, 2009). For example $W(t)$ is derived directly from (13) with the update rule $W(t) = I - \epsilon L$, with step size $\epsilon > 0$ and I the identity matrix in (Saber & Murray, 2003; Olfati-Saber & Murray, 2004; Olfati-Saber et al., 2007). For static graphs, convergence is guaranteed when W is doubly stochastic and connected. The convergence rate is exponential, determined by the topology of the graph (Fiedler, 1973), through the second-largest eigenvalue of the symmetric Laplacian $\lambda_2(L_s)$ (Fax & Murray, 2004; Olfati-Saber & Murray, 2004). Average consensus algorithms have been widely applied in multi-agent systems (Vicsek et al., 1995; Jadbabaie et al., 2003; Blondel et al., 2005; Fax & Murray, 2004) and adapted to address issues like dynamic topologies (Jadbabaie et al., 2003; Olshevsky & Tsitsiklis, 2009; Moreau, 2005) and delays (Blondel et al., 2005; Tsitsiklis et al., 1986).

- *Gossip algorithms* perform pairwise communication (Karp et al., 2000; Bawa et al., 2003). Upon a tick of its internal clock, node k communicates with node l with probability p_{kl}, and they update their values as $\pi_k(t+1) = \pi_l(t+1) = \frac{\pi_k(t)+\pi_l(t)}{2}$ (Boyd et al., 2005; 2006). The probability matrix P mirrors the role of adjacency matrices in synchronous algorithms. Gossip protocols are communication-efficient, and due to their inherent randomness, they are robust to asynchronous execution and partial observability (Jelasity et al., 2005; Denantes et al., 2008). Variations include geographic routing (Dimakis et al., 2008; Bénézit et al., 2010), weighted updates (Bénézit et al., 2010), and memory-augmented protocols (Cao et al., 2006; Liu et al., 2013).

Distributed averaging combined with local learning methods leads to decentralized aggregation. Compared to centralized aggregation, the challenge is that both the local learning and the model aggregation processes require multiple steps (as represented by respective indíces e and t). Early methods combine local gradient updates with a single averaging step (Tsitsiklis et al., 1986; Nedic & Ozdaglar, 2009). More advanced techniques incorporate multiple propagation steps (Johansson et al., 2008), handle link failures (Lobel & Ozdaglar, 2008), enforce constraints via projections (Nedic et al., 2010), or consider non-i.i.d. data (Wang et al., 2022). Gossip-based methods (Ram et al., 2009; Ormándi et al., 2013; Blot et al., 2016) enhance robustness to asynchronous participation and dynamically changing agent population(Hegedűs et al., 2019). Dual-based techniques, such as ADMM (Boyd et al., 2011), also support decentralized learning, where dual variable averaging is performed synchronously (Nesterov, 2009; Wei & Ozdaglar, 2012) or via gossip (Vanhaesebrouck et al., 2017).

The methods above rely on *distributed averaging*, where each agent's value $\pi_k(t)$ converges to a global consensus $\pi^* = \bar{\pi}(0)$. However, when the network topology reflects data heterogeneity, aggregation can instead balance global and local consensus. To preserve local information, the objective (13) can be modified to combine global agreement with local smoothness, as in $\frac{1}{2}\Pi(t)^T L\Pi(t) + \mu\|\Pi(t) - \Pi(0)\|_2^2$, as proposed in (Zhou et al., 2003). If this formulation admits a closed-form solution, it can be computed in a decentralized way using a broad class of *diffusion* (also known as *propagation* or *smoothing*) algorithms (Zhu & Ghahramani, 2002; Zhou et al., 2003). These mechanisms integrate naturally with local learning by initializing $\pi_k(t) = \theta_k(t)$, as in (Vanhaesebrouck et al., 2017).

Hierarchical aggregation. Both centralized and distributed aggregation schemes have limitations in scalability in terms of the number of agents and in terms of the size of the geographic area the agents are spread over. Hierarchical architectures are proposed to overcome these limitations by introducing multiple layers of aggregation. The lower layers typically follow a star topology with local aggregators. In wireless networks, these are often implemented at the base stations or edge servers. Aggregation at the higher layers may use a central server (Luo et al., 2020; Gupta et al., 2016; Wu et al., 2023; Azimi-Abarghouyi & Fischione, 2025), or be performed over a mesh structure (Castiglia et al., 2021; Zhang et al., 2022).

Hierarchical architectures are often motivated by retransmission constraints, where the wide geographic distribution of agents requires the model parameters to be transmitted over long transmission paths (Gupta et al., 2016). In wireless networks, the hierarchical structure improves the transmission quality over wireless links (Wen et al., 2022) and increases the efficiency of the learning process by localizing certain operations, which conserves communication resources and reduces time (Gupta et al., 2016). Clustering of agents helps to manage the heterogeneity of the system, supports personalization (Castiglia et al., 2021; Wu et al., 2023), and increases privacy by limiting data traffic to specific administrative regions or social groups (Zhou et al., 2023).

2.4.2 Asynchronous Aggregation

The model aggregation approaches discussed above generally assume that all clients are fully synchronized in the learning process, that is, they are all aware of and, if selected, participate in a given global iteration t. This is often an unrealistic approach. First, due to the high number of agents which makes it is costly to manage a shared view in the entire system. Second, due to the presence of stragglers, that is, agents with limited computational capacity or slow network connectivity, which can result in the server receiving the updated model $\theta_k(t)$ from the agents (line 11) at different times. The common concept to describe stragglers is *staleness*, defined for the server as the difference between the index of the current round global model update (line 12 of Algorithm 1), and the index $t_k \leq t$ of the local model $\theta_k(t_k)$ received by a straggler k. The issue of staleness is considered both for centralized aggregation and for decentralized learning, with the usual assumption of bounded staleness.

Asynchronous centralized aggregation. In the centralized aggregation scheme of (Xie et al., 2020), the global model is updated (line 11) at each asynchronously arriving client update, with weights that consider the staleness of the client. Under a non-i.i.d. data distribution, the models of straggling agents may still hold important updates for the global learning process. Therefore, (Li et al., 2022) suggests that late model updates from stragglers should also be weighted according to the importance of their update for the learning process. However, updating the global model after individual client updates has two disadvantages. It increases the communication requirements and contradicts the efforts for privacy-preserving FL (Nguyen et al., 2022). Therefore, to allow efficient model aggregation despite stragglers, (Chai et al., 2021) introduces a hierarchical scheme, where clients are dynamically clustered according to their update speed, synchronous learning is performed within the clusters, and asynchronous updates across the clusters. The same challenge is addressed in (Wei et al., 2015; Nguyen et al., 2022; Huba et al., 2022), suggesting delaying the global model update until a given number of asynchronous client updates arrive at the central parameter server.

Asynchronous aggregation for decentralized learning. Asynchronous updates with stragglers are also addressed for decentralized learning over graph topologies. The convergence of the learning process is proved in (Lian et al., 2018), assuming bounded staleness, and an averaging process that does not take this staleness into account. Instead, the weight matrix for the averaging is optimized according to the expected staleness of the neighboring nodes in (Bornstein et al., 2023). The comprehensive work (Even et al., 2024) suggests update strategies and proves convergence, considering both staleness due to computation delays and the effect of unreliable communication.

The convergence results of the asynchronous distributed learning schemes show that under bounded staleness, the convergence rate and the accuracy of synchronous systems are preserved, though the level of staleness and delayed aggregated updates may decrease the actual convergence speed.

2.4.3 Communication Efficiency

For collaborative learning, model parameters or gradients must be communicated to and from the parameter server (line 3 and 10 of Algorithm 1) or among the agents in a peer-to-peer or mesh architecture. Since the communication links between the agents and the central server or among the agents may be heterogeneous, unreliable, and costly (as described in Section 2.2.3), the transmission of the information needs to be organized by taking the state of the communication links and the cost of communication into account. The communication requirements of collaborative learning differ from traditional data transmission, as exact,

error-free information is usually not required for learning, and the model aggregation, as (11), needs the function of the inputs, rather than the individual values. These motivate the design of novel, FL-specific communication schemes, including the efficient transmission of information, the organization of medium access in wireless networks, and client scheduling (Hellström et al., 2022).

Transmission efficiency. Transmission efficiency measures the importance of the transmitted information versus the invested resources, such as bandwidth and transmission power. Bandwidth usage and learning efficiency in distributed learning can be balanced through quantization, that is, the number of bits representing a model parameter, and model sparsification, that is, the number of model parameters transmitted (Oh et al., 2024; Shlezinger et al., 2021; Xie et al., 2015). Similarly, in wireless networks, power control can achieve a tradeoff between energy efficiency and learning rate (Wang et al., 2020). As randomization can help the learning process, noise and wireless interference are proposed to be tuned dynamically to accelerate learning in (Zhang et al., 2022). To ensure that important model updates are transmitted successfully, (Liu et al., 2019) suggests importance-aware power and retransmission control.

Wireless communication for model aggregation. FL often assumes that agents communicate through wireless channels. Several works propose to utilize the superposition property of the wireless multiple access channel to allow all nodes in a wireless cell upload the model parameters at the same time, performing the averaging step of (11) over the air in the analog (Liu et al., 2020; Hellström et al., 2022) or in the digital (Razavikia et al., 2024) domain, including hierarchical federated learning (Azimi-Abarghouyi & Fodor, 2024). Over-the-air computation ensures that the communication costs are independent of the number of nodes. In fully decentralized learning, the broadcast nature of the medium can be used for efficient model exchange among neighbors (Pérez Herrera et al., 2025).

Client scheduling. To limit wireless resource usage while ensuring learning accuracy, (Nishio & Yonetani, 2019; Yang et al., 2019) suggests scheduling clients with good channel quality for each global update (line 3 in Algorithm 1). Recognizing that some gradient or model updates have higher importance for the learning process, combined importance and channel-aware client selection is proposed in several works, with different importance metrics (Goetz et al., 2019; Ren et al., 2020; Liu et al., 2021; Leng et al., 2022). Typically, client selection methods aim to maximize the learning accuracy in one global round, under a time constraint. Instead, client selection for minimizing the total learning time is proposed in (Chen et al., 2020; Nishio & Yonetani, 2019), balancing the time of a global round and the improvements achieved.

2.5 Privacy preserving Collaborative Learning

Collaborative learning techniques were initially motivated by their potential to preserve the privacy of local data S_k, by keeping it local (Bonawitz et al., 2022; Daly et al., 2024). Unfortunately, collaborative learning has been shown to still suffer from the exposure of partial features (Fredrikson et al., 2015; Phong et al., 2017; Melis et al., 2019) or even the complete leakage of training data (Fredrikson et al., 2015; Phong et al., 2017; Melis et al., 2019). Therefore, different levels of privacy protection are developed to protect training and inference data from other agents and the central server.

Blind computations. If training data is not shared, Algorithms 1 and 2 still require agents to exchange intermediate computations, which may need to be protected from information leakage. These intermediate values, such as local models, gradients, or intermediate representations, still carry information on S_k and can potentially be exploited to infer private information (Fredrikson et al., 2015; Phong et al., 2017; Chai et al., 2020). VFL scenarios are more subject to this information leakage since they require the exchange of intermediate representations of samples, whereas horisontal FL scenarios only aggregate models or gradients (Liu et al., 2024). One approach to prevent such an attack is to run the algorithm using *blind computation* techniques. They provide cryptographic guarantees that the operations on the private variables of the agents are performed without the possibility to read them. Notable techniques in the literature (Kairouz et al., 2021; Liu et al., 2024) include Trusted Execution Environments (TEEs) (Mo et al., 2021), Homomorphic Encryption (HE) (Chai et al., 2020), and Secure Multi-Party Computation (SMPC) (Liu et al., 2024).

Differential privacy. Even without access to raw data or intermediate computations, agents can observe the evolution of the global model parameters $\theta(t)$, which may reveal information about S_k. Prior knowledge on the training algorithm can be used to extract private information on S_k, for example, by querying the final model $F_{\theta(t)}$ (Fredrikson et al., 2015), or querying the collaborative learning algorithm (Melis et al., 2019; Zhu et al., 2019). To quantify and address these risks, the concept of *Differential Privacy (DP)* (Dwork et al., 2006; Dwork, 2006) can be applied. Originally proposed for centralized databases, DP offers rigorous formal guarantees against information leakage (Dwork & Roth, 2014).

While DP is a general-purpose privacy framework, it can be effectively adapted to any operations in machine learning that involve the release of aggregate statistics (Duchi et al., 2018; Liu et al., 2022). In particular, it is widely used to protect the training process, where repeated queries, like stochastic gradient computations, are privatized using *Differentially Private Stochastic Gradient Descent (DP-SGD)* and its variants (Song et al., 2013; Bassily et al., 2014; Abadi et al., 2016). These methods ensure strong privacy guarantees for the final model parameters by applying a sequence of privacy-preserving operations: subsampling (Kasiviswanathan et al., 2011; Bassily et al., 2014; Abadi et al., 2016) and random shuffling (Bittau et al., 2017; Erlingsson et al., 2019; Balle et al., 2019; Feldman et al., 2022) of the training data, clipping the gradients per-training sample (Andrew et al., 2021), and adding calibrated noise to the gradients (Dwork et al., 2006) before aggregation.

Differential privacy for collaborative learning. Distributed settings introduce new challenges for achieving global DP. Privacy-preserving primitives can be integrated at different levels of the Algorithm 1, which also depends on the risk model.

- *Local DP (LDP).* A first approach involves applying DP techniques directly to the local training algorithm to anonymize the gradient or model parameters $\theta_k(t)$ transmitted to the aggregator, for example, using DP-SGD and its variants. These techniques are designed for the case when participants aim to protect their data from an "honest-but-curious" central aggregator. While LDP provides strong privacy guarantees, it has high communication (Smith et al., 2017; Duchi & Rogers, 2019) and computation demands (Kasiviswanathan et al., 2011; Choi et al., 2018), and leads to the degradation of model estimates (Duchi et al., 2013; 2018; Bhowmick et al., 2018). This limited the practical adoption of LDP techniques (Bonawitz et al., 2022).
- *Distributed DP.* A second approach to protecting the privacy of local data from the central coordinator is to use Distributed DP techniques. These methods combine mild LDP mechanisms with server-side primitives for aggregating local statistics, which amplify this privacy level (Evfimievski et al., 2003). *Privacy amplification* is typically achieved through randomized shuffling and subsampling of the received LDP records (Bittau et al., 2017; Erlingsson et al., 2019; Cheu et al., 2019; Balle et al., 2020). In the context of Federated Learning, the goal is to perform distributed training while preserving the privacy of agents' training data. However, naively distributing DP-SGD is not possible, as the server lacks control over the local training data and local model updates. Therefore, DP variants of *FedAvg* have been proposed that aim to balance effective learning, privacy, and communication efficiency (McMahan et al., 2018; Agarwal et al., 2018). These variants incorporate server-side primitive such as clipping (Andrew et al., 2021), scaling (Geyer et al., 2017), shuffling (Geyer et al., 2017; Pihur et al., 2022), and subsampling (Geyer et al., 2017; Pihur et al., 2022; McMahan et al., 2018; Talwar et al., 2024) of the received gradients or model parameters. These primitives are coupled with local DP mechanisms, including additive noise and gradient quantization (McMahan et al., 2018). The current state-of-the-art technique, *Differential Privacy-Follow The Regularized Leader (DP-FTRL)* (Kairouz et al., 2021) proposes a different direction. It avoids sampling and shuffling by combining tree aggregation (Dwork et al., 2010; Chan et al., 2011) with correlated noise through online estimate regularization (Duchi et al., 2011). These techniques are demonstrated for large-scale federated learning deployments at Google (Ramaswamy et al., 2020; Xu et al., 2023), Apple (Apple, 2023), and Microsoft (Ding et al., 2017).

While most existing work focuses on privacy for centralized model aggregation, (El Mrini et al., 2024; Cyffers et al., 2024; Biswas et al., 2025) extends privacy mechanisms to decentralized architectures with

gossip learning. To enable DP in real systems, recent research proposes DP solutions with transparent aggregation (Talwar et al., 2024; Daly et al., 2024) and inference (Shumailov et al., 2025), and verifiable privacy mechanisms (Daly et al., 2024).

3 Collaborative Learning on Graph-structured Data

We have reviewed the main design choices of collaborative learning on Euclidean data in Section 2. We now extend these concepts to the case of *graph-structured data*. Since graph-structured data introduces unique characteristics and challenges, we first provide an overview of state-of-the-art solutions for *learning on graph-structured data* in Section 3.1. Building on this foundation, we apply the same analytical structure used for *Euclidean data* to the *graph domain*. We begin by formulating the *collaborative learning problem* for graph-structured data in Section 3.2, then examine the different types of *heterogeneity* that emerge in Section 3.3. Subsequently, we structure our analysis of design choices along the same *trilemma* identified in Section 2 and formalized in Figure 7: *learning effectiveness* in Section 3.4, *efficiency* in Section 3.5, and *privacy preservation* in Section 3.6. Table 2 summarizes the papers surveyed along these three lines.

3.1 Machine Learning on Graphs and GNNs

Problem definition. Machine Learning on graphs aims to develop models that operate on graph-structured data to solve tasks defined on graphs and generalize to unseen graph instances. Graphs are a fundamental data structure for modeling relationships between entities, which arise naturally in numerous domains such as social networks, molecular chemistry, and networked systems.

Formally, a graph is defined as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of *vertices*, also called *nodes*, and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ the set of *edges*, also called *links*. Given a fixed indexing of the nodes $\mathcal{V} = \{v_1 \ldots, v_N\}$, the graph structure can be conveniently encoded in an equivalent[5] matrix form as an adjacency matrix $A \in \mathbb{R}^{N \times N}$, where $a_{ij} = 1 \iff (v_i, v_j) \in \mathcal{E}$.

In practice, graphs often include auxiliary information. Most commonly, nodes are associated with features $\{x_v \in \mathcal{X}\}_{v \in \mathcal{V}}$, organized into a feature matrix $X \in \mathbb{R}^{N \times d_x}$, where d_x denotes the feature dimension. Depending on the task, label information may also be available: node-level labels for classification (e.g., user characteristics), edge-level labels for link prediction (e.g., friend recommendation), and graph-level labels for graph classification (e.g., molecule property).

A *graph sample* can then be expressed as (G, y) where $G = (A, X)$ represents the *graph data* in its matrix form and y the label associated with the task. When learning on graph data, we consider that we observe n graph data samples drawn from an unknown distribution $\mathcal{D}$, which together form a dataset $S = \{G^i, y^i\}_{i=1}^n$. The goal is then to learn a model F among our class of models $\mathcal{F}$ which minimizes the expected loss

$$\mathcal{L}(F) := \mathbb{E}_{(G,y)\sim\mathcal{D}} \ell(F(G), y). \tag{14}$$

Graph representation learning. Learning from graph data is achieved through the lens of *Graph Representation Learning* (Hamilton et al., 2017). The goal is to learn a mapping that embeds the vertices, edges, or the full graph into an Euclidean space, so that the embeddings capture useful information about the graph (see Figure 11). This mapping is usually formed by a composition of two functions. First, a learnable node encoder ϕ_{node} maps each node v in a graph G to an embedding $z_v := \phi_{node}(G, v) \in \mathbb{R}^{d_z}$ that captures both structural and feature information about v. These node embeddings $\{z_v; v \in \mathcal{V}\}$ can then be used to generate embeddings for edges and the graph, depending on the task:

- The edge embeddings can be computed by an encoder $\phi_{link}(z_u, z_v)$ (e.g., a simple concatenation), which can also be learnable (e.g., via an MLP).
- The graph embedding can be computed by an encoder on the multiset of all node representations $\phi_{graph}(\{\{z_v, v \in \mathcal{G}\}\})$ (e.g., a mean), which can be learnable (e.g., via an attention layer).

[5] A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ corresponds to an equivalence class of adjacency matrices $A \in \mathbb{R}^{N \times N}$, where two matrices are equivalent if they are related by a permutation of vertex indices; that is, for any permutation matrix on $\Pi \in \mathbb{R}^N$, we have $\Pi A \Pi^\top \sim A$.

Case	Subgraph								Multiple Graph Instances	
Graph type	Homogeneous				Heterogeneous				Homogeneous or heterogeneous	Heterogeneous
Feature partition	Horiz.		Vert.		Horiz.				Horiz.	Vert.
Type partition	Not applicable		Not applicable		No		Yes		Possible	Yes
Topology partition	Horiz.	Horiz.	Vert.	Vert.	Horiz.	Horiz.	Horiz.	Horiz.	No partitioning	Vert.
Granularity	Subgraph	Ego	edge-full	edge-free	Subgraph	Ego	Subgraph	Ego	Not applicable	Not applicable
Resulting graph cut	Edge	Edge	Vertex	Vertex	Vertex	Vertex	Vertex	Vertex	Not applicable	Not applicable
Representative work	407	357	271	67	155	373	290	335	159	88
Improving Learning and Inference Effectiveness (Section 3.4)										
Data-based techniques (Section 3.4.1)										
Data Augmentation	407; 292; 289	224; 137				63; 4; 238; 300 54	246		147; 409; 129	
Local Regularization			126				64		11; 246; 48	
Subsampling	56; 233	282; 90					290		48; 340; 324	
Model-based techniques (Section 3.4.2)										
Personalization of embeddings or parameters	221				155	63; 4; 364; 301			351; 409; 11; 383 112; 324	
Architecture Personalization		257; 201	126; 145; 246						409; 112; 324 159	
Training and Inference under vertical partitions (Section 3.4.3)										
Aggregating node representations	206	414	271; 64	67	155	63; 372; 406; 373 300; 301; 4 364			67	
Aggregating graph representations										88
Transfer techniques				145; 246						
Improving Communication and Computation Efficiency (Section 2.4)										
Aggregation patterns (Section 2.4.1)										
Centralized (Models)	394		53			63; 4; 238; 300 54; 290; 372; 373 233; 387; 364; 301 155; 406	64; 246; 290		351; 409; 11; 159 112; 324; 383; 147 129; 48; 340 67	
Centralized (Embeddings/Features)	394	90	53; 271; 61			63; 290; 372; 373 233; 54; 4; 387 364				
Decentralized Aggregation (Models)	319; 233; 414; 366 148; 128	137; 256; 357; 278							159	
Decentralized (Embeddings)	30; 401; 206; 366 301; 148	132; 416; 203; 30 281; 367; 137; 205 331; 282; 131; 130 201; 133; 265; 278 256; 357						335		
Hierarchical (Models)	148									
Hierarchical (Embeddings)		282; 265								
Asynchronicity (Section 3.5.2)										
Stale model aggregation										
Stale Message Passing		108; 331 249; 128; 367								
Communication (Section 3.5.3)										
Transmission Efficiency		203; 132; 357								
Wireless communication for aggregation	30	367; 130; 205; 131 123								
Client scheduling and topology management	233; 401; 265; 319	357								
Privacy (Section 3.6)										
For model parameters	394					238; 373; 387				
For exchanged embeddings	289	224; 282; 205	53; 271; 61; 62			372; 373; 300; 387 414; 238			409	
For topology	56; 148; 160	224; 282; 289	271			372; 63; 373; 300 387; 155; 4; 238 406; 403; 64			409	

Table 2: **The landscape of collaborative learning techniques on graph data organized by graph partition.** The table presents the design choices and techniques employed by different approaches across various graph partitioning schemes. Each row represents a specific technique or design choice, while columns correspond to different graph partitioning strategies categorized by graph type, feature partition, and granularity. Citations are placed in columns according to the partitioning approach used by each work.

The complete encoder model ϕ thus depends on the task at stake. In most models, only the node encoder is learnable, and ϕ_{node} is parametrized by these learnable weights $\theta \in \Theta$, expressed by the notation ϕ_θ. When labels are available, a predictor f_ψ, parameterized by $\psi \in \Psi$, is added to the output of the encoder such that $F_{\psi,\theta}$ is the complete model. The problem is then to identify F_{ψ^*,θ^*} which minimizes

$$\mathcal{L}(F_{\psi,\theta}) := \mathbb{E}_{(G,y)\sim\mathcal{D}}\ell(F_{\psi,\theta}(G), y). \tag{15}$$

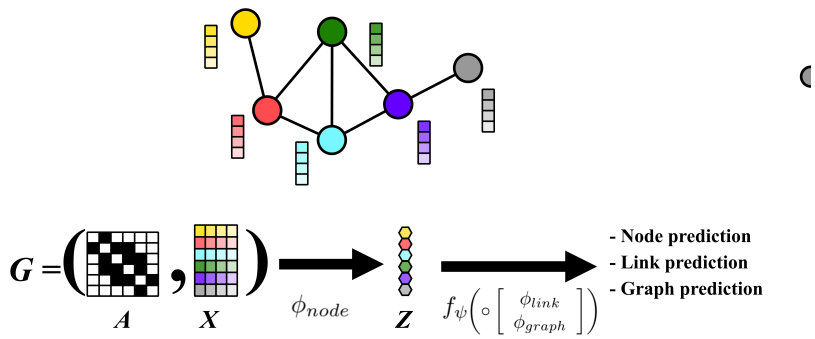

Figure 11: **Graph representation learning.** The objective is to learn a node encoder ϕ_{node} that embed the graph data $G = (A, X)$, consisting of the adjacency matrix $A \in \mathbb{R}^{N\times N}$ and the feature matrix $X \in \mathbb{R}^{N\times d_x}$, into a node representation matrix $Z \in \mathbb{R}^{N\times d_z}$. These node representations are then used for predictive tasks at the node, link, or graph level, with the use of a predictor f_ψ and, depending on the task a link ϕ_{link} or graph ϕ_{graph} encoder.

Graph Neural Networks (GNNs). Graph Neural Networks represent the dominant node encoder ϕ_{node} architecture. Two architectural paradigms are commonly distinguished (Wu et al., 2020):

- *Spectral GNNs* apply spectral filters to the graph. They operate on the full spatial domain and a truncated spectral domain (Geisler et al., 2024).
- *Spatial GNNs* apply spatial filters to the graph. They operate on the truncated spatial domain and the full spectral domain (Geisler et al., 2024).

Modern GNN architectures predominantly adopt the *Message Passing Neural Network (MPNN)* framework (Gilmer et al., 2017) (see Figure 12), which is a spatial approach to define ϕ_{node}. To encode a node $v \in \mathcal{V}$ in a representation z_v, an MPNN follows an iterative process that aggregates the information of v's L-hop neighborhood. This neighborhood, denoted as G^v, represents the so-called L-hop *ego-graph* centered on node v. The iterative process is as follows:

- *Feature encoding:*
$$h_v^{(0)} \leftarrow f_{Enc}(x_v) \tag{16}$$
- *Message Passing:* for each layer $l = 0, \ldots, L-1$,
$$m_{v,u}^{(l+1)} = f_{Mes}^{(l+1)}(h_v^{(l)}, h_u^{(l)}) \qquad \text{(message computation)} \tag{17}$$
$$m_v^{(l+1)} = f_{Agg}^{(l+1)}(\{\{m_{v,u}^{(l+1)}, u \in \mathcal{N}(v)\}\}) \qquad \text{(message aggregation)} \tag{18}$$
$$h_v^{(l+1)} = f_{Up}^{(l+1)}(h_v^{(l)}, m_v^{(l+1)}) \qquad \text{(state update)} \tag{19}$$

  where $\mathcal{N}(v)$ denotes the set of 1-hop neighbours of v in $\mathcal{G}$, $f_{Enc}, f_{Mes}^{(l+1)}$, $f_{Up}^{(l+1)}$ are neural functions (e.g., MLPs), and $f_{Agg}^{(l+1)}$ is a permutation-invariant operator on multisets (e.g., sum, mean, attention).

- *Final Embedding:*

$$z_v := h_v^{(L)} = \phi_{node}(G^v, v). \tag{20}$$

The full encoder ϕ_θ can be considered to be parametrized by θ, representing all learnable weights across message, aggregation, and update functions and across all layers $l = 0, \ldots, L$. The framework is modular: different architecture choices for f_{Mes}, f_{Agg} and f_{Up} define the most famous GNN architectures (e.g., GCN (Kipf & Welling, 2017), GraphSAGE (Hamilton et al., 2017), GAT (Veličković et al., 2018)). The resulting node embeddings $\{z_v\}_{v\in\mathcal{V}}$ can then be fed to the predictor for solving downstream tasks. In the remainder of this paper, we assume that, unless specified, MPNN is the considered GNN architecture.

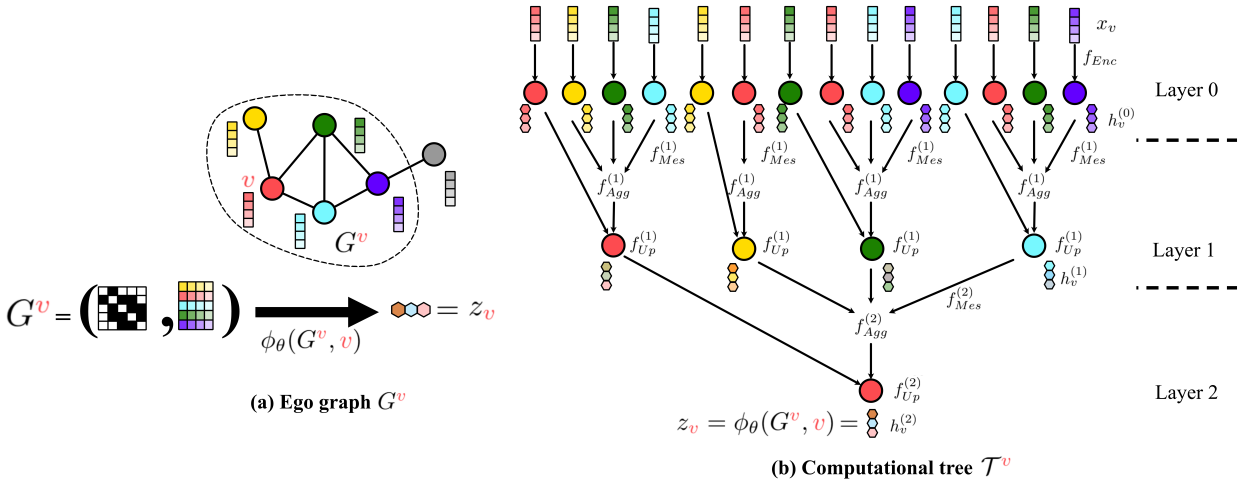

Figure 12: **Computation in a Message Passing Neural Network (MPNN) ϕ_θ.** To infer the representation z_v of a node v, the MPNN takes as input the ego-graph G^v centered on v, illustrated in (a) for the case of a 2-layer MPNN. The inference process consists of a sequence of hierarchical steps, given by (16),(17),(18), and (19), which together form the computational tree $\mathcal{T}^v$ shown in (b). In practice, these operations can be efficiently parallelized to infer the representations z_v of a batch of nodes.

GNN training. MPNNs can be trained to approximate the solution of (15) on empirical samples by minimizing the loss

$$\hat{\mathcal{L}}(F_{\psi,\theta}) = \frac{1}{|S|} \sum_{(G,y)\in S} \ell(F_{\psi,\theta}(G), y). \tag{21}$$

Minimizing (21) is commonly done using the SGD algorithm and backpropagation techniques (Rumelhart et al., 1986; Hamilton et al., 2017). In practice, most approaches minimize the empirical risk (21) as a tractable proxy for the true risk (15), assuming that the resulting model generalizes well, i.e., that the generalization gap is small (Jegelka, 2022).

A notable characteristic of the MPNN architecture is that both the model parameters and the computed gradients need to be propagated along the edges of the graph. MPNNs share strong similarities with distributed message passing algorithms (Papp & Wattenhofer, 2022), which iteratively aggregate messages from their neighbors, update their state, and then pass on the new states to their neighbors again. This suggests that MPNNss are a natural fit for Machine Learning tasks on networked systems.

3.2 Collaborative Learning on Graph-structured Data Problem Definition

As pointed out for Euclidean data in Section 2, the term distributed ML often encapsulates two techniques: parallel training on distributed machines and training on private distributed data. Our settings focus on the

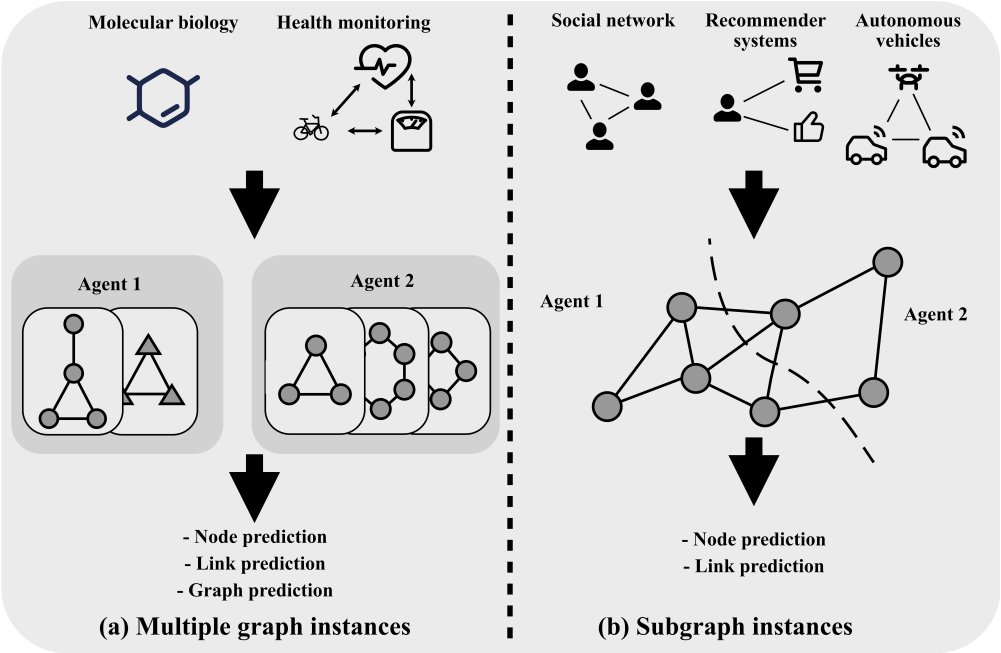

Figure 13: Collaborative learning on graph data can be broadly categorized into two settings: (a) agents holding multiple graph instances, and (b) agents holding subgraphs.

latter, where the graph data is collected locally by different agents. For the parallel training case, we refer the reader to (Shao et al., 2024).

We now extend the collaborative learning framework introduced in Section 2.1 to learning on graph data. Two main problem settings are distinguished: *learning from multiple graph instances* and *learning from subgraphs*, as shown in Figure 13.

Learning from multiple graph instances. We consider a set of agents $\mathcal{K} = \{1, \ldots, K\}$, where each agent k holds a dataset S_k containing n_k graphs sampled from an unknown local distribution $\mathcal{D}_k$. For example, the dataset can be made of numerous molecule graphs known by different labs. Usually, graph-level tasks are considered for this case, for example, predicting a property of a molecule graph (Xie et al., 2021). Each agent aims to minimize the local objective

$$\mathcal{L}_k(F_{\psi,\theta}) = \mathbb{E}_{(G,y)\sim\mathcal{D}_k} \ell(F_{\psi,\theta}(G), y). \tag{22}$$

Agents can compute local estimates of the local objectives

$$\hat{\mathcal{L}}_k(F_{\psi,\theta}) = \frac{1}{n_k} \sum_{(G,y)\in S_k} \ell(F_{\psi,\theta}(G), y). \tag{23}$$

One can notice that the local objectives (22) (23) are formulated in the same form as for Euclidean data in (1) (2), with the difference here that the data distribution $\mathcal{D}_k$ operates on graph data. As all the other aspects are the same as in Euclidean settings, if the distributions $\mathcal{D}_k$ are i.i.d., then the *FedAvg* algorithm can be used to effectively identify the optimal GNN parameters, the solution of (3).

Learning from subgraphs. We now consider the subgraph case. We consider a global graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where each agents $k \in \mathcal{K}$ observes a set of N_k nodes and the edges connecting these nodes $\mathcal{E}_k = \{(u, v) \in \mathcal{E} | u, v \in \mathcal{V}_k\}$ $\mathcal{V}_k \subset \mathcal{V}$, forming a subgraph topology $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E}_k)$. Each agent collects subgraph data G_k related to its observed nodes $\mathcal{V}_k$, and some labels y associated with the task. Learning for recommender systems, large knowledge graphs, social networks, or large-scale infrastructures, but also learning in groups

of interactive objects, falls into this category. The learning tasks are typically on the node, link, or subgraph level.

Recall that the GNN requires, for each node v, the data from the L-hop neighborhood of v as formulated in (20). The local empirical objective for agent k is then expressed not on the graph data G_k, but on the graph data supported on $\mathcal{G}_k$ expanded by an L-hop neighborhood within the global graph. We denote this expanded topology by $\tilde{\mathcal{G}}_k$, and the corresponding data supported on it by $\tilde{G}_k$

$$\hat{\mathcal{L}}_k(F_{\psi,\theta}) := \ell(F_{\psi,\theta}(\tilde{G}_k), y). \tag{24}$$

If agent k does not observe $\tilde{G}_k$, then it does not have the necessary data to minimize its local loss (24). This makes the collaboration of the agents that hold parts of $\tilde{G}_k$ necessary, a phenomenon illustrated in Figure 14.

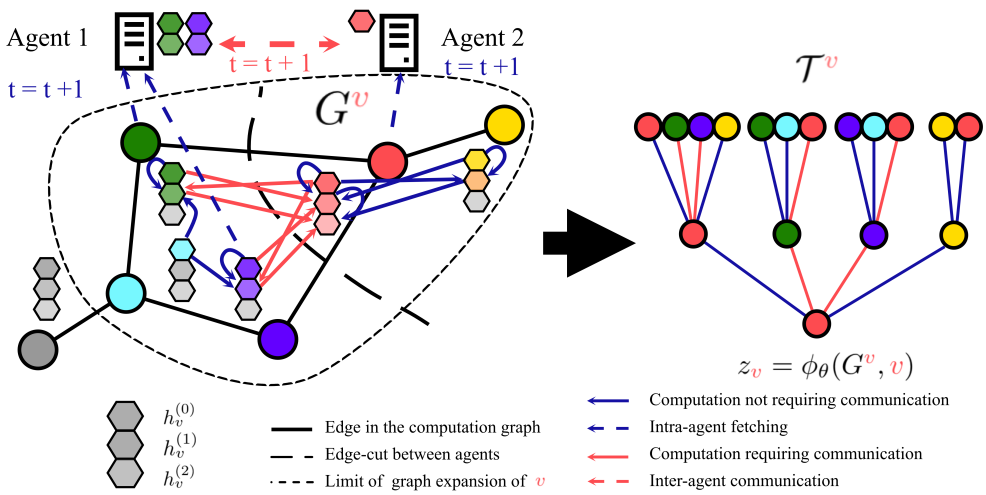

Figure 14: **Collaborative inference in a subgraph.** In the subgraph setting, inference becomes collaborative as described in Algorithm 3, illustrated here for the inference of the red node v's representation based on its L-hop ego-graph G^v. The computational tree $\mathcal{T}^v$ shows how computations can be distributed across multiple agents, requiring the exchange of intermediate results between them, hence the term *collaborative inference*.

This need for collaborative inference resembles the setting of vertical partitions described in Section 2.3.3 and leads to the following observation: *Collaborative training of GNN borrows from the principles of VFL.* If we assume that $(\tilde{G}_k, y)$ are sampled from an unknown distribution $\mathcal{D}_k$, the true local objective can be formulated in a collaborative form

$$\mathcal{L}_k(F_{\psi,\theta}) := \mathbb{E}_{(\tilde{G}_k y)\sim\mathcal{D}_k} \ell\left(F_{\psi,\theta}(\tilde{G}_k), y\right). \tag{25}$$

An efficient solution to minimize the local objectives (9) is to collaboratively train a common GNN $F_{\psi,\theta}$ to minimize the average of local objectives across agents

$$\frac{1}{K}\sum_{k\in\mathcal{K}} \mathcal{L}_k(F_{\psi,\theta}). \tag{26}$$

This formulation is analogous to the Horizontal FL problem formulation (2) with the difference that the local losses $\mathcal{L}_k$ are collaborative, as thus the subgraph FL problem can also *borrow from the principles of horisontal FL.*

A GNN can be trained to minimize (26) in collaborative settings by following two algorithmic directions:

- With *isolated inference*, inference is performed on the local subgraph only. This case borrows directly from the Horizontal FL Algorithm 1.
- With *collaborative inference*, inference requires the information exchange between the agents. This case borrows directly from the Vertical FL Algorithm 2.

A possible unified implementation of federated training from subgraph data is described in Algorithm 3, following the concepts from (Chen et al., 2021; Wu et al., 2022; Du & Wu, 2022; Chen et al., 2022; Yao et al., 2023; Han et al., 2024). It is a round-based algorithm which outputs both the model parameters $\psi(T), \theta_k(T)$, and the node representations $\{z_v(T)\}_{v\in\mathcal{V}}$ (line 25) at the end of the training round. The algorithm builds on the centralized aggregation of Algorithm 1, where model parameters iterate $\psi(t)$ and $\theta_k(t)$ are aggregated by a central server (line 22). In the subgraph setting, however, the server also assumes additional responsibilities: it stores the node representation iterates $z_v(t)_{v\in\mathcal{V}}$ (line 23) and allows the exchange of node embeddings between agents during collaborative inference (lines 11 and 12).

In the algorithm, for agents $k \in \mathcal{K}$, G_k is the local subgraph data defined on $\mathcal{V}_k$, the set of nodes held by the agent k. G_k^v is the part of the ego-graph data held by agent k and related to v. We also need to define boundary nodes, that is, nodes that are connected to nodes held by different agents. For each agent k, the set of foreign boundary nodes is the set of nodes held by a different agent and connected to a node seen by k: $\mathcal{V}_{f,k} = \{v; \exists u \in \mathcal{V}_k \mid (u,v) \in \mathcal{E}, v \in \mathcal{V} \setminus \mathcal{V}_k\}$. The set of private boundary nodes is the set of nodes private to k that are foreign nodes of other agents: $\mathcal{V}_{p,k} = \bigcup_{l\in\mathcal{K}\setminus k}\{v \in \mathcal{V}_k \mid v \in \mathcal{V}_{f,l}\}$.

A round of Algorithm 3 begins with the broadcast of common model parameters to all agents (line 2), which then initialize their local model parameters accordingly (line 4). Two inference modes are considered:

- *Isolated inference:* each agent performs inference solely on its local subgraph data $\mathcal{G}_k$ (line 6).
- *Collaborative inference:* agents collectively replicate the inference process of a message-passing neural network (MPNN) in a distributed fashion. This is executed iteratively over the L layers of the GNN. At each iteration l, agents communicate (line 11), and request (line 12) the l-th layer embeddings of their foreign boundary nodes $\mathcal{V}_{f,k}$ through the server, then compute the $(l+1)$-th embeddings using the MPNN inference equations (17), (18), and (19).

The remaining steps are common to both inference modes. Node representations are sent to the server for aggregation (lines 17–23). Local model parameters are updated using gradient descent based on the computed loss (line 19), and subsequently sent to the server for aggregation (line 11).

3.3 Graph Data Heterogeneity

The homogeneity of graph data collected across agents is central to the problem setup in (22)–(25). However, data heterogeneity, which is well-studied for Euclidean data in Section 2.2, also challenges collaborative learning in the graph setting. We now describe graph-specific heterogeneity from two angles: data partitioning and statistical heterogeneity.

3.3.1 Graph Partition Taxonomy

A key feature of distributed learning is how the data is partitioned across agents. We propose subsequently a characterization of the different partitions for graph data appearing in the literature. To structure our analysis, we build on the taxonomy for Euclidean data partitions (Section 2.2.1) and adapt it to graph data, following prior efforts in (Zhang et al., 2021; He et al., 2021; Liu et al., 2025). We continue to discuss the two main cases introduced in Section 3.2: (i) agents holding multiple graph instances, and (ii) agents holding a subgraph as described in Figure 13.

[7] In practice, this step can be efficiently computed in parallel using vectorization libraries for GNN such as PyTorch Geometric (Fey et al., 2025).

Algorithm 3 Federated training of Subgraph GNN

Require: $\mathcal{K} = \{1, \ldots, K\}$ set of agents , $F_{\psi,\theta}$: a GNN parametrized by $\psi \in \Psi, \theta \in \Theta$, global graph topology $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $\{\mathcal{V}_k\}_{k\in\mathcal{K}}$ subset of nodes $\mathcal{V}$ held by each agent, foreign boundary nodes $\{\mathcal{V}_{p,k}\}_{k\in\mathcal{K}}$ known by each agent, private boundary nodes $\{\mathcal{V}_{f,k}\}_{k\in\mathcal{K}}$ known by each agent, boundary edges $\{\mathcal{E}_k\}_{k\in\mathcal{K}}$ known by each agent, learning rate η, initial model parameters $\{\psi(0), \theta(0)\}$ and initial node embeddings $\{z_v(0)\}_{v\in\mathcal{V}}$, local graph data $\{G_k\}_{k\in\mathcal{K}}$ held by each agent, number of training rounds T, boolean *Isolated inference*, *Collaborative inference*

— **Training loop** —

for each round $t = 1, \ldots, T$ **do**

Server:

Broadcast current model parameters $\theta(t-1), \psi(t-1)$ to all agents

for each agent $k \in \mathcal{K}$ in parallel **do**

— **Local Model Initialization** —

Initialize $F_{\psi_k(t-\frac{1}{2}),\theta_k(t-\frac{1}{2}),}$ such as $\theta_k(t-\frac{1}{2}) \leftarrow \theta(t-1)$, $\psi_k(t-\frac{1}{2}) \leftarrow \psi(t-1)$

— **Inference Phase** —

if *Isolated Inference* **then**

— **Local Inference on Isolated Subgraph** —

Decompose G_k into $\{G_k^v\}_{v\in\mathcal{V}_k}$ and compute $z_v(t-\frac{1}{2}) = \phi_{\theta_k(t-\frac{1}{2})}(G_k^v, v)$ for each node[6] $v \in \mathcal{V}_k$ as in (20)

end if

if *Collaborative inference* **then**

— **Collaborative Inference** —

Compute $\{h_v^{(0)}\}_{v\in\mathcal{V}_k}$ as in (16) and set $\forall v \in \mathcal{V}_k, h_v^{(0)}(t) \leftarrow h_v^{(0)}$

for $l = 0, \ldots, L-1$ **do**

Send $\{h_v^{(l)}(t)\}_{v\in\mathcal{V}_{p,k}}$ to the server

Request $\{h_v^{(l)}(t-1)\}_{v\in\mathcal{V}_{f,k}}$ from the server

Upon reception, compute (17) (18) (19) for each[7] $v \in \mathcal{V}_k$

end for

$z_v(t-\frac{1}{2}) \leftarrow h_v^{(L)}$ for all $v \in \mathcal{V}_k$

end if

Send $\{z_v(t-\frac{1}{2})\}_{v\in\mathcal{V}_k}$ to the server

— **Loss Computation** —

Compute $\hat{\mathcal{L}}_k\left(F_{\psi_k(t-\frac{1}{2}),\theta_k(t-\frac{1}{2})}\right)$

— **Local Model Update** —

$\theta_k(t) \leftarrow \theta_k(t-\frac{1}{2}) - \eta \cdot \frac{\partial \hat{\mathcal{L}}_k}{\partial \theta_k^{(l)}(t-\frac{1}{2})}, \quad \psi_k(t) \leftarrow \psi_k(t-\frac{1}{2}) - \eta \cdot \frac{\partial \hat{\mathcal{L}}_k}{\partial \psi_k^{(l)}(t-\frac{1}{2})}$

Send $\theta_k(t), \psi_k(t)$ to the server

end for

— **Server Aggregation** —

Server:

$\theta(t) \leftarrow \sum_{k\in\mathcal{K}} \frac{N_k}{N}\theta_k(t), \psi(t) \leftarrow \sum_{k\in\mathcal{K}_t} \frac{N_k}{N}\psi_k(t)$

$\{z_v(t)\}_{v\in\mathcal{V}} \leftarrow \bigcup_{k\in\mathcal{K}}\{z_v(t-\frac{1}{2})\}_{v\in\mathcal{V}_k}$

end for

— **Final Output** —

return $\psi(T), \theta(T), \{z_v(T)\}_{v\in\mathcal{V}}$

Multiple graph instances. We examine the case where agents observe full graph instances as introduced in Section 3.2. To formalize, we assume the existence of a global dataset S made of full graph instances $S = \{(G^i, y^i)\}_{i\in\mathcal{I}}$. We distinguish two ways in which graph instances are typically distributed across agents, illustrated in Figure 15:

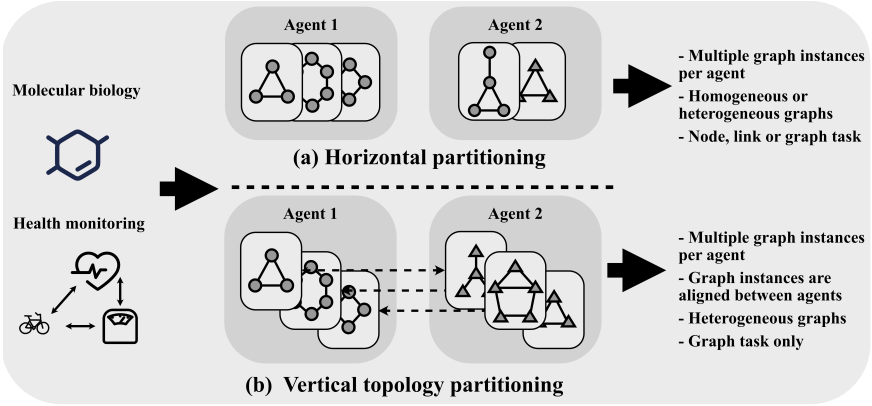

Figure 15: **Illustration of the different partitioning schemes for the case of multiple graph instances**. Multiple graph instances can be partitioned (a) horizontally, where agents hold unrelated graph samples, or (b) vertically, where graph instances associated with the same label are divided among agents.

- *Horizontal topology partition* analogous to the taxonomy of *inter-graph FL* (Xie et al., 2021) or *graph level FL* (He et al., 2021). Similar to horizontal partitioning seen in Section 2.2.1, the index $\mathcal{I}$ of S is partitioned across agents into local indexes $\mathcal{I}_k \subset \mathcal{I}$ such as $\cup_{k \in \mathcal{K}} \mathcal{I}_k = \mathcal{I}$, yielding datasets for agent k:

$$S_k = \left\{((A^i, X^i), y^i)\right\}_{i \in \mathcal{I}_k}.$$

  That is, agents possess data about distinct graph samples indexed by $\mathcal{I}_k$, for example, laboratories observe distinct sets of protein graphs (Xie et al., 2021; He et al., 2022).

- *Vertical topology partition* For each sample of graph data G^i, the topology $\mathcal{G}^i$ is partitioned into K disconnected components $\mathcal{G}^i_k$ held by different agents. The node features X^i are split so that agent k observes the features of the nodes $\mathcal{G}^i_k$. The label, usually a graph-label for these cases, is held by agent k^*, yielding datasets:

$$S_k = \{G^i_k\}_{i \in \mathcal{I}}; \quad S_{k^*} = \{G^i_k, y^i\}_{i \in \mathcal{I}}.$$

  This partitioning is typical in multimodal scenarios where the graph G^i_k represents one modality k for the same sample i. These modalities are combined to solve a task related to the same sample G^i. An example is to diagnose patients based on the combination of different physiological observations, represented as graphs and held by different hospitals (Dong et al., 2021). This vertical topology partition is studied for centralized training (Qian et al., 2022), but has not yet been implemented in distributed training settings.

Subgraphs of a global graph. We now consider a setting where agents observe parts of a single, usually large-scale, global graph. In such scenarios, agents have only partial local views of the global graph, making learning over isolated graph data a central challenge.

In the most general case, the global, possibly heterogeneous graph, is partitioned across agents, such as each agent k observes a local heterogeneous subgraph $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E}_k, \mathcal{O}_k, \mathcal{R}_k)$. That is, each agent observes a subset of nodes $\mathcal{V}_k \subseteq \mathcal{V}$ of a subset of types $\mathcal{O}_k \subseteq \mathcal{O}$, and links $\mathcal{E}_k \subseteq \mathcal{E}$ of subset of relation types $\mathcal{R}_k \subseteq \mathcal{R}$. These may induce a restricted type and relation mapping $\rho_k : \mathcal{V}_k \to \mathcal{O}_k$ and $\rho_k : \mathcal{E}_k \to \mathcal{R}_k$. For example, in a citation network, we might have object types $\mathcal{O} = \{\text{Author}, \text{Paper}, \text{Venue}\}$ and relation types $\mathcal{R} = \{\text{Writes}, \text{PublishedIn}, \text{Cites}\}$.

A graph is characterized by three main components: its topology, the types of nodes and relations along with their mappings, and the node features and potential labels. Accordingly, partitions can be defined along these dimensions, giving rise to the following taxonomy of graph partitions illustrated in Figures 16– 17.

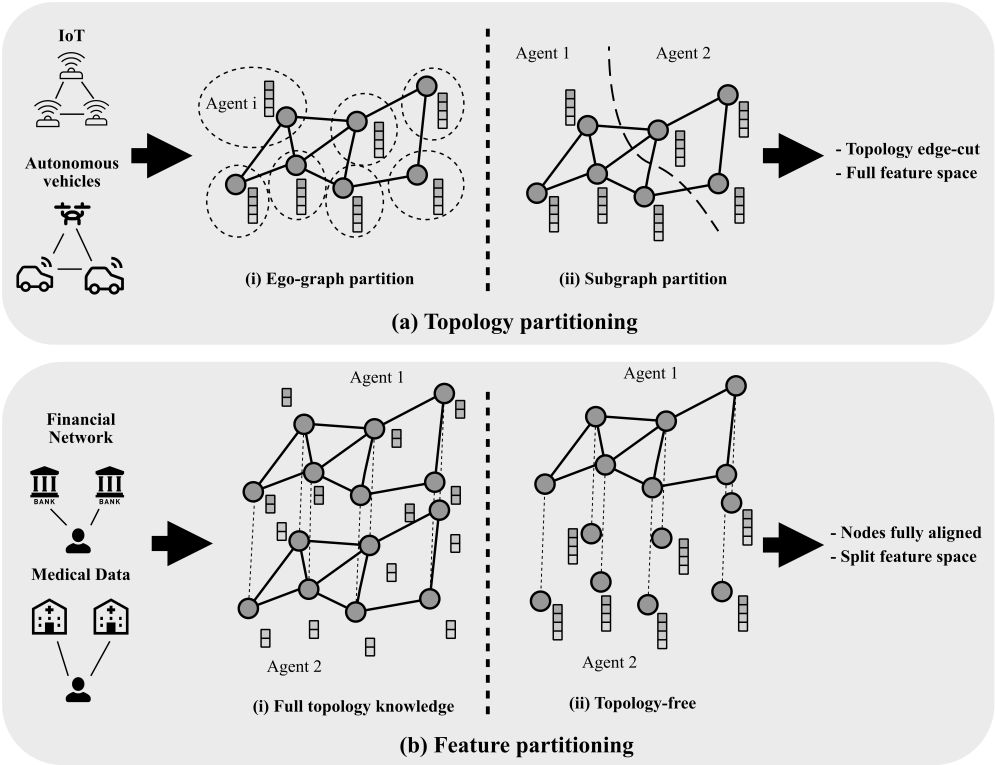

Figure 16: **Illustration of different subgraph partitioning schemes in a homogeneous graph.** In *topology partitioning* (a), the graph topology is divided across agents by edge cuts. Each agent observes the complete set of node features for the nodes assigned to it. Two cases can be distinguished based on the subgraph size: *ego-graph partition* (i), where an agent only observes a node together with its direct neighbors, thus forming an ego-graph; and *subgraph partition* (ii), where an agent holds a larger portion of the graph consisting of multiple, aligned nodes. In *feature partitioning* (b), agents share the same set of nodes but access different subsets of features. Two scenarios arise: in *full topology knowledge* (i), all agents know the complete set of graph links; whereas in the *topology-free* case (ii), only a subset of agents has access to link information.

- *Topology partition.* The partition of the graph is done along the graph structure $\mathcal{H}_{\mathcal{G}} = (\mathcal{V}, \mathcal{E})$, without considering the node types, and is achieved by :
  - *Edge-cut:* Nodes are divided into (possibly overlapping) sets $\bigcup_{k\in\mathcal{K}} \mathcal{V}_k$, inducing local edge sets $\mathcal{E}_{\mathcal{V}_k\times\mathcal{V}_k} = \{(u,v) \in \mathcal{E} \mid (u,v) \in \mathcal{V}_k \times \mathcal{V}_k\}$. For example, consider a network of IoT devices related by a global graph. Telecom base stations monitor different geographical areas, resulting in a different set of devices observed (Zeng et al., 2022). Usually, agents also know about their foreign nodes $\mathcal{V}_{f,k} = \{v; \exists u \in \mathcal{V}_k \mid (u,v) \in \mathcal{E}, v \in \mathcal{V} \setminus \mathcal{V}_k\}$ and the related edges cut $\mathcal{E}_{cut,k} := \mathcal{E}_{\mathcal{V}_k\times\mathcal{V}_{f,k}} = \{(u,v) \in \mathcal{E} \mid u \in \mathcal{V}_k \; v \in \mathcal{V}_{f,k}\}$ but not their features.
  - *Vertex-cut:* Edges are divided into (possibly overlapping) sets $\bigcup_{k\in\mathcal{K}} \mathcal{E}_k$, inducing node sets $\mathcal{V}_k = \mathcal{V}_{|\mathcal{E}_k} = \{u \in \mathcal{V} \,|\, \exists v \in \mathcal{V},\ (u,v) \in \mathcal{E}_k \text{ or } (v,u) \in \mathcal{E}_k\}$. For example, a device records the items browsed by a single user in a marketplace, forming an interaction graph private to the device (Wu et al., 2022; Han et al., 2024). We define the set of *vertex-cut*, vertices which are shared between agents: $\mathcal{V}_{\text{cut},k} = \{v \in \mathcal{V}_k \,|\, \exists k' \neq k,\ v \in \mathcal{V}_{k'}\}$. The interactions represent the edges private to the agent, and some items represent the vertex cut between agents.

- *Feature partition.* Analogous to vertical partitioning for Euclidean data seen in Section 2.2.1, the feature space $\mathcal{X}$ is split among agents $\mathcal{X} = \oplus_{k\in\mathcal{K}}\mathcal{X}_k$, resulting in *vertical feature partitioning* of the features matrix X. Specifically, for each object type $o \in \mathcal{O}$, the feature matrices X_o is partitioned

across feature dimensions: for each node v of object type o, $x_v = (x_{v,k})_{k \in \mathcal{K}}$, where $x_{v,k} \in \mathcal{X}_k$ is the part of vector x_v observed by agent k.. Feature partition may result topology partition as well, which can be classified as a vertical topology partition, with a vertex cut.

- *Schema partition.* In a heterogeneous graph $\mathcal{G}$, partitioning can be defined along the network schema $\mathcal{T}_\mathcal{G} = (\mathcal{O}, \mathcal{R})$, by partitioning:
  - *The object types:* Partitioning the set of object types $\mathcal{O}$ into disjoint subsets $\{\mathcal{O}_k\}_k$, i.e., $\mathcal{O} = \bigcup_k \mathcal{O}_k$, induces a corresponding partition of the node set in the original heterogeneous graph via the inverse mapping of the type function $\tau : \mathcal{V} \to \mathcal{O}$, namely $\mathcal{V}_k = {\mathcal{O}_k}^{-1} := \tau^{-1}(\mathcal{O}_k) = \{v \in \mathcal{V} \mid \tau(v) \in \mathcal{O}_k\}$. Object type based partition often leads to edge cuts.
  - *The relation types:* Similarly, partitioning the set of relation types $\mathcal{R}$ into disjoint subsets $\{\mathcal{R}_k\}_k$, i.e., $\mathcal{R} = \bigcup_k \mathcal{R}_k$, induces a partition of the edge set via the inverse mapping of the relation function $\rho : \mathcal{E} \to \mathcal{R}$, yielding $\mathcal{E}_k = {\mathcal{R}_k}^{-1} := \rho^{-1}(\mathcal{R}_k) = \{e \in \mathcal{E} \mid \rho(e) \in \mathcal{R}_k\}$. Relation type based partition often leads to vertex cut.

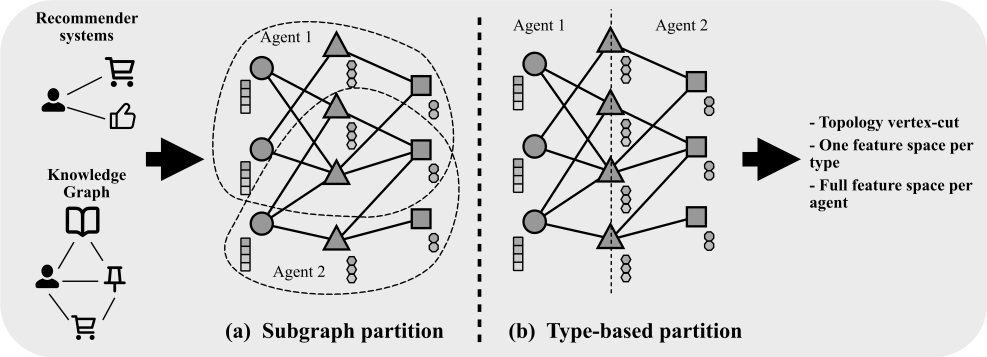

Figure 17: **Illustration of different subgraph partitioning schemes in a heterogeneous graph**. The subgraph can be partitioned at two levels: either by dividing the node set $\mathcal{V}$ among agents, or by partitioning the schema graph, partitioning the object types $\mathcal{O}$ across agents.

Graph partition of topology, schema, or features can be combined freely, leading to a high number of possible combinations of data distribution across the agents. In Table 2 we list the combinations that are most prominent in the literature.

- Considering *homogeneous graphs*, *topology partition*, resulting horizontal feature partition, is the most frequently studied case. Agents may hold a *subgraph* of the global graph due to geographic limitations, for example, in mobile networks, or due to collaborating system providers in IoT systems (Zhang et al., 2021). The other typical case of topology partition is when agents hold a single node and perform inference through the *ego-graph*. This is the typical case of interacting autonomous vehicles, robots, or drones (Pan et al., 2023). Traffic flow forecasting (Zeng et al., 2022), location-based user recommendation (Zeng et al., 2022), vehicle motion forecasting (Casas et al., 2020), or swarm navigation (Zhou et al., 2021; Solodova et al., 2025; Gao et al., 2021; Blumenkamp et al., 2022) are prominent tasks.

- *Vertical feature partition* of *homogeneous graphs* arises when agents hold partial information about the nodes, for example, spending habits of customers are known partially by e-commerce sites (Ni et al., 2021; Chen et al., 2024; He et al., 2024), and hospitals have partial records of patients (Fu et al., 2024). In most scenarios, *all agents hold a graph* (Ni et al., 2021; Chen et al., 2024; He et al., 2024), but alternatively, some agents may be topology-free, holding only Euclidean data (Mei et al., 2019; Cheung et al., 2021; Fu et al., 2024).

- Key examples of *heterogeneous graphs* are knowledge graphs (Chen et al., 2022; Wang et al., 2024) and recommender systems (Agrawal et al., 2024; Han et al., 2024; Liu et al., 2022; Qiu et al., 2022;

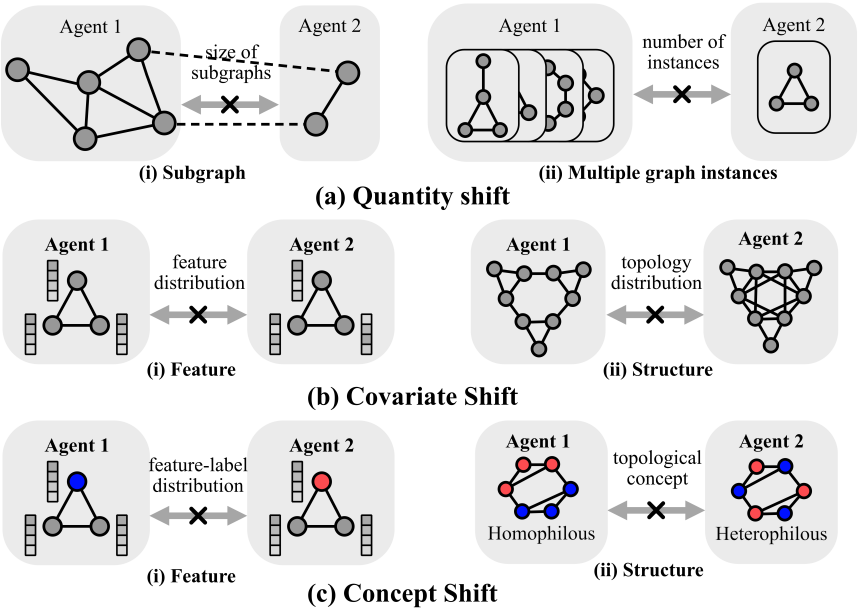

Figure 18: **Different types of statistical heterogeneity in partitioned graph data.** (a) *Quantity shift* refers to differences in dataset size across agents sampling from the same distribution, reflected in varying (i) subgraph sizes or (ii) numbers of graph instances. (b) *Covariate shift* refers to differences in input graph distributions, reflected in (i) node or edge features or (ii) graph structures. (c) *Concept shift* refers to differences in the mapping from input graphs to labels (colored on the figure), reflected in (i) *feature–concept* or (ii) *structure–concept* shifts.

Yan et al., 2024). Heterogeneous graphs can be distributed without considering the type of nodes, resulting in *topology partition*. Within that, *subgraph partition* emerges when companies observe the ego-graphs of a group of customers (Han et al., 2024). The requirement to preserve node privacy, on the other hand, may lead to *ego-graph* partition, where one agent knows about the ego graph of a single customer (Agrawal et al., 2024).

- *Type-based partition* in *heterogeneous graphs* is typical in knowledge graphs, for example, when one agent holds a patient-symptom database, and another one a symptom-disease database (Peng et al., 2021), with a common goal of training a patient-to-disease predictor, leading to agents holding *type-based subgraphs*. Privacy preservation under collaborating agents, for example, in health applications, can require a combination of *type-based* and *ego-graph* partition (Sun et al., 2020).

3.3.2 Statistical Heterogeneity in Graphs

Section 2.2.2 outlines the main forms of statistical imbalance in Euclidean data, *quantity skew* and *distributional shift*, which in turn include *covariate shift* and *concept shift*. In the Euclidean setting, data is typically drawn from a joint distribution $\mathcal{D}(x, y)$. In contrast, for graph data, the feature matrix X is structured through an adjacency matrix, which together forms a graph sample $G = (A, X)$. The coupling between topology, features, and labels is captured by the joint distribution $\mathcal{D}(G, y)$, and brings new dimensions in the possible forms of statistical imbalance in the data, illustrated in Figure 18

Quantity shift can be observed in graph data when agents hold different amounts of graph information. In the subgraph case, this can be observed when the size of subgraphs, measured by the number of edges or nodes, differs across agents. In the full graph instances case, quantity shift can be observed when agents hold different numbers of graph samples.

Distributional shifts bring more new research challenges, especially in the context of out-of-distribution generalization for GNNs (Zhang et al., 2024). Below we define graph specific distributional shifts under the covariate and concepts shift categories.

Covariate shift. Covariate shift refers to differences in the input graph distribution $\mathcal{D}(G)$ across agents, while the conditional label distribution $\mathcal{D}(y \mid G)$ remains unchanged. Alternatively, one can view it from the label perspective as discussed in Section 2.2.2, where $\mathcal{D}(y)$ and $\mathcal{D}(G \mid y)$ differ across agents. We adopt the input perspective, which is more common in the literature, and introduce graph-specific covariate shifts, *feature shift* and *structure shift* (Zhang et al., 2024).

- *Feature shift* for graph data is analogous to *feature shift* for Euclidean data seen in Section 2.2.1. It occurs when, given a graph topology A, the node features distribution $\mathcal{D}(X \mid A)$ differs (Zhang et al., 2024). For example, in superpixel graphs of natural scenes (Monti et al., 2017), one agent may use grayscale images while another uses color images from the same scene. In email networks (Bojchevski & Günnemann, 2018), changes in the language used across organizations can lead to different node features.
- *Structure shift* is specific to graph data. It occurs when the topology distribution $\mathcal{D}(A)$ differs across agents (Liu et al., 2023). The factors influencing structural shifts are many, often driven by some hidden latent variable of the environment (Wu et al., 2018; Gui et al., 2022). Structural shifts have three larger categories. *Size shifts* occur when graphs or subgraphs differ in the number of vertices, e.g., sentence graphs vary with text length, and it is the most discussed in the literature (Yehudai et al., 2021). *Density shifts* are present when the edge density, measured by different metrics (node degree, clustering coefficient), differs. This is a structure shift that could prove to be the most critical to GNN generalization capabilities (Bazhenov et al., 2023). *Motif shifts* occur when the base of the graph is common, but only a portion of the graph differs in terms of typical connection patterns (Wu et al., 2022). This case is mostly theoretical, but popular to measure the generalization capabilities of GNNs (Gui et al., 2022).

Concept shift. Concept shift arises when the conditional distribution $\mathcal{D}(y \mid G)$ differs between agents.

- *Feature concept shift* in graph data is analog to *concept shift* seen in Section 2.2.2. It refers to changes in the concept distribution $\mathcal{D}(y \mid X)$, assuming the topology A is fixed. For example, in the case of full graph instances, some superpixel graphs might have a tight correlation between color and class (e.g., blue for water), while others do not (Gui et al., 2022). In the subgraph case, an email network, agents holding emails from respectively older and younger users, may observe different vocabulary-class relations. These shifts are particularly challenging for GNNs since they assume stable feature-label correlations (Duong et al., 2019; Faber et al., 2021).
- *Structure concept shift* for graph refers to differences in $\mathcal{D}(y \mid A)$ (or $\mathcal{D}(A \mid y)$(Liu et al., 2023). This structural concept has been historically modeled using using the *Stochastic Block Model* (Holland et al., 1983; Karrer & Newman, 2011; Abbe, 2018) that assumes $\mathcal{D}_u(A_{uv} \mid \{y_v, y_u\})$ uniform across the graph. More commonly, structure concept shift is connected to the well-studied principle of graph homophily (Zhu et al., 2020; Platonov et al., 2023;; Fuchsgruber et al., 2025), which states that nodes tend to connect to nodes of the same labels. Constant homophily level is critical in centralized training, as GNN tends to perform gradually worse on heterophilous graphs, referred to as *homophily bias* (Tan et al., 2025). In the case of distributed graphs, differences in the graphs homophily level between agents, referred to as *homophily conflict* (Tan et al., 2025), lead to decreased learning accuracy.

3.3.3 System Heterogeneity in Collaborative Learning over Graph Data

Collaborative learning over graphs needs to handle *device*, *link*, and *task* heterogeneity, similarly to the learning over Euclidean data, as discussed in Section 2.2.3. The heterogeneity of the communication and computing resources of the agents, and the quality and stability of their connectivity affect the model

aggregation process during learning. But the new challenge, compared to the case of Eucledian data, is that now the availability of communication and computing resources affects also the exchange of the embeddings for the message passing process, both under learning and under the inference itself.

The task heterogeneity requirement means that the learning and inference processes need to handle several tasks in the GNN (Mai et al., 2024; Meng et al., 2021; Zhang et al., 2023). The modular structure of graph representation learning in general and MPNN in particular supports this heterogeneity inherently, as described in Section 3.1.

3.4 Improving Inference and Learning Effectiveness

As per Euclidean data, methods for improving the effectiveness of learning have been proposed for graph data. They aim at two objectives: designing a tighter convergence to a global solution ψ^*, θ^* of the collaborative training objective (26), or tighter individual convergence to the optimal solution of the agent, ψ^*, θ^*, according to the local objectives (22) (25). Effectiveness is challenged by the generalization gap (seen in Section 3.1), data isolation arising from the various graph data partitions (seen in Section 3.3.1), or by statistical heterogeneity (seen in Section 3.3.2). Subsequently, we classify and describe the main strategies that address effectiveness by borrowing from the classification seen for Euclidean data in Section 2.3.

3.4.1 Data-based Techniques

Data-based techniques manipulate the data used for the learning process.

Data augmentation. The datasets S_k held locally by agents can be augmented to address three challenges: addressing the generalization gap, statistical heterogeneity, and data isolation. Specifically for graph data, we distinguish methods which **augment a graph instance**:

- *L-hop recovering* techniques are applied both for edge and vertex cut, and exchange raw node and edge information between agents, so that they complete the ego-graphs of the nodes. This simplifies the collaborative inference phase of Algorithm 3; however requires to use of cryptographic primitives to preserve privacy (Qiu et al., 2022; Agrawal et al., 2024). Consequently, the L-hop expansion can prove to be expensive and can therefore be restricted to the 1-hop neighborhood (Agrawal et al., 2024).
- *Node addition* techniques aim to extend the local graph by adding new virtual nodes and their associated edges using local graph data only. This can be done using generative models trained on the server-side and then communicated to agents to extend local graphs (Peng et al., 2022). Rewiring strategies using pairwise feature similarity have also been explored, allowing agents to infer plausible edges and labels. These are then integrated into a global graph on the server and redistributed to other agents (Chen et al., 2024).
- *Feature augmentation* techniques for graph data manipulate the feature matrix X. Averaging features of neighbors and synthetic feature generation address homophily bias and feature imbalance (Pei et al., 2023; Lin et al., 2022). Classical methods developed for Euclidean data are used to inject controlled noise (Gao et al., 2024).
- *Label augmentation* has been proposed for various learning tasks. For link prediction, one method uses a local link prediction model to identify plausible links, which are treated as virtual edges. These virtual edges can then be shared across agents to enhance their subgraphs in vertex-cut partitions (Liu et al., 2022). For node classification, label smoothing has been applied to mitigate the effects of homophily bias. In this setting, a single round of distributed averaging (as introduced in Section 2.4.1) is used to propagate label information across agents (Mai et al., 2024). An alternative approach uses contrastive learning, where negative node embeddings—chosen to be close to true node embeddings—are used to sharpen representation learning. These embeddings can be sampled from other agents, such as randomly selected agents in ego-graph datasets when the number of agents is large enough (Giaretta & Girdzijauskas, 2023).

Other methods **augment the dataset with new graph instances**. A graph generative model, trained on the server side, can be used to synthesize new graph instances on the agent side based on local graph data (Guo et al., 2023). The existing graph instances and their node labels can be combined (Wang et al., 2021) to generate new plausible graph and label instances. These techniques can be applied to full graph instances or to the collection of ego graphs and node label (G_k^v, y_v) in the subgraph case (Zhang et al., 2023). Instead of generating new plausible graphs, one can focus on generating non-plausible graphs. This contrastive data augmentation is a frequent approach to provide good generalization capabilities to GNN (Veličković et al., 2019; You et al., 2020). Perturbations include node dropping, edge perturbation, attribute masking, random reindexing of the features, or subgraph sampling (Veličković et al., 2019; You et al., 2020). In collaborative settings, these techniques can be applied to the local graph data for better generalization as in (Chen et al., 2021; 2022; Mai et al., 2024).

Local regularization. Local regularization techniques have been proposed to reduce the discrepancies between agents, either in their model parameters or in the vertex-cut embeddings.

- *Regularizing local models* is ubiquitous for FL as described in Section 2.3.1. For graph data, it can be achieved by modifying the local loss (line 18) of the agents with an additional regularization term. This additional term can be a contrastive term against perturbed graphs (Mai et al., 2024), or a loss on trimmed subgraphs (Ceyani et al., 2025), thus making use of the augmented data techniques described above.
- *Regularizing the vertex cut representations* is specific to the subgraph case with vertex cut, and so far discussed for KG tasks only (Chen et al., 2022). Regularization terms can be added to the local loss (line 18) to decrease the variance between vertex cut representations, respectively on the agent-side (line 15) and server-side (line 23), or to decrease the similarity of the consecutive updates.

Subsampling strategies. In contrast to data augmentation, which uses more data than available locally for the training, subsampling techniques use less data than available locally. This subsampling can be done at multiple levels: agents are sampled in the training round (line 2 of Algorithm 1 and line 3 of Algorithm 3) while nodes are sampled in the local inference (lines 6 and 13 of Algorithm 3). The subsampling can address both learning effectiveness and efficiency. The techniques aimed at improving the effectiveness include:

- *Agent selection* for graph data is inspired by agent selection techniques in traditional FL, seen in Section 2.3.1. It can be used to mitigate the *homophily conflict* through server-side weighted aggregation of the models, where the weight assigned to each agent depends on the homophily level of its local dataset. Learning effectiveness is increased if agents with high-homophily graphs receive higher aggregation weights (Tan et al., 2025). As a form of agent selection for vertex-cut subgraphs, Peng et al. (2021) suggests to average $z_v(t)$ only if it leads to lower local loss.
- *Graph subsampling* has the potential to improve the model accuracy, as shown in (Ceyani et al., 2025). For effective sampling in large graphs, (Bengio et al., 2023; Ceyani et al., 2025) suggests the use of advanced generative models. Training on graphs that are perturbed using subsampling techniques like random node dropping or subgraph sampling is a classical technique that allows GNN to better generalize (You et al., 2020). For collaborative learning, same technique is applied to the local datasets in (Chen et al., 2021; 2022). Node and model parameter sampling is proposed in (Shi et al., 2025) to decrease the computational and communication burden under training. In the subgraph collaborative inference case, node subsampling can be applied to the local expanded graph $\tilde{G}_k$. For instance, (Liu et al., 2021) proposes to select foreign vertex boundary nodes only if their contribution improves the local loss (23). Subsampling strategies are also designed to balance efficiency and effectiveness, using greedy heuristics (Pan et al., 2023), reinforcement learning (Chen et al., 2021), or Markov chain methods (Pan et al., 2023).

Open challenges. The design of data-based techniques to increase learning effecitveness are closely tied to the development of machine learning on graph data itself. However, several research directions are specific to the collaborative learning scenario:

- *Towards more effective graph topology.* Recent research has explored the decoupling of the input graph (i.e., the observed graph) from the graph employed for message passing (Wang et al., 2019; Fatemi et al., 2021; Kazi et al., 2022; Topping et al., 2022; Gutteridge et al., 2023). Such approaches, commonly referred to as *graph rewiring*, provide a novel paradigm to augment the graph structure, such that the neighbor information is used more effectively for increased model accuracy.

- *Towards investigating graph-specific regularization.* Although regularization has been extensively studied in collaborative learning on Euclidean data, it remains underexplored for graph data. Investigation of generalization bounds for regularized graph models (Jegelka, 2022) and adaptation of advanced regularization strategies from centralized GNN research may enhance robustness under data heterogeneity. This aspect has been identified as a priority direction for the theory of graph machine learning (Morris et al., 2024).

- *Towards sampling under partial availability.* Partial availability of agents has recently attracted attention in federated learning research (see Section 2.3.1). Existing studies focus on the impact of unavailability on model parameter evolution. However, for graph data, agent participation affects not only model parameters but also other components, such as the node embeddings passed. Such constraints are further exacerbated in collaborative inference scenarios, where nodes must store representations from other agents.

3.4.2 Model-based Techniques

Model-based solutions with *parameter personalization* and *architecture personalization* are proposed to find optimal models despite data and system heterogeneity in GNNs. Some of these build on the model-based techniques seen for Euclidean data in Section 2.3.2. Others are more specific to GNNs.

Parameter personalization. Parameter personalization techniques have been proposed in the form of:

- *Hypernetwork-based* Specifically for graph data, (Liang et al., 2024) extend the hypernetwork personalized framework of (Shamsian et al., 2021) by defining agent k representation vector as the mean of the node representations within the local subgraph: $\frac{1}{N_k}\sum_{v\in\mathcal{V}_k} z_v(t)$.

- *Local adaptation* The classical local adaptation strategies, previously reviewed in Section 2.3.2, can be extended to distributed GNNs. In (Zhang et al., 2023), the local model is formulated as a learnable interpolation between the global and the local GNN parameters, while (Baek et al., 2023; Fang et al., 2025; Han et al., 2024) suggest masking the weights of the global models using learnable parameters. A meta-learning framework is suggested for node-level tasks on graph data in (Wang et al., 2022; Han et al., 2024), where the support and query sets are defined as subgraph partitions.

- *Cluster-based* Classical techniques can be used to cluster agents based on gradient similarities (Xie et al., 2021). If the agents observe a single ego-graph only, then they can be clustered based on their node embeddings. As suggested in (Qu et al., 2023), an effective way to support learning within clusters is to construct virtual local graphs.

- *Similarities-based* The similarity-based techniques seen in Section 2.3.2 have been adapted specifically to graph data. The similarity matrix $\Omega = [\omega_{kl}]_{k,l\in\mathcal{K}^2}$ can be interpreted as defining a latent *inter-agent graph*, distinct from the local graph data of each agent. Similarity scores, inspired by (Zec et al., 2022), are evaluated for different GNN models on randomly generated subgraphs in (Baek et al., 2023). (Liang et al., 2024) use subgraph representations c_k to compute similarities between agents, which are then diffused over the inter-agent graph using learnable *sheaf diffusion* (Bodnar et al., 2022) algorithms trained on the server. A further contribution is the proposal of a novel metric based on the optimal transport distance between model parameters (Shi et al., 2025).

Architecture personalization. Architecture personalization modifies the model architecture and the parameters returned at the end of the training round:

- *Model decoupling* As described in Section 3.1, a GNN $F_{\psi,\theta}$ can be decomposed into an encoder ϕ_θ decomposed into sublayers $f_{Enc}, f_{Mes}^{(l)}$, $f_{Agg}^{(l)}$, $f_{Up}^{(l)}$ for l=1,...,L, and a predictor f_ψ. This layered architecture allows for splitting the model *layer-wise* into different components. The split means that one component is shared while another is personalized to the agents (Mai et al., 2024). More precisely, the shared component can either be trained on the agent side and is then aggregated periodically on the server, or it is trained on the server (Meng et al., 2021; Zhang et al., 2023) using alternative optimization techniques (Meng et al., 2021). Decoupling can occur between the encoder ϕ_θ and predictor f_ψ (Mai et al., 2024), or between the feature encoder f_{Enc} and subsequent layers, where either the early layers are personalized and later ones shared (Meng et al., 2021), or vice versa (Zhang et al., 2023).

  Local GNNs have also been pruned *width-wise* in (Shi et al., 2025), where each agent removes local learnable weights based on a layer-specific, learnable threshold. To mitigate severe performance degradation during aggressive pruning, a gradual rollback strategy is employed. Aggregating width-wise pruned models remains challenging and has been proposed to be performed at lower bit resolution to promote generalization across agents.

- *Multimodal personalization* The server can maintain M specialized GNNs parameters $\{\theta_m^c\}_{m=1}^M$ for each graph or subgraph type. The framework for multimodal personalization (Chen & Zhang, 2022; Smith et al., 2017) has been applied to full graph instances associated when one agent that can collect graphs of different modalities or tasks (He et al., 2022). In vertically partitioned subgraph settings, predictions from agents' local models for the same nodes can be combined using a confidence-based aggregation scheme known as *knowledge voting* (Mai et al., 2024).

- *Knowledge distillation* techniques allow us to build more effective personalized models for local inference. They consist of distilling the knowledge of the private topology observed by each agent into personalized models. The main benefit is that agents can locally build models that do not require collaborative inference, yet these models implicitly incorporate information about other agents' topologies. This case is particularly useful in the case of vertical topology-free partition, where these techniques can be applied to the feature encoder f_{Enc} to mimic the GNN performance (Fu et al., 2024). In practice, an alignment term is added to the local loss to align the vertex-cut embeddings generated by the different models on their respective private data.

- *Decoupling the message passing operations.* To address data isolation, the message computations in (17) are adapted based on the type of edge (Lei et al., 2023). In particular, messages along cut edges $\mathcal{E}_{\text{cut}}$ may carry only partial information, for instance, when connected to stale embeddings. To compensate for this, these messages are processed using more expressive message and aggregation functions, f_{Mes} and f_{Agg}, designed to better capture missing dependencies. In contrast, messages over local edges (i.e., those within G_k) rely on simpler functions¨(Wu et al., 2019), as they benefit from more complete and up-to-date local information. This separation effectively decouples the message passing process according to edge type.

Open challenges. Given the current state-of-the-art of collaborative learning and the recent development of ML on graphs, open challenges include:

- *Towards more advanced MPNNs.* More expressive extensions of the standard MPNN framework have been introduced under the umbrella of *Topological Neural Networks*(Hajij et al., 2022; Papamarkou et al., 2024). These architectures preserve the core message-passing mechanism while generalizing it to broader topological structures beyond graphs. Understanding how these structures integrate into current collaborative settings constitutes an open research direction. Separately, pruning and layer selection strategies, extensively studied in CNN-based federated learning(Park & Joe-Wong, 2024), can be adapted to GNNs to enhance scalability and efficiency, as demonstrated in (Kummer et al., 2025); however, these techniques remain underexplored in collaborative GNN settings.

- *Towards more personalized models.* Promising directions include the exploration of GNN-specific metrics in collaborative settings. For example, the moving tree distances (Chuang & Jegelka, 2022)

have been tied to generalization capabilities both empirically and theoretically (Maskey et al., 2025; Southern et al., 2025). We also envision multimodal embeddings arising when the schema graph is split across agents, with each agent selecting from specialized local models. Furthermore, knowledge distillation techniques, although well studied in the GNN literature (Tian et al., 2025), remain underexplored in collaborative settings.

- *Towards the utilization of foundation models.* A significant paradigm revolutionizing ML is the use of *foundation models (FMs)*, that is, generalist pre-trained models that can be effectively adapted to downstream tasks with a relatively small amount of new data. Successful FMs abound in the textual and visual domains, but the development of FMs on graphs remains an open problem (Wang et al., 2025), and its potential integration in collaborative settings is not yet evaluated.

3.4.3 Training and Inference under Vertical Partitions

The different vertical partitioning cases introduced in Section 3.3.1 present requirements on the learning process that are not captured by the standard Algorithm 1, which is suitable for horizontal partitions of full graph instances, or by Algorithm 3, which is designed for horizontal subgraph partitioning. Specifically, on top of the message aggregation (18) already integrated in Algorithm 1, the vertical partitions require aggregations of node or graph representations between the agents to improve the effectiveness of the learning.

Aggregating node representations. In the case of vertex-cut partitions, node representations for the same node are generated at different agents, and therefore need to be combined. Indeed, the server receives (line 17) a collection of node representation $\{z_{v,k}(t-\frac{1}{2})\}_{v\in\mathcal{V}_k}$ from each agent $k \in \mathcal{K}$. So for a node cut $v \in \mathcal{V}_{cut}$, the different representations $\{z_{v,k}(t-\frac{1}{2})\}_{k\in\mathcal{K}}$ have to be aggregated (line 23). Different choices exist for this aggregation step, which aim to increase the quality of the final node representation $\{z_v(T)\}_{v\in\mathcal{V}}$:

- *Averaging node representations.* The standard approach is to compute $\bar{z}_v(t)$, the average on $\mathcal{K}$ of the representations $\{z_{v,k}(t-\frac{1}{2}), k \in \mathcal{K}\}$ from the different agents to obtain more generalizable representations (Chen et al., 2021; 2022; Wu et al., 2022; 2021; Zhang et al., 2022; Qu et al., 2023; Han et al., 2024; Agrawal et al., 2024; Ni et al., 2021). The averaging can be restricted to agents with low local loss (Mai et al., 2024). Alternatively, the embeddings $z_v(t)$ can be updated by taking a small step toward $\bar{z}_v(t)$ using a gradient-based formulation, which provides more stable updates throughout the training (Wu et al., 2022; 2021; Zheng et al., 2023).

- *Personalizing node representations.* Personalization techniques return a different node representation $z_{v,k}(T)$ for each agent at the end of the training instead of a common one $z_v(T)$. The personalization techniques used for the model parameters $\theta_k(t-\frac{1}{2})$ in Section 2.3.2 can be applied to the node representations $z_{v,k}(t-\frac{1}{2})$ to address statistical shifts for graph data and task heterogeneity. $z_{v,k}(t)$ is constructed by combining $z_{v,k}(t-\frac{1}{2})$ and $\bar{z}_v(t)$ using MLP (Chen et al., 2021) or a learnable linear interpolation (Wang et al., 2024).

Transfer learning. Rather than aggregating the node representations $\{z_{v,k}(t-\frac{1}{2})\}_{k\in\mathcal{K}}$, their difference between the agents can be used to design *transfer learning* for graph in a federated setting. Typically, we assume that a model has been trained by "source" agents s on a sufficient amount of vertically partitioned graph data, and we want to transfer the knowledge to a target agent t. Aligning these representations while t trains its local model allows to adapt the frozen model of s to t's domain. This alignment can be achieved by training against an adversarial network (Shen et al., 2020), trained to distinguish between source and target domains (Guan et al., 2021), or using a knowledge distillation loss (Hinton et al., 2014; Mai et al., 2024).

Aggregating graph representations. Lastly, in the case of vertical partition of full graph instances, multiple graphs held by different agents can characterize the same item or person. To increase the generalization capabilities of graph-level tasks, (Dong et al., 2021) suggests averaging the graph embeddings to increase prediction accuracy, considering the specific use case of health applications.

Open challenges. Vertical partitions raise challenges that go beyond the standard assumptions of federated learning for graphs, and future work can address them along several dimensions:

- *Towards more personalized aligned representations.* Aligned node or graph representations can be adapted to local distributions (Section 3.3.2). For example, the customer behavior on two marketplaces might be different depending on the platform.

- *Towards more advanced transfer techniques* Transfer techniques have the benefits of decreasing training costs, allowing fast deployment, and in the case of vertical graph partitioning also decreasing communication needs. While there are few results on transfer learning in collaborative scenarios, further advanced solutions could be designed, for example by adapting centralized solutions (Frasca et al., 2024; Wang et al., 2025) to collaborative settings.

3.5 Improving Inference and Learning Efficiency

While Section 2.4 reviewed general design principles to improve efficiency in collaborative learning under realistic resource constraints, we now focus on how these principles translate to graph data specifically. In particular, we discuss techniques inherited from Section 2.4, as well as methods specific to graph-structured data. Due to message passing, efficiency now needs to be considered both for training and for inference.

3.5.1 Aggregation over a network

The centralized, decentralized and hierarchical aggregation patterns, introduced in Section 2.4.1, are all applicable also for learning with graph data. They can be employed to aggregate the local models $\theta_k(t)$ during *collaborative training* and to aggregate node embeddings in the message passing equation (18), during *collaborative inference*.

Under collaborative learning with Eucledian data, the communication graph W for the model aggregation is determined mainly by the physical network topology, that is, by the possible communication links or paths among the agents. This is true as well for collaborative learning over multiple graph instances. However, in the subgraph or ego-graph partition cases the computational graph $\mathcal{G}$ (or its matrix representation A) has to be taken into account as well.

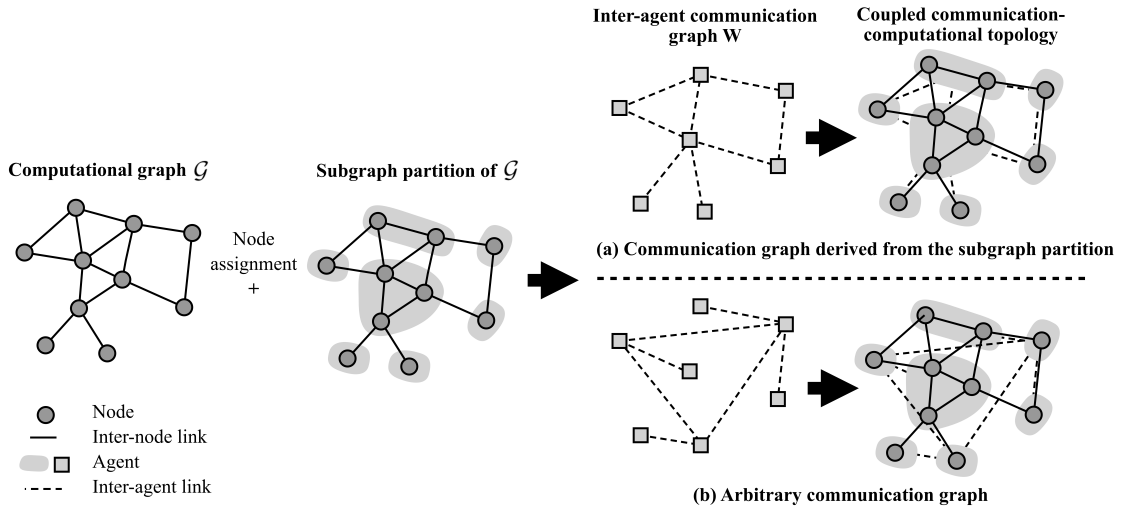

Figure 19: Illustration of the coupling between computation and communication topologies in the subgraph case. The computational topology, represented by graph $\mathcal{G}$, and the communication topology, represented by matrix W, can be: (a) *derived from* the partitioning of $\mathcal{G}$ between agents, such as an edge-cut partition (see Section 3.3.1), or (b) *independent from* $\mathcal{G}$, representing an arbitrary communication network.

The coupling of the computation graph $\mathcal{G}$ and the communication graph given by W typically follows one of two paradigms, as illustrated in Figure 19.

- *Communication graph derived from subgraph partitioning of $\mathcal{G}$.* In this case, the communication topology is consistent with the computational one, meaning that certain edges in $\mathcal{G}$ also represent communication links in W. For example, in a swarm of robots (Blumenkamp et al., 2022; Nazzal et al., 2024), $\mathcal{G}$ connects physically adjacent robots, thereby defining both computation and communication topologies. In ego-graph settings, the communication graph is identical to the computation graph. For edge-cut subgraph partitions, communication links correspond to edges cut between agents, as in IoT systems where agents monitor different subsets of nodes (Pan et al., 2023; Nazzal et al., 2024; Zeng et al., 2022). The node-to-agent assignments in $\mathcal{G}$ may be given as problem constraints, for example, determined by agent locations (Nazzal et al., 2024; Zeng et al., 2022), or optimized given the problem constraints (Pan et al., 2023; Zeng et al., 2022).

- *Communication and computational graph decoupled.* Here, computation and communication are decoupled. Typically, when peer-to-peer communications among all agents are allowed (Giaretta & Girdzijauskas, 2023; Qu et al., 2023) (W forming a fully connected graph), the communications graph can be further refined adaptively to accommodate learning efficiency (Wang et al., 2025).

Typically, agents aggregate their model parameters and embeddings through the same topology. However, distinct communication topologies can be employed for model and embedding aggregation, allowing the topology to differ between training and inference to leverage both paradigms (Lei et al., 2023; Giaretta & Girdzijauskas, 2023; Guo et al., 2024; Wang et al., 2025) Hierarchical architectures have also been proposed, where centralized or decentralized aggregation patterns may occur at different levels of the hierarchy (Guo et al., 2024).

Embedding aggregation. In the case of subgraph partitions, embeddings $h_v^{(l)}$ generated by different agents must be aggregated either to perform message passing (line 13 of Algorithm 3) or for vertex-cut processing (see Section 3.4.3). Embedding aggregation techniques aim to adapt to constrained communication topologies while reducing both the communication cost of transmitting embeddings (line 12) and the computational cost of aggregating them (line 13) during collaborative inference in Algorithm 3. solutions are proposed for centralized and for decentralized aggregation topologies:

- *Centralized communications* are the default approach inherited from FL algorithms and Algorithm 3. The server collects and redistributes embeddings and is sometimes responsible for maintaining them (Liu et al., 2022; Wu et al., 2021; 2022; Yan et al., 2024).

- *Decentralized communications,* leading to truly decentrized inference, means that embeddings are exchanged directly between the agents. Decentralized inference is the default for edge-cut ego-graph partitions, where the communication pattern naturally aligns with the computation pattern (Gao et al., 2021; Lee et al., 2021; Blumenkamp et al., 2022; Wang et al., 2022; Giaretta & Girdzijauskas, 2023; Krasanakis et al., 2022; Pan & Zhu, 2022; Lee et al., 2023; Solodova et al., 2025; Pan et al., 2023; Gao & Gündüz, 2025; 2023; Gao et al., 2023; Guo et al., 2024; Nazzal et al., 2024; Meng et al., 2025; Wang et al., 2025). However, decentralized inference also appears in subgraph edge-cut partitions, for example, when agents represent institutions or edge servers monitoring a subset of nodes (Lei et al., 2023; Pan et al., 2023; Nazzal et al., 2024; Zeng et al., 2022; Wang et al., 2023). In such cases, the communication graph can be inferred from the cut of $\mathcal{G}$ (Zeng et al., 2022) or from the communication coverage area of the servers (Nazzal et al., 2024), as shown in Figure 19.

The embedding aggregation techniques for increased efficiency modify the time of the aggregation, the topology, or the content of the exchanged data, as summarized in Figure 20.

The first line of work focuses on modifying the *timing* of embedding aggregation to reduce the communication cost using:

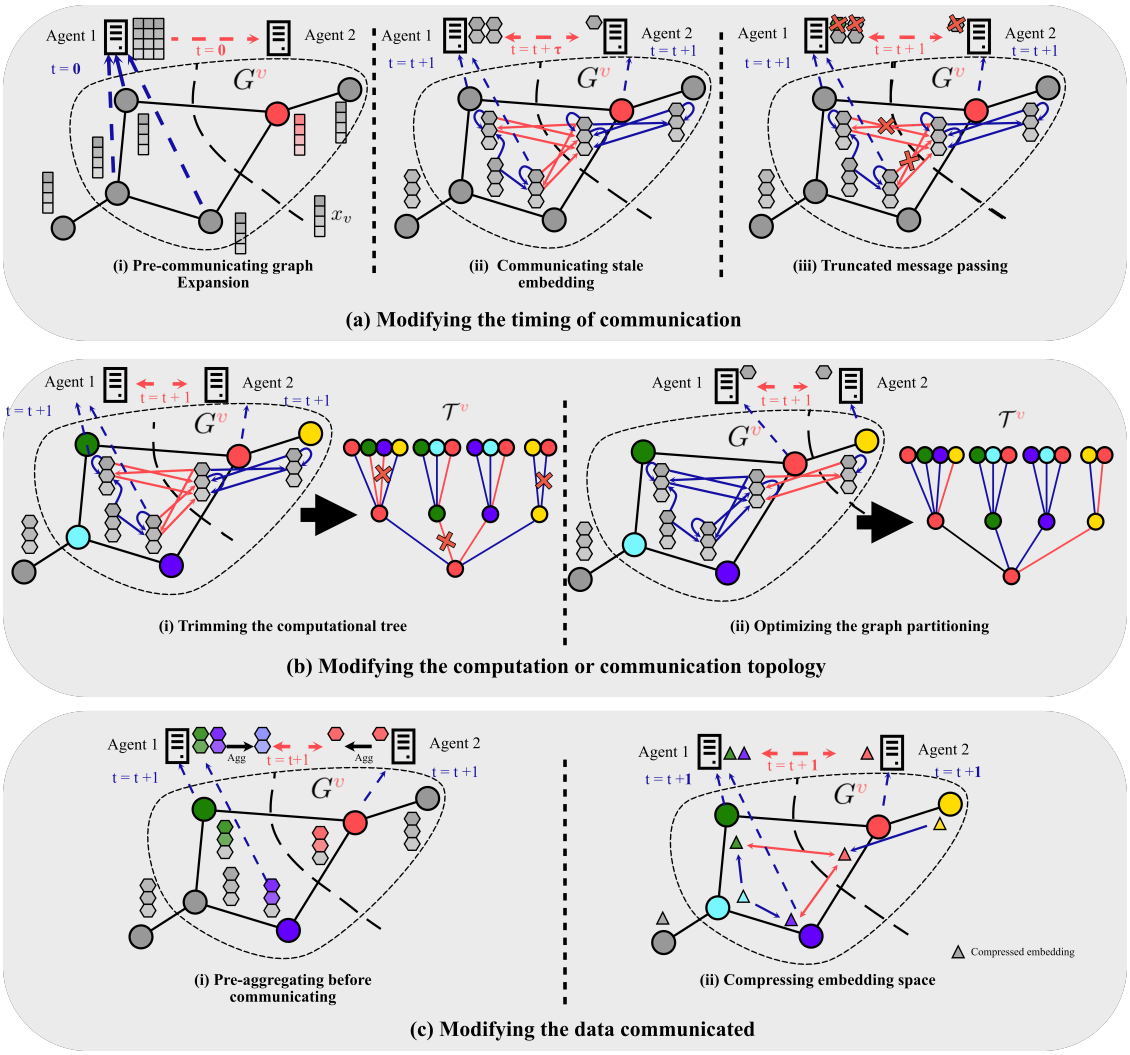

Figure 20: Embedding aggregation techniques focus on modifying: (a) the timing of communications, (b) the computation or communication topology, and (c) the data communicated

- *Pre-communication.* Communicating sufficient information before the start of the training can avoid the need for a collaborative inference. Typically, the expanded subgraph data $\tilde{G}_k$ is communicated to relevant agents or the server prior to the first training round, so that agents perform isolated inference on $\tilde{G}_k$, proposed for centralized aggregation (Zhang et al., 2022; Chen et al., 2021; 2022) or through a decentralized protocol (Agrawal et al., 2024; Qiu et al., 2022). This uphill communication cost can prove worth it if the algorithm takes many training rounds to converge in the collaborative inference, ultimately outweighing it.

- *Stale embeddings.* Letting node embeddings $h_v^{(l)}(t)$, requested normally in every round in the default Algorithm 3 (line 12), become stale reduces the communication cost of collaborative inference. Therefore, $h_v(t)$ is proposed to be requested periodically in (Du & Wu, 2022). To avoid affecting learning effectiveness due to this staleness, (Du & Wu, 2022) derives an optimal sampling rate w.r.t. the tradeoff between a communication budget and the convergence speed of training.

- *Truncated message passing.* Alternatively, to reduce the amount of information to be communicated in each round, (Liu et al., 2021) suggests communicating only a part of the embeddings, specifically only the later layers l, eliminating layers that have less contribution to the learning process.

The literature focuses on modifying the computation and communication topology to reduce communication overhead. During inference, the representation z_v for node v is obtained by propagating and aggregating messages along the edges of the input graph, following the procedures in (17), (18), and (19). This sequence of computations induces a *computational tree* $\mathcal{T}^v$ rooted in v, directly derived from its ego-graph G^v. In the case of subgraph partitioning, edge cuts in the input graph translate into inter-agent communication links in the computational tree. Reducing communication during collaborative inference thus often amounts to modifying $\mathcal{T}_v$ by:

- *Trimming the computational tree* Trimming the computational tree $\mathcal{T}^v$ can significantly reduce the communication needed. Tree-trimming strategies proposed in the literature (Liu et al., 2021; Chen et al., 2021; Pan & Zhu, 2022; Pan et al., 2023) remove selected branches of $\mathcal{T}_v$ leading to node v, thereby reducing both computation and communication, and notably the communication when the removed branches involve inter-agent communication. Such strategies can be determined locally by each agent (Liu et al., 2021), applied to its local expanded topology $\tilde{\mathcal{G}}_k$, assumed to be known; or centrally by a server with knowledge of the global graph $\mathcal{G}$ (Chen et al., 2021; Pan & Zhu, 2022; Pan et al., 2023). The communication workload can be balanced, either at the global level, e.g., minimizing the maximum workload across agents (Pan et al., 2023), or at the individual level by balancing each agent's workload (Liu et al., 2021; Chen et al., 2022). While trimming $\mathcal{T}^v$ reduces the computational and communication load, it also discards part of the knowledge contained in the subgraph G^v, potentially degrading model performance. The impact of the computational tree depth on prediction accuracy has been demonstrated in real-world applications such as robot swarm localization (Zhou et al., 2021; Blumenkamp et al., 2022). The trade-off between efficiency and accuracy can be balanced by various techniques, including selecting nodes that contribute only marginally to loss reduction (Liu et al., 2021) or using reinforcement learning to guide branch removal (Chen et al., 2021). In homophilous graphs, shortcut jumps within $\mathcal{T}^v$ can be used to reach a target node v in fewer communication steps (Pan & Zhu, 2022).

- *Graph partitioning.* Another approach is to optimize graph partitioning to minimize communication cost shown by $\mathcal{T}^v$. Partition refinement can be achieved by reassigning nodes dynamically to agents to reduce the number of edge cuts (Nazzal et al., 2024). Node assignment can also account for other system-level constraints such as data acquisition cost or the number of agents (Zeng et al., 2022).

Modifying the nature of the data exchanged can significantly reduce communication costs. This can be achieved by employing simpler architectures than the generic MPNN described in Section 3.4.3, achieved with techniques such as:

- *Pre-aggregating before communicating* Embeddings can be pre-aggregated before being communicated, thus reducing the amount of communication. When the expanded subgraph $\tilde{G}_k$ is transmitted for isolated inference, communicating the full L-hop neighborhood may still involve a prohibitive amount of node data. A practical solution is to leverage the properties of simple GNNs such as GCN (Kipf & Welling, 2017) or SGC (Wu et al., 2019), which use non-learnable f_{Mes} and f_{Agg}, to pre-aggregate the required embeddings locally before transmission. For example, in a GCN, the data of an l-hop neighborhood of a private boundary node $v \in \mathcal{V}_{p,k}$ can be pre-aggregated as $\sum_{u \in \mathcal{V}_k | (u,v) \in \mathcal{E}_k} x_u$ instead of transmitting the set $\{x_u \mid u \in \mathcal{V}_k, (u,v) \in \mathcal{E}_k\}$.(Yao et al., 2023) This reduces the cost of communicating an l-hop neighborhood to that of an $(l-1)$-hop neighborhood for GCN, and to that of the node v itself for SGC. For SGC, the local inference cost can be further reduced by pre-computing the non-learnable operations (Lei et al., 2023).
- *Squashing the embedding space* The dimensionality of node embeddings d_x directly impacts communication cost. In node-level prediction tasks, an agent can locally generate pseudo-labels $\hat{y}_v$ by applying a feature encoder f_{Enc} followed by a local predictor f_ψ. Instead of transmitting full embeddings, only the scalar pseudo-labels $\hat{y}_v$ are sent, substantially reducing the communication load. Although these pseudo-labels do not initially incorporate graph topology, they can be refined via non-learnable diffusion algorithms (see Section 2.4.1), which propagate labels across the graph in a topology-aware manner. Under the homophily assumption, this allows even unlabeled nodes not involved in the collaborative inference to receive meaningful predictions y_v (Huang et al., 2021). This approach has been implemented in decentralized settings for node-label prediction, achieving minimal communication cost while maintaining competitive performance (Krasanakis et al., 2022).

Model aggregation. For model aggregation, centralized, decentralized and hierarchical patterns are proposed, with the choice largely determined by the type of the graph partition:

- *Centralized model aggregation*, inherited from Algorithms 1 and 3, is the dominant approach in cross-silo scenarios where network constraints are moderate. It is commonly used for Multiple Graph Instance settings (He et al., 2022), Vertical Subgraph settings (Chen et al., 2021), and, less frequently, for cross-silo subgraph cases (Yao et al., 2023). As in Euclidean FL, periodic aggregation trades off convergence speed and communication cost. For Euclidean data, prior work typically assumed generic data distributions and derived convergence bounds from the local objectives (2) (Stich, 2019; Yu et al., 2019). In the graph case, structural assumptions such as graph homophily allow a refined convergence characterization (Yao et al., 2023) as a function of the input graphs $\{G_k\}_{k \in \mathcal{K}}$. In particular, (Yao et al., 2023) uses a stochastic random graph model (Keriven et al., 2020) with fixed homophily for node-level tasks to analyze convergence in the subgraph edge-cut configuration using a GCN $\mathcal{F}$. Their analysis covers both i.i.d. and label-shift scenarios, showing that convergence depends on the structural shift between each G_k and the global graph G. Higher graph homophily, lower label shift, and smaller structural shift, achieved via higher-hop graph expansion, improve convergence. As the number of agents K increases, distributional shifts intensify, requiring proportionally larger hop expansions to sustain performance.
- *Decentralized synchronous model aggregation* is typically used in ego-graph scenarios (Olshevskyi et al., 2025) and has been extensively analyzed for subgraph edge-cut cases (Scardapane et al., 2020). The mixing matrix $W = [w_{kl}]_{(k,l) \in \mathcal{K}^2}$ corresponds to the adjacency matrix of the communication graph derived from the subgraph edge-cut. Convergence proofs rely on standard FL assumptions of loss similarity and training stability (Stich, 2019; Yu et al., 2019). Stability requires W to be doubly stochastic and connected. Dual-based distributed aggregation inspired by ADMM (Boyd et al., 2011) has also been proposed (Scardapane et al., 2020). Consistent with Euclidean FL, frequent aggregation accelerates convergence and improves resilience to link loss. Additional robustness to link failures is achieved in (Gao et al., 2022) by centrally training a GNN to perform distributed averaging that tolerates both link loss and partial participation. Other works propose to use a different mixing matrix at each round, subsampled from the original communication graph to limit communications. This subsampling optimizes the mixing matrix at every round using an actor-critic model, trained to minimize convergence time. During training, backpropagating the gradient

(19 of Algorithm 3) requires communication between agents for collaborative inference and local gradient aggregation. This communication challenge can be mitigated (Olshevskyi et al., 2025) through batching and message piggybacking: gradients are backpropagated along the graph during each sample's backward pass with small messages, while expensive parameter consensus occurs only once per mini-batch, amortizing the communication cost across multiple samples.

- *Decentralized asynchronous model aggregation* uses gossip-based protocols (Zheng et al., 2023; Giaretta & Girdzijauskas, 2023), though the derivation of the formal convergence rates remains an open problem. Giaretta & Girdzijauskas (2023) propose asynchronously exchanging parameters at the layer level—aggregating different layers at different times—to enhance robustness to asynchronicity.
- *Hierarchical aggregation*, discussed (Guo et al., 2024) in improves the scalability of learning by combining local peer-to-peer aggregation with occasional centralized coordination, while collaborative inference is proposed to remanin peer-to-peer, to avoid central bottlenecks.

Beyond topology choice, communication cost can be reduced by adjusting the timing of model aggregation:

- *Pre-training* performs part of the training centrally before collaborative learning, avoiding communication during early training stages (Krasanakis et al., 2022; Gao et al., 2022). A hybrid approach initializes parameters from a centrally pre-trained model (Chen et al., 2022) before fine-tuning collaboratively with Algorithm 3.
- *Local SGD* runs several local epochs between aggregation rounds. As in standard FL (Section 2.4.1), this trades reduced communication for slower model synchronization. This trade-off has also been studied in the subgraph setting (Yao et al., 2023).

Open challenges. Distributed aggregation has received significant attention in the literature, particularly because it is closely tied to the subgraph setting and to seminal diffusion algorithms on graphs. Yet, more results are needed to understand the role of the physical network infrastructure, to ensure efficient inference and learning.

- *Towards decoupling input and computational graphs.* The benefits of *graph rewiring*, that is, separating input graphs from computational graphs to improve scalability and robustness, are well-established in centralized settings (Wang et al., 2019; Fatemi et al., 2021; Kazi et al., 2022; Topping et al., 2022; Gutteridge et al., 2023). However, this line of research is largely unexplored in distributed environments, where the resource efficiency of learning could be improved with the alignment of input, computational, and physical communication graphs.
- *Towards highly decentralized cases.* The viability of collaborative inference in highly decentralized cases, in typically asynchronous and resource-constrained environments, remains uncertain, with open questions about speed, stability, and overall utility compared to centralized approaches (Scardapane et al., 2020; Blumenkamp et al., 2022). Separate from resolving these performance questions, practical progress will require hardware alignment (Hooker, 2021), as current deployments still predominantly favor centralized GPU clusters, over more distributed architectures like edge computing.
- *Towards topology information as a complement.* While end-to-end training of deep MPNNs proves unnecessary for simple graph relations, particularly in homophilous graphs (Krasanakis et al., 2022), fundamental questions remain about when this simplification is justified. The gain of using large MPNN models (mapping (A, X) to Z) compared to models that ignore edge information (mapping X to Z) is poorly understood, a limitation exposed by current graph benchmarks (Bechler-Speicher et al., 2025). This knowledge gap suggests an underexplored alternative where agents infer local representations on private data using Euclidean-only models, then refine results via lightweight graph-based post-processing (Zhu & Ghahramani, 2002), substantially reducing communication costs. Such modular frameworks open up for the integration between FMs on Euclidean data (ex, images, text) and graph-based solutions.

3.5.2 Asynchronous Aggregation and Delay Control

The synchronous standard Algorithms 1 and 3 assume synchronicity of agents for model aggregation and message aggregation. However, in practice, communication delays can differ from agent to agent and can also vary in time. This leads to stale model parameters and embeddings, which challenges the efficiency of the learning. Therefore solutions that address the communication delays have been explored. Conversely, asynchronicity in the form of introduced delay has also been exploited to improve the effectiveness of models. The solutions proposed for asynchronous model aggregation and inference are:

Model aggregation under asynchronicity. In collaborative training of GNNs, as in Algorithms 1 and 3, the model aggregation is similar to traditional FL. Consequently, techniques and convergence results from Section 2.4.2 apply to GNN. With subgraph partitioning, the communication graph often matches the computation graph, making the system topology-dependent. This ties aggregation issues to asynchronous aggregation over graphs, as studied in (Lian et al., 2018).

Message passing under asynchronicity. Standard GNN inference assumes synchronous message passing between nodes (17),(18), and (19). Considering message passing under asynchronicity enables two complementary approaches: mitigate the effect of the delays, or conversely, exploit delays to improve model performance.

This staleness is typically reflected in the index t of the node embeddings $h_v^{(l)}(t)$, arising due to imperfect synchronization steps (lines 12 and 13 of Algorithm 3) during collaborative inference.

Some works suggest to introduce topological delays or acceleration, that changes the hop count index l in the message aggregation (18), either at the level of the aggregated messages $m_{v,u}^{(l)}$ or the neighborhood $\mathcal{N}^{(l)}(v)$ (Gutteridge et al., 2023; Chen et al., 2024). The main solutions focus on:

- *Modifying the collaborative inference protocol* to maintain model performance despite embedding staleness. GNN training with stale embeddings is studied in (Fey et al., 2021; Peng et al., 2022). To reduce communication costs during training, these works propose to skip certain embedding transmissions and derive convergence bounds that account for the resulting staleness. Message passing under asynchronicity is also addressed in (Yu et al., 2024), where the communication and processing burden of stale embeddings is reduced by sequentially performing updates along predefined paths in the graph, referred to as update chains.

- *Modifying the MPNN to mitigate staleness.* Specifically, (17), (18), and (19) are modified to improve robustness to embedding staleness during collaborative inference. For example, Energy GNNs (Gu et al., 2020; Di Giovanni et al., 2022), inspired by diffusion algorithms on networks (see Section 2.4.1), allow for asymptotic inference that is provably robust to stale embeddings (Solodova et al., 2025). Alternatively, asynchronous MPNNs (Faber & Wattenhofer, 2024) support scenarios in which a feature update at an origin node triggers asynchronous updates throughout the graph. These models propose integrating the origin node's state into the messages (17) to mitigate the effects of stale embeddings (Faber & Wattenhofer, 2024). This also allows more expressive variants of MPNNs by leveraging unique node identifiers. Another asynchronous MPNN class is proposed in (Mathys et al., 2024), where node updates are propagated outward from the origin up to distance L (flood phase) and then back (echo phase), a process found to be more effective than standard MPNNs.

- *Modifying the MPNN to improve expressiveness.* Variants of the MPNN have also been proposed to incorporate topological delays for enhanced expressiveness. These delays can be controlled by learnable node actions (e.g., listening, broadcasting, or remaining inactive) (Finkelshtein et al., 2024). Embedding updates may be delayed or accelerated based on topological considerations (Gutteridge et al., 2023; Chen et al., 2024), typically by modifying the aggregation step (18). Delays arise when l-hop embeddings from 1-hop neighbors are used, whereas accelerations occur when embeddings from nodes beyond 1-hop are included. These adjustments help address oversmoothing (Chen et al., 2020) and oversquashing (Topping et al., 2022), two fundamental limitations of GNNs.

Open challenges. As we have seen, asynchronicity is a key constraint of collaborative systems. It is considered by some novel MPNN architectures, but several challenges need to be addressed in realistic collaborative systems, including:

- *Towards temporal/time-aware MPNNs in collaborative asynchronous settings.* Some of the approaches reviewed above (Barros et al., 2021; Cini et al., 2025; Mathys et al., 2024; Gutteridge et al., 2023; Finkelshtein et al., 2024; Chen et al., 2024; Faber & Wattenhofer, 2024) are designed for centralized settings and do not account for asynchronicity arising in collaborative environments, due to heterogeneous agent capabilities, variable communication latency, and computational resource constraints. Integrating temporal-dependent MPNNs into collaborative frameworks under such asynchronous conditions, particularly understanding how agent-induced delays interact with temporal mechanisms, affect convergence guarantees, and impact idle time during distributed inference, remains largely unexplored.

- *Towards capability-aware MPNNs.* Current collaborative GNN approaches typically assume homogeneous agents with static capabilities. However, practical deployments involve agents whose computational power, memory constraints, and communication bandwidth fluctuate dynamically due to varying workloads, network conditions, and resource availability. While prior work has integrated capability considerations for balancing load or cost constraints (Liu et al., 2021; Chen et al., 2022; Zeng et al., 2022), these approaches assume capabilities remain static throughout execution. How computation graphs should adapt to runtime capability changes, by dynamically adjusting message passing depth, embedding dimensions, or aggregation responsibilities while agents operate at varying paces, is an exciting direction for future research.

- *Towards asynchronous aggregation in extreme environments.* Collaborative GNN deployment in extreme communication environments, such as satellite networks, underwater sensor systems (e.g., autonomous underwater vehicles), or highly mobile vehicular networks, introduces challenges that differ fundamentally from standard distributed settings. These environments exhibit prolonged and highly variable delays (orders of magnitude longer than typical network latency), intermittent connectivity with unpredictable outages, and rapidly changing network topologies. Understanding how asynchronous aggregation protocols can maintain model effectiveness under such extreme conditions remains a critical open problem for allowing collaborative graph learning in these emerging domains.

3.5.3 Communication Efficiency

For collaborative GNNs with topology partitioning, embedding vectors must be communicated among agents during the message-passing phase. These exchanges can occur in centralized, peer-to-peer, or mesh architectures and are subject to heterogeneous, unreliable, and costly communication links. Similar to model aggregation in traditional FL, the efficiency and reliability of embedding aggregation depend critically on the state of the communication links and the cost of transmission.

However, communication for message passing differs from conventional data exchange: embeddings are intermediate representations, and the downstream learning task typically depends on aggregated functions of these embeddings rather than their exact values. This observation motivates the design of communication schemes tailored to collaborative GNNs, optimizing both transmission efficiency and robustness to channel impairments.

Transmission efficiency. Transmission efficiency quantifies the utility of the transmitted embeddings relative to the consumed communication resources, such as bandwidth and transmission power. In collaborative GNNs, this is particularly critical when each agent holds only a small portion of the graph, potentially a single node, as in fully decentralized settings. In (Lee et al., 2021), a coding and retransmission scheme is proposed to improve the reliability of embedding aggregation, aiming to maximize inference accuracy under constrained resources. The work in (Wang et al., 2025) proposes a node subsampling strategy for selecting communication edges that accommodates link bandwidth constraints, to minimize convergence time during training. Theoretical analysis for the same scenario is provided in (Gao et al., 2021), where disruptions

in embedding communication are modeled by a probability of failed message passing due to wireless impairments or temporary disconnections. The results show that accurate inference remains possible if such stochastic failures are accounted for in the learning process.

Wireless communication for embedding aggregation. The principles of over-the-air computation, originally proposed for FL model aggregation, have been extended to embedding aggregation in collaborative GNNs (Gao & Gündüz, 2023; Lee et al., 2023; Gao & Gündüz, 2025). In this setting, the superposition property of the wireless channel is exploited to aggregate embeddings directly during transmission, reducing communication costs and allowing simultaneous updates from multiple agents. These methods have been shown to maintain robustness against channel impairments and, in some cases, operate without channel state information. Extensions to more complex and dynamic topological structures are explored in (Fiorellino et al., 2024). Practical demonstrations, such as (Blumenkamp et al., 2022), further highlight that variable delays and unreliable communication links must be considered in system design, while resource allocation in wireless mesh networks (Wang et al., 2022) has been shown to significantly affect application performance.

Client scheduling and topology management. Client scheduling and topology management for collaborative GNNs are mainly considered for subgraphs with topology partitioning in learning tasks over wide-area networks, where the message-passing operation is costly. In (Liu et al., 2021), collaborative GNNs in wide-area networks are considered. To reduce traffic congestion, client selection is performed based on the importance of the communicated embeddings and the communication costs. Scheduling algorithms are designed to allow interleaved message passing and local updates. A multi-tier wide-area network with edge servers is considered in (Zeng et al., 2022; Nazzal et al., 2024). The topology of the GNN is given by the application, and these papers address the problem of GNN node assignment to edge servers and edge-network topology design for communication-efficient training, for both static (Zeng et al., 2022) and dynamic (Nazzal et al., 2024) GNN topologies.

Most works on communication efficiency in collaborative GNN implementations assume that the communication graph given by the possible connections between agents is identical to the computing graph given by the adjacency matrix A. This assumption is not necessary. An algorithm to find the optimal communication graph for subgraphs with topology partitioning is proposed in (Scardapane et al., 2020); however, this area of research remains largely unexplored. Conversely, (Wang et al., 2025) assumes a fully connected communication graph (i.e., agents are free to communicate with anyone in a peer-to-peer network) but samples a communication topology at each round. The authors design an adaptive algorithm that samples communication topologies adaptively throughout training to accommodate bandwidth constraints while maintaining fast inference and training.

Open challenges. Existing work has started to address the challenges of communicating embeddings and model parameters over unreliable or costly communication channels or transmission paths. However, since both inference and learning involve communication, several interesting research questions remain open:

- *Towards co-designing communication and computing.* While recent methods consider the effects of unreliable communication channels (Lee et al., 2021; Gao et al., 2021), they treat communication and computation topologies independently. Co-designing these topologies while explicitly accounting for channel quality represents an avenue for further efficiency improvements. Additionally, existing embedding aggregation methods operate on static neighbor selection policies and could be extended to dynamically exploit time-varying wireless channel quality by scheduling or selecting neighbors with temporarily favorable channel conditions.

- *Towards communication-native inference.* The wireless medium itself, with its superposition and broadcast properties, natively supports *many-to-one* and *one-to-many* communications. These properties could allow efficient embedding and model aggregation in a single operation among clusters of agents or agents connecting to the same base station or edge computing server. While initial results exist in (Fiorellino et al., 2024), further research is needed to fully exploit these network environment properties through complex topological structures beyond simple graphs.

- *Towards safe training and inference.* Training and inference on graph data require communication in both phases. Consequently, communication impairments deteriorate inference quality even when the model itself is optimal, potentially leading to unsafe decisions. This raises the question of whether safety requirements should be incorporated into the learning process or considered only during inference, a design choice that requires further investigation in networked systems where communication reliability varies.

3.6 Privacy Preservation Techniques for Graph Data

A high variety of privacy-preserving techniques are proposed for distributed graph data, aiming to protect training and inference data from *honest-but-curious* agents and central servers. Anonymization can target three main levels: *(i)* model parameters, *(ii)* exchanged node embeddings, and *(iii)* exchanged edge information. Many of the techniques build on blind computation methods and LDP introduced in Section 2.5.

Anonymizing model parameters. Model parameters or gradients exchanged under GNN training can be anonymized directly with the techniques proposed for FL in Section 2.5 (Liu et al., 2022; Wu et al., 2022; Yao et al., 2023; Yan et al., 2024). These techniques can be further refined to improve the stability–privacy trade-off, such as *adaptive noise scaling* (Wu et al., 2022), in which the injected noise magnitude is adjusted according to the gradient norm to maintain a consistent perturbation throughout training.

Anonymizing exchanged embeddings. In the case of graph data, collaborative learning and inference often require the exchange of intermediate data such as embeddings or representations. Early works like (Chen et al., 2021) incorrectly assumed that if local computations are performed entirely in isolation, including local model updates and isolated inference, then the raw node features X_k would be inherently protected from privacy leakage. However, it is now clear that even when only intermediate embeddings $h_v^{(l)}$ are shared and aggregated, they must still be protected through some form of differential privacy (Chen et al., 2024).

- *Blind computation* methods securely aggregate embeddings (e.g., computing f_{Agg} in (18)) or average representations $z_v(t)$ without revealing their values. For example, secure multiparty computation can be applied to jointly perform specific GNN functions such as aggregating f_{Agg} (18) or computing matrix multiplications. Due to their high computational cost, these methods are often applied only to a subset of layers (Chen et al., 2021), with the remaining computations delegated to the server in a split-learning setup. Homomorphic encryption offers another option, allowing the server to aggregate results while keeping them encrypted end-to-end (Ni et al., 2021), using techniques such as secure multiparty computation (Chen et al., 2021) or homomorphic encryption for server-side aggregation (Ni et al., 2021).

- *LDP* techniques can be applied directly to the communicated vectors z_v or $h_v^{(l)}$. The simplest approach adds Gaussian noise to the vectors (Zhang et al., 2023; Lee et al., 2023; Ni et al., 2021). More advanced methods adapt classical LDP protocols (Dwork & Roth, 2014; Erlingsson et al., 2014; Choi et al., 2018; Ding et al., 2017) to the GNN setting (Lee et al., 2023; Wu et al., 2022; Liu et al., 2022; Wu et al., 2021; Chen et al., 2021; Qiu et al., 2022; Yan et al., 2024). These generally follow a pipeline of noise injection, vector clipping, dimension subsampling, and privacy budget adjustment. Noise can follow a Laplace distribution (Wu et al., 2022; 2021), a Gaussian distribution (Chen et al., 2021), or a Bernoulli distribution (Zheng et al., 2023; Pan et al., 2023). Clipping may be applied to vectors or gradients (Wu et al., 2022; 2021; Chen et al., 2021; Liu et al., 2022; Yan et al., 2024) and can be combined with stochastic quantization (Lin et al., 2022; Zheng et al., 2023; Pan et al., 2023). Dimensionality reduction can be performed by random subsampling (Lin et al., 2022; Pei et al., 2023; Zheng et al., 2023; Pan et al., 2023) or top-k sampling (Chen et al., 2022). In some cases, these operations are unified through complex noise functions such as piecewise noise injection (Pei et al., 2023; Wang et al., 2019). The privacy level can also be balanced against other objectives or system constraints. The noise added in LDP is known for degrading learning effectiveness (Lee et al., 2023), mitigated while aggregating high degree nodes (Guo et al., 2024). Therefore, adaptive

noise strategies propose to scale the noise by the gradient norm during training to stabilize learning citep wu2022federated. When communication occurs over wireless channels, the trade-off between signal-to-noise ratio and differential privacy can be jointly optimized (Lee et al., 2023). LDP can also be applied to node labels in transfer learning with adversarial losses to protect the training data of the source agent (Guan et al., 2021).

Anonymizing edge (or topology) information. Inferring the edges between nodes, so-called *link stealing* (Pei et al., 2023), can mean a serious privacy attack when the existence of the edge itself represents sensitive information, or when the edges encode most of the information, as is often the case in heterogeneous graphs (Liu et al., 2022; Wu et al., 2021; 2022; Yan et al., 2024). Therefore, the edge or topology information needs to be protected in many scenarios. The concept of *Edge Local Differential Privacy (Edge LDP)* is formalized in (Wu et al., 2022; Han et al., 2024; Lin et al., 2022) and is extended to heterogeneous graphs in (Yan et al., 2024). In some cases, protection also covers statistics on the edge set, such as degree distributions (Pan et al., 2023). Protection strategies build both on blind computations and on LDP.

- *Blind computations* can be employed for *privacy-preserving subgraph expansion.* Each agent obtains its L-hop expansion $\tilde{G}_k$ of the local graph using blind computation techniques. The process begins with agents collectively running a private set intersection (PSI) protocol to identify the vertex cut $mathcalV_{cut}$ between agents. Given the knowledge of $\mathcal{V}_{cut} \cap \mathcal{V}_k$, each agent collects and communicates the missing expanded subgraphs, typically via a central server (Zhang et al., 2022; Chen et al., 2021; 2022; Agrawal et al., 2024; Qiu et al., 2022). Blind computation techniques can also be applied to aggregate useful graph statistics, such as node degree (Pan et al., 2023), using methods as secure group aggregation (Bonawitz et al., 2017) or zero-knowledge protocols (Goldwasser et al., 1985). Such statistics may be required to compute auxiliary losses, for example, fairness losses (Agrawal et al., 2024), or may be used by the local learning algorithm (Pan et al., 2023).

- *Standard LDP* (Dwork & Roth, 2014) can be applied to edge information in set form $\mathcal{E}_k$ (Yan et al., 2024; Lin et al., 2022) or in adjacency matrix form A_k (Han et al., 2024; Qiu et al., 2022). For example, (Zhang et al., 2021) projects A_k into a lower-dimensional Gaussian random subspace. However, LDP introduces topological noise, which is known to degrade GNN performance even at low noise levels (Zügner et al., 2020). Excessive noise can densify the graph, leading to *oversmoothing* (Chen et al., 2020). Mitigation strategies include regularization (Lin et al., 2022), degree preservation (Yan et al., 2024), or allocating part of the privacy budget to maintain adjacency sparsity (Qiu et al., 2022).

Protecting edge information is often formulated alongside protecting the raw feature X_k, thus protecting the local graph data $G_k = (A_k, X_k)$. The concept of node-level privacy is introduced for the ego-graph case (Guo et al., 2024), consisting of protecting the raw features of nodes v as well as the 1-hop topology typical of ego-graph partitions. This concept is further extended to heterogeneous graphs, specifically focusing on protecting nodes that correspond to users (Yan et al., 2024). Other works focus on protecting collections of ego-graphs sampled from G_k (Zhang et al., 2023; Wu et al., 2021; 2022). The concept of subgraph-level privacy is defined in (Guo et al., 2024), protecting the subgraph data G_k obtained from aggregating the private ego-graphs of the agents. Techniques that anonymize A_k together with X_k have been proposed in the following forms:

- *Blending ego-graphs together.* The embeddings of ego-graphs are averaged across all ego-graphs subsampled from G_k in the same batch, respectively between nodes at the same depth in the computational tree, to anonymize the ego-graphs and their embeddings (Zhang et al., 2023).

- *Inferring pseudo-edges.* This common technique is used for anonymizing A_k for heterogeneous graphs, usually in the vertex cut case (Liu et al., 2022; Wu et al., 2021; 2022; Yan et al., 2024). It consists of constructing unobserved yet highly probable edges based on node representation similarity, and adding them to the observed edges in G_k.

Hierarchical scenarios. In hierarchical setups (e.g., device–edge and server–central server) (Guo et al., 2024), privacy protections can be considered between higher-rank and lower-rank entities or among peers of the same rank. The above anonymization techniques, focusing on model parameters, embeddings, and topology information, can then be selectively combined depending on the nature of the data exchanged at each hierarchy level.

Open challenges. Privacy-preserving techniques for collaborative Graph Neural Networks (GNNs) operate at the intersection of two research domains: privacy mechanisms for GNNs and privacy-preserving collaborative training and inference. This convergence presents unique challenges and opportunities, as collaborative settings introduce multi-agent dynamics that both complicate and potentially enhance privacy guarantees. Two critical research directions emerge from this intersection.

- *Towards multi-agent privacy amplification.* Most existing work applying DP to GNNs relies on LDP, where privacy-preserving techniques are applied locally to individual training or inference data. However, LDP techniques have been shown to compress communicated information in ways that significantly degrade learning effectiveness (Smith et al., 2017; Bonawitz et al., 2022). In collaborative settings, privacy amplification, seen in Section 2.5, offers a promising alternative. Rather than applying hard anonymization, this approach allows agents to sequentially amplify protection guarantees along the chain of collaborative computations, thereby limiting information loss. Privacy amplification fits well with the sequential operation of MPNNs and is worth exploring in future research.
- *Towards tighter integration with defense mechanisms.* While anonymization techniques are designed to protect privacy against *honest-but-curious* agents, this protection proves insufficient against advanced threat scenarios involving *malicious agents*. Consequently, privacy mechanisms are frequently motivated as the *first* defense layer against adversarial inference attacks, such as node label manipulation and inference (Zügner et al., 2020; Chen et al., 2022), link stealing (Pei et al., 2023), node embedding extraction (Chen et al., 2024), and graph rewiring (Chen et al., 2024). Given these threats, privacy mechanisms are proposed to incorporate additional defense mechanisms, such as embedding filtering (He et al., 2024). However, how privacy mechanisms should integrate with defenses against increasingly sophisticated threat scenarios remains an open research question.

4 Conclusion

Collaborative learning has unlocked access to previously unavailable distributed data sources, leading to the success of federated learning and decentralized learning on Euclidean data, and more recently, to emerging solutions for graph-structured data. Extending collaborative learning to graph-structured data opens up opportunities that exceed those of centralized graph ML through its ability to mine data across domains ranging from sensor networks to molecular biology.

With this survey, we systematically map the design choices from Euclidean-structured to graph-structured collaborative learning. We first consolidate the foundational principles of collaborative learning on Euclidean data (Section 2), then extend this mapping to graph-structured data (Section 3). For both cases, we formulate the fundamental problem, discuss the challenges of data and system heterogeneity, and analyze how the trilemma of effectiveness, efficiency, and privacy preservation manifests in the design choices for addressing these challenges. Specifically for graph-data, we identify emerging research challenges that define its future development.

Our analysis reveals that collaborative learning on graph-structured data has given rise to a nascent, yet rich, landscape of problems and solutions. This landscape can be broadly divided into two paradigms: *isolated inference* and *collaborative inference*. Notably, collaborative inference emerges as a more central paradigm for graph-structured data due to the message passing mechanism of GNNs. Beyond challenges that persist in Euclidean collaborative learning, combining distributed relational data with collaborative training and/or inference introduces unique constraints that require novel solutions. This includes the joint consideration of topologies, communication, privacy, and heterogeneities.

By systematically bridging collaborative learning on Euclidean and graph-structured data, this survey establishes a coherent vision to support researchers in developing the next generation of collaborative graph learning methods and to enable their practical impact across scientific and industrial domains, leading to societally important breakthroughs in drug discovery and healthcare, sustainable industries, and smart cities.

References

[1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *ACM SIGSAC Conference on Computer and Communications Security*, 2016.

[2] Emmanuel Abbe. Community detection and stochastic block models: Recent developments. *Journal of Machine Learning Research*, 18(177):1–86, 2018.

[3] Naman Agarwal, Ananda Theertha Suresh, Felix Xinnan X Yu, Sanjiv Kumar, and Brendan McMahan. cpSGD: Communication-efficient and differentially-private distributed SGD. *Advances in Neural Information Processing Systems*, 31, 2018.

[4] Nimesh Agrawal, Anuj Kumar Sirohi, and Sandeep Kumar. No prejudice! Fair federated graph neural networks for personalized recommendation. In *AAAI Conference on Artificial Intelligence*, 2024.

[5] Galen Andrew, Om Thakkar, Brendan McMahan, and Swaroop Ramaswamy. Differentially private learning with adaptive clipping. *Advances in Neural Information Processing Systems*, 34:17455–17466, 2021.

[6] Apple. Private cloud compute. Technical report, Apple, 2023.

[7] Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv:1912.00818*, 2019.

[8] Seyed Mohammad Azimi-Abarghouyi and Carlo Fischione. Multi-layer hierarchical federated learning with quantization. *arXiv:2505.08145*, 2025.

[9] Seyed Mohammad Azimi-Abarghouyi and Viktoria Fodor. Scalable hierarchical over-the-air federated learning. *IEEE Transactions on Wireless Communications*, 2024.

[10] Stefan Baack. A critical analysis of the largest source for generative ai training data: Common crawl. In *ACM Conference on Fairness, Accountability, and Transparency*, 2024.

[11] Jinheon Baek, Wonyong Jeong, Jiongdao Jin, Jaehong Yoon, and Sung Ju Hwang. Personalized subgraph federated learning. In *International Conference on Machine Learning*, 2023.

[12] Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. The privacy blanket of the shuffle model. In *Annual International Cryptology Conference*, 2019.

[13] Borja Balle, Peter Kairouz, Brendan McMahan, Om Thakkar, and Abhradeep Guha Thakurta. Privacy amplification via random check-ins. *Advances in Neural Information Processing Systems*, 33:4623–4634, 2020.

[14] Claudio D. T. Barros, Matheus R. F. Mendonça, Alex B. Vieira, and Artur Ziviani. A survey on embedding dynamic graphs. *ACM Computing Surveys*, 55(1):1–37, 2021.

[15] Raef Bassily, Adam Smith, and Abhradeep Thakurta. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *IEEE Annual Symposium on Foundations of Computer Science*, 2014.

[16] Mayank Bawa, Hector Garcia-Molina, Aristides Gionis, and Rajeev Motwani. Estimating aggregates on a peer-to-peer network. *Stanford InfoLab*, 2003.

[17] Gleb Bazhenov, Denis Kuznedelev, Andrey Malinin, Artem Babenko, and Liudmila Prokhorenkova. Evaluating robustness and uncertainty of graph models under structural distributional shifts. *Advances in Neural Information Processing Systems*, 36:75567–75594, 2023.

[18] Maya Bechler-Speicher, Ben Finkelshtein, Fabrizio Frasca, Luis Müller, Jan Tönshoff, Antoine Siraudin, Viktor Zaverkin, Michael M. Bronstein, Mathias Niepert, Bryan Perozzi, Mikhail Galkin, and Christopher Morris. Position: Graph learning will lose relevance due to poor benchmarks. In *International Conference on Machine Learning*, 2025.

[19] Florence Bénézit, Vincent Blondel, Patrick Thiran, John Tsitsiklis, and Martin Vetterli. Weighted gossip: Distributed averaging using non-doubly stochastic matrices. In *IEEE International Symposium on Information Theory*, 2010.

[20] Florence Bénézit, Alexandros G Dimakis, Patrick Thiran, and Martin Vetterli. Order-optimal consensus through randomized path averaging. *IEEE Transactions on Information Theory*, 56(10):5150–5167, 2010.

[21] Yoshua Bengio, Salem Lahlou, Tristan Deleu, Edward J Hu, Mo Tiwari, and Emmanuel Bengio. Gflownet foundations. *Journal of Machine Learning Research*, 24(210):1–55, 2023.

[22] Dimitri Bertsekas and John Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 2015.

[23] Dimitri P Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. *Optimization for Machine Learning*, 2010(1-38):3, 2011.

[24] Abhishek Bhowmick, John Duchi, Julien Freudiger, Gaurav Kapoor, and Ryan Rogers. Protection against reconstruction and its applications in private federated learning. *arXiv:1812.00984*, 2018.

[25] Ilai Bistritz, Ariana Mann, and Nicholas Bambos. Distributed distillation for on-device learning. *Advances in Neural Information Processing Systems*, 33:22593–22604, 2020.

[26] Sayan Biswas, Davide Frey, Romaric Gaudel, Anne-Marie Kermarrec, Dimitri Lerévérend, Rafael Pires, Rishi Sharma, and François Taïani. Low-cost privacy-preserving decentralized learning. In *Proceedings on Privacy Enhancing Technologies Symposium*, 2025.

[27] Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Symposium on Operating Systems Principles*, 2017.

[28] Vincent D Blondel, Julien M Hendrickx, Alex Olshevsky, and John N Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *IEEE Conference on Decision and Control*, 2005.

[29] Michael Blot, David Picard, Matthieu Cord, and Nicolas Thome. Gossip training for deep learning. In *Conference on Neural Information Processing Systems*, 2016.

[30] Jan Blumenkamp, Steven Morad, Jennifer Gielis, Qingbiao Li, and Amanda Prorok. A framework for real-world multi-robot systems running decentralized GNN-based policies. In *IEEE International Conference on Robotics and Automation*, 2022.

[31] Cristian Bodnar, Francesco Di Giovanni, Benjamin Chamberlain, Pietro Lio, and Michael Bronstein. Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in GNNs. *Advances in Neural Information Processing Systems*, 35:18527–18541, 2022.

[32] Aleksandar Bojchevski and Stephan Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. In *International Conference on Learning Representations*, 2018.

[33] Kallista Bonawitz, Peter Kairouz, Brendan McMahan, and Daniel Ramage. Federated learning and privacy. *Communications of the ACM*, 19(5):87–114, 2022.

[34] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *ACM SIGSAC Conference on Computer and Communications Security*, 2017.

[35] Marco Bornstein, Tahseen Rabbani, Evan Wang, Amrit S. Bedi, and Furong Huang. SWIFT: rapid decentralized federated learning via wait-free model communication. In *International Conference on Learning Representations*, 2023.

[36] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of Computational Statistics*, 2010.

[37] Léon Bottou and Olivier Bousquet. The tradeoffs of large scale learning. *Advances in Neural Information Processing Systems*, 20, 2007.

[38] Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.

[39] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Gossip algorithms: Design, analysis and applications. In *IEEE Infocom Conference on Computer Communications*, 2005.

[40] Stephen Boyd, Arpita Ghosh, Balaji Prabhakar, and Devavrat Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.

[41] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.

[42] Christopher Briggs, Zhong Fan, and Peter Andras. Federated learning with hierarchical clustering of local updates to improve training on non-iid data. In *IEEE International Joint Conference on Neural Networks*, 2020.

[43] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. *Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges*. arXiv:2104.13478, 2021.

[44] Duc Bui, Kshitiz Malik, Jack Goetz, Honglei Liu, Seungwhan Moon, Anuj Kumar, and Kang G Shin. Federated user representation learning. *arXiv:1909.12535*, 2019.

[45] Ming Cao, Daniel A Spielman, and Edmund M Yeh. Accelerated gossip algorithms for distributed computation. In *The Annual Allerton Conference on Communication, Control, and Computation*, 2006.

[46] Sergio Casas, Cole Gulino, Renjie Liao, and Raquel Urtasun. Spagnn: Spatially-aware graph neural networks for relational behavior forecasting from sensor data. In *IEEE International Conference on Robotics and Automation*, 2020.

[47] Timothy Castiglia, Anirban Das, and Stacy Patterson. Multi-level local SGD for heterogeneous hierarchical networks. In *International Conference on Learning Representations*, 2021.

[48] Emir Ceyani, Han Xie, Baturalp Buyukates, Carl Yang, and Salman Avestimehr. Fedgrains: Personalized subgraph federated learning with adaptive neighbor sampling. In *SIAM International Conference on Data Mining*, 2025.

[49] Di Chai, Leye Wang, Kai Chen, and Qiang Yang. Secure federated matrix factorization. *IEEE Intelligent Systems*, 36(5):11–20, 2020.

[50] Zheng Chai, Ahsan Ali, Syed Zawad, Stacey Truex, Ali Anwar, Nathalie Baracaldo, Yi Zhou, Heiko Ludwig, Feng Yan, and Yue Cheng. Tifl: A tier-based federated learning system. In *International Symposium on High-Performance Parallel and Distributed Computing*, 2020.

[51] Zheng Chai, Yujing Chen, Ali Anwar, Liang Zhao, Yue Cheng, and Huzefa Rangwala. Fedat: A high-performance and communication-efficient federated learning system with asynchronous tiers. In *International Conference for High Performance Computing, Networking, Storage and Analysis*, 2021.

[52] T-H Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *ACM Transactions on Information and System Security*, 14(3):1–24, 2011.

[53] Chaochao Chen, Jun Zhou, Longfei Zheng, Huiwen Wu, Lingjuan Lyu, Jia Wu, Bingzhe Wu, Ziqi Liu, Li Wang, and Xiaolin Zheng. Vertically federated graph neural network for privacy-preserving node classification. In *International Joint Conference on Artificial Intelligence*, 2021.

[54] Chuan Chen, Ziyue Xu, Weibo Hu, Zibin Zheng, and Jie Zhang. Fedgl: Federated graph learning framework with global self-supervision. *Information Sciences*, 657:119976, 2024.

[55] Deli Chen, Yankai Lin, Wei Li, Peng Li, Jie Zhou, and Xu Sun. Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *AAAI Conference on Artificial Intelligence*, 2020.

[56] Fahao Chen, Peng Li, Toshiaki Miyazaki, and Celimuge Wu. Fedgraph: Federated graph learning with intelligent sampling. *IEEE Transactions on Parallel and Distributed Systems*, 33(8):1775–1786, 2021.

[57] Fengwen Chen, Guodong Long, Zonghan Wu, Tianyi Zhou, and Jing Jiang. Personalized federated learning with a graph. In *International Joint Conference on Artificial Intelligence*, 2022.

[58] Jialong Chen, Tianchi Liao, Chuan Chen, and Zibin Zheng. Improving message-passing GNNs by asynchronous aggregation. In *ACM International Conference on Information and Knowledge Management*, 2024.

[59] Jiayi Chen and Aidong Zhang. Fedmsplit: Correlation-adaptive federated multi-task learning across multimodal split networks. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.

[60] Jingdi Chen, Hanqing Yang, Zongjun Liu, and Carlee Joe-Wong. The five Ws of multi-agent communication: Who talks to whom, when, what, and why-a survey from MARL to emergent language and LLMs. *Authorea Preprints*, 2025.

[61] Jinyin Chen, Guohan Huang, Haibin Zheng, Shanqing Yu, Wenrong Jiang, and Chen Cui. Graph-fraudster: Adversarial attacks on graph neural network-based vertical federated learning. *IEEE Transactions on Computational Social Systems*, 10(2):492–506, 2022.

[62] Jinyin Chen, Ximin Zhang, and Haibin Zheng. Adversarial attacks on GNN-based vertical federated learning. In *Attacks, Defenses and Testing for Deep Learning*. Springer, 2024.

[63] Mingyang Chen, Wen Zhang, Zonggang Yuan, Yantao Jia, and Huajun Chen. Fede: Embedding knowledge graphs in federated setting. In *International Joint Conference on Knowledge Graphs*, 2021.

[64] Mingyang Chen, Wen Zhang, Zonggang Yuan, Yantao Jia, and Huajun Chen. Federated knowledge graph completion via embedding-contrastive learning. *Knowledge-Based Systems*, 252:109459, 2022.

[65] Mingzhe Chen, H. Vincent Poor, Walid Saad, and Shuguang Cui. Convergence time optimization for federated learning over wireless networks. *IEEE Transactions on Wireless Communications*, 20: 2457–2471, 2020.

[66] Albert Cheu, Adam Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, 2019.

[67] Tsz-Him Cheung, Weihang Dai, and Shuhan Li. Fedsgc: Federated simple graph convolution for node classification. In *International Workshop on Federated and Transfer Learning for Data Sparsity and Confidentiality*, 2021.

[68] Yae Jee Cho, Jianyu Wang, and Gauri Joshi. Client selection in federated learning: Convergence analysis and power-of-choice selection strategies. *arXiv:2010.01243*, 2020.

[69] Yae Jee Cho, Pranay Sharma, Gauri Joshi, Zheng Xu, Satyen Kale, and Tong Zhang. On the convergence of federated averaging with cyclic client participation. In *International Conference on Machine Learning*, 2023.

[70] Woo-Seok Choi, Matthew Tomei, Jose Rodrigo Sanchez Vicarte, Pavan Kumar Hanumolu, and Rakesh Kumar. Guaranteeing local differential privacy on ultra-low-power systems. In *ACM/IEEE Annual International Symposium on Computer Architecture*, 2018.

[71] Ching-Yao Chuang and Stefanie Jegelka. Tree mover's distance: Bridging graph metrics and stability of graph neural networks. *Advances in Neural Information Processing Systems*, 35:2944–2957, 2022.

[72] Andrea Cini, Ivan Marisca, Daniele Zambon, and Cesare Alippi. Graph deep learning for time series forecasting. *ACM Computing Surveys*, 57(12):1–34, 2025.

[73] Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared representations for personalized federated learning. In *International Conference on Machine Learning*, 2021.

[74] Jorge Cortés. Finite-time convergent gradient flows with applications to network consensus. *Automatica*, 42(11):1993–2000, 2006.

[75] Edwige Cyffers, Aurélien Bellet, and Jalaj Upadhyay. Differentially private decentralized learning with random walks. In *International Conference on Machine Learning*, 2024.

[76] Ali Dadras, Sebastian U Stich, and Alp Yurtsever. Personalized federated learning via low-rank matrix factorization. In *Optimization for Machine Learning*, 2024.

[77] Rong Dai, Li Shen, Fengxiang He, Xinmei Tian, and Dacheng Tao. DisPFL: Towards communication-efficient personalized federated learning via decentralized sparse training. In *International Conference on Machine Learning*, 2022.

[78] Katharine Daly, Hubert Eichner, Peter Kairouz, H Brendan McMahan, Daniel Ramage, and Zheng Xu. Federated learning in practice: Reflections and projections. In *IEEE International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications*, 2024.

[79] Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc'aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc V. Le, and Andrew Y. Ng. Large scale distributed deep networks. *Advances in Neural Information Processing Systems*, 25, 2012.

[80] Morris H DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345): 118–121, 1974.

[81] Patrick Denantes, Florence Bénézit, Patrick Thiran, and Martin Vetterli. Which distributed averaging algorithm should i choose for my sensor network? In *IEEE Infocom Conference on Computer Communications*, 2008.

[82] Yongheng Deng, Weining Chen, Ju Ren, Feng Lyu, Yang Liu, Yunxin Liu, and Yaoxue Zhang. TailorFL: Dual-personalized federated learning under system and data heterogeneity. In *ACM Conference on Embedded Networked Sensor Systems*, 2022.

[83] Yuyang Deng, Mohammad Mahdi Kamani, and Mehrdad Mahdavi. Adaptive personalized federated learning. *arXiv:2003.13461*, 2020.

[84] Francesco Di Giovanni, James Rowbottom, Benjamin P Chamberlain, Thomas Markovich, and Michael M Bronstein. Graph neural networks as gradient flows: Understanding graph convolutions via energy. *arXiv:2206.10991*, 2022.

[85] Enmao Diao, Jie Ding, and Vahid Tarokh. HeteroFL: Computation and communication efficient federated learning for heterogeneous clients. In *International Conference on Learning Representations*, 2021.

[86] Alexandros DG Dimakis, Anand D Sarwate, and Martin J Wainwright. Geographic gossip: Efficient averaging for sensor networks. *IEEE Transactions on Signal Processing*, 56(3):1205–1216, 2008.

[87] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. *Advances in Neural Information Processing Systems*, 30, 2017.

[88] Guimin Dong, Mingyue Tang, Lihua Cai, Laura E Barnes, and Mehdi Boukhechba. Semi-supervised graph instance transformer for mental health inference. In *IEEE International Conference on Machine Learning and Applications*, 2021.

[89] Guimin Dong, Mingyue Tang, Zhiyuan Wang, Jiechao Gao, Sikun Guo, Lihua Cai, Robert Gutierrez, Bradford Campbel, Laura E Barnes, and Mehdi Boukhechba. Graph neural networks in IoT: A survey. *ACM Transactions on Sensor Networks*, 19(2):1–50, 2023.

[90] Bingqian Du and Chuan Wu. Federated graph learning with periodic neighbour sampling. In *IEEE/ACM International Symposium on Quality of Service*, 2022.

[91] Moming Duan, Duo Liu, Xianzhang Chen, Renping Liu, Yujuan Tan, and Liang Liang. Self-balancing federated learning with global imbalanced data in mobile systems. *IEEE Transactions on Parallel and Distributed Systems*, 32(1):59–71, 2020.

[92] Moming Duan, Duo Liu, Xinyuan Ji, Renping Liu, Liang Liang, Xianzhang Chen, and Yujuan Tan. Fedgroup: Efficient federated learning via decomposed similarity-based clustering. In *IEEE Internatioanal Conference on Parallel & Distributed Processing with Applications, Big Data & Cloud Computing, Sustainable Computing & Communications, Social Computing & Networking*, 2021.

[93] John Duchi and Ryan Rogers. Lower bounds for locally private estimation via communication complexity. In *Conference on Learning Theory*, 2019.

[94] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 2011.

[95] John C Duchi, Michael I Jordan, and Martin J Wainwright. Local privacy and statistical minimax rates. In *IEEE Annual Symposium on Foundations of Computer Science*, 2013.

[96] John C Duchi, Michael I Jordan, and Martin J Wainwright. Minimax optimal procedures for locally private estimation. *Journal of the American Statistical Association*, 113(521):182–201, 2018.

[97] Chi Thang Duong, Thanh Dat Hoang, Ha The Hien Dang, Quoc Viet Hung Nguyen, and Karl Aberer. On node features for graph neural networks. *arXiv:1911.08795*, 2019.

[98] Cynthia Dwork. Differential privacy. In *International Colloquium on Automata, Languages, and Programming*, 2006.

[99] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

[100] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography*, 2006.

[101] Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *ACM Symposium on Theory of Computing*, 2010.

[102] Abdellah El Mrini, Edwige Cyffers, and Aurélien Bellet. Privacy attacks in decentralized learning. In *International Conference on Machine Learning*, 2024.

[103] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *ACM SIGSAC Conference on Computer and Communications Security*, 2014.

[104] Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In *Annual ACM-SIAM Symposium on Discrete Algorithms*, 2019.

[105] European Union. General data protection regulation. Technical report, European Union, 2016.

[106] Mathieu Even, Anastasia Koloskova, and Laurent Massoulie. Asynchronous SGD on graphs: A unified framework for asynchronous decentralized and federated optimization. In *International Conference on Artificial Intelligence and Statistics*, 2024.

[107] Alexandre Evfimievski, Johannes Gehrke, and Ramakrishnan Srikant. Limiting privacy breaches in privacy preserving data mining. In *ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, 2003.

[108] Lukas Faber and Roger Wattenhofer. GwAC: GNNs with asynchronous communication. In *Learning on Graphs Conference*, 2024.

[109] Lukas Faber, Yifan Lu, and Roger Wattenhofer. Should graph neural networks use features, edges, or both? *arXiv:2103.06857*, 2021.

[110] Alireza Fallah, Aryan Mokhtari, and Asuman Ozdaglar. Personalized federated learning: A meta-learning approach. *arXiv:2002.07948*, 2020.

[111] Zhenan Fan, Huang Fang, and Michael P Friedlander. A dual approach for federated learning. *arXiv:2201.11183*, 2022.

[112] Hui Fang, Yang Gao, Peng Zhang, Jiangchao Yao, Hongyang Chen, and Haishuai Wang. Large language models enhanced personalized graph neural architecture search in federated learning. In *AAAI Conference on Artificial Intelligence*, 2025.

[113] Bahare Fatemi, Layla El Asri, and Seyed Mehran Kazemi. Slaps: Self-supervision improves structure learning for graph neural networks. *Advances in Neural Information Processing Systems*, 34:22667–22681, 2021.

[114] J Alexander Fax and Richard M Murray. Information flow and cooperative control of vehicle formations. *IEEE Transactions on Automatic Control*, 49(9):1465–1476, 2004.

[115] Vitaly Feldman, Audra McMillan, and Kunal Talwar. Hiding among the clones: A simple and nearly optimal analysis of privacy amplification by shuffling. In *IEEE Annual Symposium on Foundations of Computer Science*, 2022.

[116] Siwei Feng. Vertical federated learning-based feature selection with non-overlapping sample utilization. *Expert Systems with Applications*, 208:118097, 2022.

[117] Siwei Feng, Boyang Li, Han Yu, Yang Liu, and Qiang Yang. Semi-supervised federated heterogeneous transfer learning. *Knowledge-Based Systems*, 252:109384, 2022.

[118] Matthias Fey, Jan E. Lenssen, Frank Weichert, and Jure Leskovec. GNNAutoscale: Scalable and expressive graph neural networks via historical embeddings. In *International Conference on Machine Learning*, 2021.

[119] Matthias Fey, Jinu Sunil, Akihiro Nitta, Rishi Puri, Manan Shah, Blaž Stojanovič, Ramona Bendias, Alexandria Barghi, Vid Kocijan, Zecheng Zhang, Xinwei He, Jan Eric Lenssen, and Jure Leskovec. Pyg 2.0: Scalable learning on real world graphs. In *KDD Worshop on Temporal Graph Learning*, 2025.

[120] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.

[121] Ben Finkelshtein, Xingyue Huang, Michael Bronstein, and İsmail İlkan Ceylan. Cooperative graph neural networks. In *International Conference on Machine Learning*, 2024.

[122] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International Conference on Machine Learning*, 2017.

[123] Simone Fiorellino, Claudio Battiloro, and Paolo Di Lorenzo. Topological neural networks over the air. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2024.

[124] Fabrizio Frasca, Fabian Jogl, Moshe Eliasof, Matan Ostrovsky, Carola-Bibiane Schönlieb, Thomas Gärtner, and Haggai Maron. Towards foundation models on graphs: An analysis on cross-dataset transfer of pretrained GNNs. In *NEURIPS Workshop on Symmetry and Geometry in Neural Representations*, 2024.

[125] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *ACM SIGSAC Conference on Computer and Communications Security*, 2015.

[126] Xingbo Fu, Song Wang, Yushun Dong, Binchi Zhang, Chen Chen, and Jundong Li. Federated graph learning with graphless clients. *Transactions on Machine Learning Research*, 2024.

[127] Dominik Fuchsgruber, Tom Wollschläger, Johannes Bordne, and Stephan Günnemann. Uncertainty estimation for heterophilic graphs through the lens of information theory. In *International Conference on Machine Learning*, 2025.

[128] Huiguo Gao, Mengyuan Lee, Guanding Yu, and Zhaolin Zhou. A graph neural network based decentralized learning scheme. *Sensors*, 22(3):1030, 2022.

[129] Rufei Gao, Zhaowei Liu, Chenxi Jiang, Yingjie Wang, Shenqiang Wang, and Pengda Wang. BI-fedGNN: Federated graph neural networks framework based on bayesian inference. *Neural Networks*, 169:143–153, 2024.

[130] Zhan Gao and Deniz Gündüz. Airgnns: Graph neural networks over the air. In *IEEE European Signal Processing Conference*, 2023.

[131] Zhan Gao and Deniz Gündüz. Graph neural networks over the air for decentralized tasks in wireless networks. *IEEE Transactions on Signal Processing*, 2025.

[132] Zhan Gao, Elvin Isufi, and Alejandro Ribeiro. Stochastic graph neural networks. *IEEE Transactions on Signal Processing*, 69:4428–4443, 2021.

[133] Zhan Gao, Yulin Shao, Deniz Gündüz, and Amanda Prorok. Decentralized channel management in WLANs with graph neural networks. In *IEEE International Conference on Communications*, 2023.

[134] Simon Markus Geisler, Arthur Kosmala, Daniel Herbst, and Stephan Günnemann. Spatio-spectral graph neural networks. *Advances in Neural Information Processing Systems*, 37:49022–49080, 2024.

[135] Robin C Geyer, Tassilo Klein, and Moin Nabi. Differentially private federated learning: A client level perspective. In *NEURIPS Workshop on Machine Learning on the Phone and Other Consumer Devices*, 2017.

[136] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for clustered federated learning. *Advances in Neural Information Processing Systems*, 33:19586–19597, 2020.

[137] Lodovico Giaretta and Sarunas Girdzijauskas. Fully-decentralized training of GNNs using layer-wise self-supervision. *diva2:1745180*, 2023.

[138] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, 2017.

[139] Jack Goetz, Kshitiz Malik, Duc Bui, Seungwhan Moon, Honglei Liu, and Anuj Kumar. Active federated learning. Technical report, University of Michigan, 2019.

[140] Shafi Goldwasser, Silvio Micali, and Chales Rackoff. The knowledge complexity of interactive proof-systems. In *ACM Symposium on Theory of Computing*, 1985.

[141] Yonghai Gong, Yichuan Li, and Nikolaos M Freris. FedADMM: A robust federated deep learning framework with adaptivity to system heterogeneity. In *IEEE International Conference on Data Engineering*, 2022.

[142] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. `http://www.deeplearningbook.org`.

[143] Sven Gronauer and Klaus Diepold. Multi-agent deep reinforcement learning: A survey. *Artificiell Intelligence Review*, 55(2):895–943, 2022.

[144] Fangda Gu, Heng Chang, Wenwu Zhu, Somayeh Sojoudi, and Laurent El Ghaoui. Implicit graph neural networks. *Advances in Neural Information Processing Systems*, 33:11984–11995, 2020.

[145] Zeli Guan, Yawen Li, Zhe Xue, Yuxin Liu, Hongrui Gao, and Yingxia Shao. Federated graph neural network for cross-graph node classification. In *IEEE International Conference on Cloud Computing and Intelligent Systems*, 2021.

[146] Shurui Gui, Xiner Li, Limei Wang, and Shuiwang Ji. Good: A graph out-of-distribution benchmark. *Advances in Neural Information Processing Systems*, 35:2059–2073, 2022.

[147] Jiayan Guo, Shangyang Li, and Yan Zhang. An information theoretic perspective for heterogeneous subgraph federated learning. In *International Conference on Database Systems for Advanced Applications*, 2023.

[148] Zhuoning Guo, Duanyi Yao, Qiang Yang, and Hao Liu. Hifgl: A hierarchical framework for cross-silo cross-device federated graph learning. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024.

[149] Otkrist Gupta and Ramesh Raskar. Distributed learning of deep neural network over multiple agents. *Journal of Network and Computer Applications*, 116:1–8, 2018.

[150] Suyog Gupta, Wei Zhang, and Fei Wang. Model accuracy and runtime tradeoff in distributed deep learning: A systematic study. In *IEEE International Conference on Data Mining*, 2016.

[151] Benjamin Gutteridge, Xiaowen Dong, Michael M Bronstein, and Francesco Di Giovanni. Drew: Dynamically rewired message passing with delay. In *International Conference on Machine Learning*, 2023.

[152] Mustafa Hajij, Ghada Zamzmi, Theodore Papamarkou, Nina Miolane, Aldo Guzmán-Sáenz, Karthikeyan Natesan Ramamurthy, Tolga Birdal, Tamal K. Dey, Soham Mukherjee, Shreyas N. Samaga, Neal Livesay, Robin Walters, Paul Rosen, and Michael T. Schaub. Topological deep learning: Going beyond graph data. *arXiv:2206.00606*, 2022.

[153] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in Neural Information Processing Systems*, 30, 2017.

[154] William L Hamilton, Rex Ying, and Jure Leskovec. Representation learning on graphs: Methods and applications. *IEEE Data Engineering Bulletin*, 2017.

[155] Zhaoxing Han, Chengyu Hu, Tongyaqi Li, Qingqiang Qi, Peng Tang, and Shanqing Guo. Subgraph-level federated graph neural network for privacy-preserving recommendation with meta-learning. *Neural Networks*, 179:106574, 2024.

[156] Andrew Hard, Kanishka Rao, Rajiv Mathews, Swaroop Ramaswamy, Françoise Beaufays, Sean Augenstein, Hubert Eichner, Chloé Kiddon, and Daniel Ramage. Federated learning for mobile keyboard prediction. Technical report, Google, 2018.

[157] Chaoyang He, Murali Annavaram, and Salman Avestimehr. Group knowledge transfer: Federated learning of large cnns at the edge. *Advances in Neural Information Processing Systems*, 33:14068–14080, 2020.

[158] Chaoyang He, Keshav Balasubramanian, Emir Ceyani, Carl Yang, Han Xie, Lichao Sun, Lifang He, Liangwei Yang, Philip S. Yu, Yu Rong, Peilin Zhao, Junzhou Huang, Murali Annavaram, and Salman Avestimehr. Fedgraphnn: A federated learning benchmark system for graph neural networks. In *Workshop on Distributed and Private Machine Learning*, 2021.

[159] Chaoyang He, Emir Ceyani, Keshav Balasubramanian, Murali Annavaram, and Salman Avestimehr. Spreadgnn: Decentralized multi-task federated learning for graph neural networks on molecular data. In *AAAI Conference on Artificial Intelligence*, 2022.

[160] Guanghui He, Yanli Ren, Jingyuan Jiang, Guorui Feng, and Xinpeng Zhang. Privacy-enhanced federated GNN inference against adversarial example attack. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2024.

[161] Shiwen He, Shaowen Xiong, Yeyu Ou, Jian Zhang, Jiaheng Wang, Yongming Huang, and Yaoxue Zhang. An overview on the application of graph neural networks in wireless networks. *IEEE Open Journal of the Communications Society*, 2:2547–2565, 2021.

[162] Yuanqin He, Yan Kang, Xinyuan Zhao, Jiahuan Luo, Lixin Fan, and Yuxing Han. A hybrid self-supervised learning framework for vertical federated learning. *IEEE Transactions on Big Data*, 2025.

[163] István Hegedűs, Gábor Danner, and Márk Jelasity. Gossip learning as a decentralized alternative to federated learning. In *IFIP International Conference on Distributed Applications and Interoperable Systems*, 2019.

[164] Henrik Hellström, José Mairton B da Silva Jr, Mohammad Mohammadi Amiri, Mingzhe Chen, Viktoria Fodor, H Vincent Poor, and Carlo Fischione. Wireless for machine learning: A survey. *Foundations and Trends® in Signal Processing*, 15(4):290–399, 2022.

[165] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2014.

[166] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109–137, 1983.

[167] Sara Hooker. The hardware lottery. *Communications of the ACM*, 64(12):58–65, 2021.

[168] Charlie Hou, Kiran Koshy Thekumparampil, Giulia Fanti, and Sewoong Oh. Fedchain: Chained algorithms for near-optimal communication cost in federated learning. In *International Conference on Learning Representations*, 2022.

[169] Kevin Hsieh, Amar Phanishayee, Onur Mutlu, and Phillip Gibbons. The non-iid data quagmire of decentralized machine learning. In *International Conference on Machine Learning*, 2020.

[170] Chung-ju Huang, Leye Wang, and Xiao Han. Vertical federated knowledge transfer via representation distillation for healthcare collaboration networks. In *ACM Web Conference*, 2023.

[171] Li Huang, Andrew L Shea, Huining Qian, Aditya Masurkar, Hao Deng, and Dianbo Liu. Patient clustering improves efficiency of federated machine learning to predict mortality and hospital stay time using distributed electronic medical records. *Journal of Biomedical Informatics*, 99:103291, 2019.

[172] Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R Benson. Combining label propagation and simple models out-performs graph neural networks. In *International Conference on Learning Representations*, 2021.

[173] Tiansheng Huang, Shiwei Liu, Li Shen, Fengxiang He, Weiwei Lin, and Dacheng Tao. Achieving personalized federated learning with sparse local models. *arXiv:2201.11380*, 2022.

[174] Yujin Huang, Han Hu, and Chunyang Chen. Robustness of on-device models: Adversarial attack to deep learning models on android apps. In *IEEE/ACM International Conference on Software Engineering*, 2021.

[175] Yutao Huang, Lingyang Chu, Zirui Zhou, Lanjun Wang, Jiangchuan Liu, Jian Pei, and Yong Zhang. Personalized cross-silo federated learning on non-iid data. In *AAAI Conference on Artificial Intelligence*, 2021.

[176] Dzmitry Huba, John Nguyen, Kshitiz Malik, Ruiyu Zhu, Mike Rabbat, Ashkan Yousefpour, Carole-Jean Wu, Hongyuan Zhan, Pavel Ustinov, Harish Srinivas, Kaikai Wang, Anthony Shoumikhin, Jesik Min, and Mani Malek. Papaya: Practical, private, and scalable federated learning. In *Machine Learning and Systems*, 2022.

[177] Ali Jadbabaie, Jie Lin, and A Stephen Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.

[178] Stefanie Jegelka. Theory of graph neural networks: Representation and learning. In *The International Congress of Mathematicians*, 2022.

[179] Márk Jelasity, Alberto Montresor, and Ozalp Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Transactions on Computer Systems*, 23(3):219–252, 2005.

[180] Eunjeong Jeong, Seungeun Oh, Hyesung Kim, Jihong Park, Mehdi Bennis, and Seong-Lyun Kim. Communication-efficient on-device machine learning: Federated distillation and augmentation under non-iid private data. In *NEURIPS Workshop on Machine Learning on the Phone and Other Consumer Devices*, 2018.

[181] Fangcui Jiang and Long Wang. Finite-time information consensus for multi-agent systems with fixed and switching topologies. *Physica D: Nonlinear Phenomena*, 238(16):1550–1560, 2009.

[182] Yihan Jiang, Jakub Konečný, Keith Rush, and Sreeram Kannan. Improving federated learning personalization via model agnostic meta learning. *arXiv:1909.12488*, 2019.

[183] Bjorn Johansson, Tamás Keviczky, Mikael Johansson, and Karl Henrik Johansson. Subgradient methods and consensus algorithms for solving convex optimization problems. In *IEEE Conference on Decision and Control*, 2008.

[184] Peter Kairouz, Brendan McMahan, Shuang Song, Om Thakkar, Abhradeep Thakurta, and Zheng Xu. Practical and private (deep) learning without sampling or shuffling. In *International Conference on Machine Learning*, 2021.

[185] Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D'Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, 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(1–2):1–210, 2021.

[186] Honggu Kang, Seohyeon Cha, Jinwoo Shin, Jongmyeong Lee, and Joonhyuk Kang. NeFL: Nested model scaling for federated learning with system heterogeneous clients. *IEEE Transactions on Mobile Computing*, 24(8):6734–6746, 2025.

[187] Yan Kang, Yuanqin He, Jiahuan Luo, Tao Fan, Yang Liu, and Qiang Yang. Privacy-preserving federated adversarial domain adaptation over feature groups for interpretability. *IEEE Transactions on Big Data*, 10(6):879–890, 2022.

[188] Yan Kang, Yang Liu, and Xinle Liang. Fedcvt: Semi-supervised vertical federated learning with cross-view training. *ACM Transactions on Intelligent Systems and Technology*, 13(4):1–16, 2022.

[189] Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, 2020.

[190] Richard Karp, Christian Schindelhauer, Scott Shenker, and Berthold Vocking. Randomized rumor spreading. In *IEEE Annual Symposium on Foundations of Computer Science*, 2000.

[191] Brian Karrer and Mark EJ Newman. Stochastic blockmodels and community structure in networks. *Physical Review E*, 83(1):016107, 2011.

[192] Shiva Prasad Kasiviswanathan, Homin K Lee, Kobbi Nissim, Sofya Raskhodnikova, and Adam Smith. What can we learn privately? *SIAM Journal on Computing*, 40(3):793–826, 2011.

[193] Anees Kazi, Luca Cosmo, Seyed-Ahmad Ahmadi, Nassir Navab, and Michael M Bronstein. Differentiable graph module (dgm) for graph convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1606–1617, 2022.

[194] Nicolas Keriven, Alberto Bietti, and Samuel Vaiter. Convergence and stability of graph convolutional networks on large random graphs. *Advances in Neural Information Processing Systems*, 33:21512–21523, 2020.

[195] Ahmed Khaled, Konstantin Mishchenko, and Peter Richtárik. Tighter theory for local SGD on identical and heterogeneous data. In *International Conference on Artificial Intelligence and Statistics*, 2020.

[196] Prashant Khanduri, Pranay Sharma, Haibo Yang, Mingyi Hong, Jia Liu, Ketan Rajawat, and Pramod Varshney. Stem: A stochastic two-sided momentum algorithm achieving near-optimal sample and communication complexities for federated learning. *Advances in Neural Information Processing Systems*, 34:6050–6061, 2021.

[197] Geeho Kim, Jinkyu Kim, and Bohyung Han. Communication-efficient federated learning with accelerated client gradient. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

[198] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.

[199] Benjamin Klein, Lior Wolf, and Yehuda Afek. A dynamic convolutional layer for short range weather prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[200] Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *NIPS Workshop on Private Multi-Party Machine Learning*, 2016.

[201] Emmanouil Krasanakis, Symeon Papadopoulos, and Ioannis Kompatsiaris. p2pGNN: A decentralized graph neural network for node classification in peer-to-peer networks. *IEEE Access*, 10:34755–34765, 2022.

[202] Lorenz Kummer, Samir Moustafa, Anatol Ehrlich, Franka Bause, Nikolaus Suess, Wilfried N. Gansterer, and Nils Morten Kriege. Weisfeiler and leman go gambling: Why expressive lottery tickets win. In *International Conference on Machine Learning*, 2025.

[203] Mengyuan Lee, Guanding Yu, and Huaiyu Dai. Decentralized inference with graph neural networks in wireless communication systems. *IEEE Transactions on Mobile Computing*, 22(5):2582–2598, 2021.

[204] Mengyuan Lee, Guanding Yu, Huaiyu Dai, and Geoffrey Ye Li. Graph neural networks meet wireless communications: Motivation, applications, and future directions. *IEEE Wireless Communications*, 29 (5):12–19, 2022.

[205] Mengyuan Lee, Guanding Yu, and Huaiyu Dai. Privacy-preserving decentralized inference with graph neural networks in wireless networks. *IEEE Transactions on Wireless Communications*, 23(1):543–558, 2023.

[206] Runze Lei, Pinghui Wang, Junzhou Zhao, Lin Lan, Jing Tao, Chao Deng, Junlan Feng, Xidian Wang, and Xiaohong Guan. Federated learning over coupled graphs. *IEEE Transactions on Parallel and Distributed Systems*, 34(4):1159–1172, 2023.

[207] Jichao Leng, Zihuai Lin, Ming Ding, Peng Wang, David Smith, and Branka Vucetic. Client scheduling in wireless federated learning based on channel and learning qualities. *IEEE Wireless Communications Letters*, 11(4):732–735, 2022.

[208] Ang Li, Jingwei Sun, Xiao Zeng, Mi Zhang, Hai Li, and Yiran Chen. Fedmask: Joint computation and communication-efficient personalized federated learning via heterogeneous masking. In *ACM Conference on Embedded Networked Sensor Systems*, 2021.

[209] Daliang Li and Junpu Wang. Fedmd: Heterogenous federated learning via model distillation. In *NEURIPS Workshop on Federated Learning for Data Privacy and Confidentiality*, 2019.

[210] Li Li, Yuxi Fan, Mike Tse, and Kuo-Yi Lin. A review of applications in federated learning. *Computers & Industrial Engineering*, 149:106854, 2020.

[211] Li Li, Moming Duan, Duo Liu, Yu Zhang, Ao Ren, Xianzhang Chen, Yujuan Tan, and Chengliang Wang. Fedsae: A novel self-adaptive federated learning framework in heterogeneous systems. In *IEEE International Joint Conference on Neural Networks*, 2021.

[212] Qinbin Li, Bingsheng He, and Dawn Song. Model-contrastive federated learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[213] Qinbin Li, Yiqun Diao, Quan Chen, and Bingsheng He. Federated learning on non-iid data silos: An experimental study. In *IEEE International Conference on Data Engineering*, 2022.

[214] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *Proceedings of Machine Learning and Systems*, 2: 429–450, 2020.

[215] Wenjie Li, Qiaolin Xia, Junfeng Deng, Hao Cheng, Jiangming Liu, Kouying Xue, Yong Cheng, and Shu-Tao Xia. VFed-SSD: Towards practical vertical federated advertising. In *Trustworthy Federated Learning Workshop at the International Joint Conference on Artificial Intelligence*, 2022.

[216] Xiang Li, Kaixuan Huang, Wenhao Yang, Shusen Wang, and Zhihua Zhang. On the convergence of FedAvg on non-iid data. In *International Conference on Learning Representations*, 2020.

[217] Xingyu Li, Zhe Qu, Bo Tang, and Zhuo Lu. Stragglers are not disasters: A hybrid federated learning framework with delayed gradients. In *IEEE International Conference on Machine Learning and Applications*, 2022.

[218] Xiangru Lian, Wei Zhang, Ce Zhang, and Ji Liu. Asynchronous decentralized parallel stochastic gradient descent. In *International Conference on Machine Learning*, 2018.

[219] Gang Liang and Sudarshan S Chawathe. Privacy-preserving inter-database operations. In *International Conference on Intelligence and Security Informatics*, 2004.

[220] Paul Pu Liang, Terrance Liu, Liu Ziyin, Nicholas B Allen, Randy P Auerbach, David Brent, Ruslan Salakhutdinov, and Louis-Philippe Morency. Think locally, act globally: Federated learning with local and global representations. *arXiv:2001.01523*, 2020.

[221] Wenfei Liang, Yanan Zhao, Rui She, Yiming Li, and Wee Peng Tay. Fedsheafhn: Personalized federated learning on graph-structured data. *arXiv:2405.16056*, 2024.

[222] Haiyang Lin, Mingyu Yan, Xiaochun Ye, Dongrui Fan, Shirui Pan, Wenguang Chen, and Yuan Xie. A comprehensive survey on distributed training of graph neural networks. *Proceedings of the IEEE*, 2023.

[223] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. *Advances in Neural Information Processing Systems*, 33:2351–2363, 2020.

[224] Wanyu Lin, Baochun Li, and Cong Wang. Towards private learning on decentralized graphs with local differential privacy. *IEEE Transactions on Information Forensics and Security*, 17:2936–2946, 2022.

[225] Dongzhu Liu, Guangxu Zhu, Jun Zhang, and Kaibin Huang. Wireless data acquisition for edge learning: Importance-aware retransmission. In *IEEE International Workshop on Signal Processing Advances in Wireless Communications*, 2019.

[226] Dongzhu Liu, Guangxu Zhu, Jun Zhang, and Kaibin Huang. Data-importance aware user scheduling for communication-efficient edge machine learning. *IEEE Transactions on Cognitive Communications and Networking*, 7(1):265–278, 2021.

[227] Fengxia Liu, Zhiming Zheng, Yexuan Shi, Yongxin Tong, and Yi Zhang. A survey on federated learning: A perspective from multi-party computation. *Frontiers of Computer Science*, 18(1):181336, 2024.

[228] Ji Liu, Brian DO Anderson, Ming Cao, and A Stephen Morse. Analysis of accelerated gossip algorithms. *Automatica*, 49(4):873–883, 2013.

[229] Ji Liu, Jizhou Huang, Yang Zhou, Xuhong Li, Shilei Ji, Haoyi Xiong, and Dejing Dou. From distributed machine learning to federated learning: A survey. *Knowledge and Information Systems*, 64(4):885–917, 2022.

[230] Rui Liu, Pengwei Xing, Zichao Deng, Anran Li, Cuntai Guan, and Han Yu. Federated graph neural networks: Overview, techniques, and challenges. *IEEE Transactions on Neural Networks and Learning Systems*, 36(3):4279–4295, 2025.

[231] Shikun Liu, Tianchun Li, Yongbin Feng, Nhan Tran, Han Zhao, Qiang Qiu, and Pan Li. Structural re-weighting improves graph domain adaptation. In *International Conference on Machine Learning*, 2023.

[232] Shuchang Liu, Yingqiang Ge, Shuyuan Xu, Yongfeng Zhang, and Amelie Marian. Fairness-aware federated matrix factorization. In *ACM Conference on Recommender Systems*, 2022.

[233] Tao Liu, Peng Li, and Yu Gu. Glint: Decentralized federated graph learning with traffic throttling and flow scheduling. In *IEEE/ACM International Symposium on Quality of Service*, 2021.

[234] Wanchun Liu, Xin Zang, Yonghui Li, and Branka Vucetic. Over-the-air computation systems: Optimization, analysis and scaling laws. *IEEE Transactions on Wireless Communications*, 19(8):5488–5502, 2020.

[235] Yang Liu, Yan Kang, Chaoping Xing, Tianjian Chen, and Qiang Yang. A secure federated transfer learning framework. *IEEE Intelligent Systems*, 35(4):70–82, 2020.

[236] Yang Liu, Yan Kang, Tianyuan Zou, Yanhong Pu, Yuanqin He, Xiaozhou Ye, Ye Ouyang, Ya-Qin Zhang, and Qiang Yang. Vertical federated learning: Concepts, advances, and challenges. *IEEE Transactions on Knowledge and Data Engineering*, 36(7):3615–3634, 2024.

[237] Zhaowei Liu, Yifan Jiang, Jiaojiao Gu, Diantong Liu, Yongchao Song, Anzuo Jiang, and Peiyong Duan. Federated graph neural networks based on multiscale residuals in industrial internet of things. *IEEE Transactions on Computational Social Systems*, 2025.

[238] Zhiwei Liu, Liangwei Yang, Ziwei Fan, Hao Peng, and Philip S Yu. Federated social recommendation with graph neural network. *ACM Transactions on Intelligent Systems and Technology*, 13(4):1–24, 2022.

[239] Ilan Lobel and Asuman Ozdaglar. Convergence analysis of distributed subgradient methods over random networks. In *IEEE Allerton Conference on Communication, Control, and Computing*, 2008.

[240] Guodong Long, Ming Xie, Tao Shen, Tianyi Zhou, Xianzhi Wang, and Jing Jiang. Multi-center federated learning: Clients clustering for better personalization. *World Wide Web*, 26(1):481–500, 2023.

[241] Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph Hellerstein. Graphlab: A new framework for parallel machine learning. In *Conference on Uncertainty in Artificial Intelligence*, 2010.

[242] Linpeng Lu and Ning Ding. Multi-party private set intersection in vertical federated learning. In *IEEE International Conference on Trust, Security and Privacy in Computing and Communications*, 2020.

[243] Ruichen Luo, Sebastian U Stich, Samuel Horváth, and Martin Takáč. Revisiting localSGD and SCAFFOLD: Improved rates and missing analysis. In *International Conference on Artificial Intelligence and Statistics*, 2025.

[244] 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.

[245] Lingjuan Lyu, Han Yu, and Qiang Yang. Threats to federated learning: A survey. *arXiv:2003.02133*, 2020.

[246] Chengyuan Mai, Tianchi Liao, Chuan Chen, and Zibin Zheng. Fgtl: Federated graph transfer learning for node classification. *ACM Transactions on Knowledge Discovery from Data*, 19(1):1–20, 2024.

[247] Yishay Mansour, Mehryar Mohri, Jae Ro, and Ananda Theertha Suresh. Three approaches for personalization with applications to federated learning. *arXiv:2002.10619*, 2020.

[248] Sohir Maskey, Raffaele Paolino, Fabian Jogl, Gitta Kutyniok, and Johannes F Lutzeyer. Graph representational learning: When does more expressivity hurt generalization? *arXiv:2505.11298*, 2025.

[249] Joël Mathys, Florian Grötschla, Kalyan Varma Nadimpalli, and Roger Wattenhofer. Rethinking message passing for algorithmic alignment. In *NEURIPS Workshop on Symmetry and Geometry in Neural Representations*, 2024.

[250] Ryan Mcdonald, Mehryar Mohri, Nathan Silberman, Dan Walker, and Gideon Mann. Efficient large-scale distributed training of conditional maximum entropy models. *Advances in Neural Information Processing Systems*, 22, 2009.

[251] Ryan McDonald, Keith Hall, and Gideon Mann. Distributed training strategies for the structured perceptron. In *Annual Conference of the North American Chapter of the Association for Computational Linguistics*, 2010.

[252] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, 2017.

[253] H Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *International Conference on Learning Representations*, 2018.

[254] Guangxu Mei, Ziyu Guo, Shijun Liu, and Li Pan. Sgnn: A graph neural network based federated learning approach by hiding structure. In *IEEE International Conference on Big Data*, 2019.

[255] Luca Melis, Congzheng Song, Emiliano De Cristofaro, and Vitaly Shmatikov. Exploiting unintended feature leakage in collaborative learning. In *IEEE Symposium on Security and Privacy*, 2019.

[256] Chuiyang Meng, Ming Tang, Mehdi Setayesh, and Vincent WS Wong. Tackling resource allocation for decentralized federated learning: A GNN-based approach. *IEEE Transactions on Mobile Computing*, 2025.

[257] 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 & Data Mining*, 2021.

[258] Fan Mo, Hamed Haddadi, Kleomenis Katevas, Eduard Marin, Diego Perino, and Nicolas Kourtellis. PPFL: Privacy-preserving federated learning with trusted execution environments. In *Annual International Conference on Mobile Systems, Applications, and Services*, 2021.

[259] Bojan Mohar, Y Alavi, G Chartrand, and Ortrud Oellermann. The laplacian spectrum of graphs. *Graph Theory, Combinatorics, and Applications*, 2(871-898):12, 1991.

[260] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.

[261] Luc Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005.

[262] Jose G Moreno-Torres, Troy Raeder, Rocío Alaiz-Rodríguez, Nitesh V Chawla, and Francisco Herrera. A unifying view on dataset shift in classification. *Pattern Recognition*, 45(1):521–530, 2012.

[263] Christopher Morris, Fabrizio Frasca, Nadav Dym, Haggai Maron, İsmail İlkan Ceylan, Ron Levie, Derek Lim, Michael Bronstein, Martin Grohe, and Stefanie Jegelka. Future directions in the theory of graph machine learning. *arXiv:2402.02287*, 2024.

[264] Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. A comprehensive overview of large language models. *ACM Transactions on Intelligent Systems and Technology*, 2023.

[265] Mahmoud Nazzal, Abdallah Khreishah, Joyoung Lee, Shaahin Angizi, Ala Al-Fuqaha, and Mohsen Guizani. Semi-decentralized inference in heterogeneous graph neural networks for traffic demand forecasting: An edge-computing approach. *IEEE Transactions on Vehicular Technology*, 73(12):19400–19416, 2024.

[266] Angelia Nedic and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

[267] Angelia Nedic, Alex Olshevsky, Asuman Ozdaglar, and John N Tsitsiklis. On distributed averaging algorithms and quantization effects. *IEEE Transactions on Automatic Control*, 54(11):2506–2517, 2009.

[268] Angelia Nedic, Asuman Ozdaglar, and Pablo A Parrilo. Constrained consensus and optimization in multi-agent networks. *IEEE Transactions on Automatic Control*, 55(4):922–938, 2010.

[269] Yurii Nesterov. Primal-dual subgradient methods for convex problems. *Mathematical Programming*, 120(1):221–259, 2009.

[270] John Nguyen, Kshitiz Malik, Hongyuan Zhan, Ashkan Yousefpour, Mike Rabbat, Mani Malek, and Dzmitry Huba. Federated learning with buffered asynchronous aggregation. In *International Conference on Artificial Intelligence and Statistics*, 2022.

[271] Xiang Ni, Xiaolong Xu, Lingjuan Lyu, Changhua Meng, and Weiqiang Wang. A vertical federated learning framework for graph convolutional network. *arXiv:2106.11593*, 2021.

[272] Takayuki Nishio and Ryo Yonetani. Client selection for federated learning with heterogeneous resources in mobile edge. In *IEEE International Conference on Communications*, 2019.

[273] Thomas E Nisonger. The "80/20 rule" and core journals. *The Serials Librarian*, 55(1-2):62–84, 2008.

[274] Yongjeong Oh, Yo-Seb Jeon, Mingzhe Chen, and Walid Saad. Fedvqcs: Federated learning via vector quantized compressed sensing. *IEEE Transactions on Wireless Communications*, 23(3):1755–1770, 2024.

[275] Reza Olfati-Saber and Richard M Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520–1533, 2004.

[276] Reza Olfati-Saber, J Alex Fax, and Richard M Murray. Consensus and cooperation in networked multi-agent systems. *IEEE Conference on Decision and Control*, 95(1):215–233, 2007.

[277] Alex Olshevsky and John N Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM Journal on Control and Optimization*, 48(1):33–55, 2009.

[278] Rostyslav Olshevskyi, Zhongyuan Zhao, Kevin Chan, Gunjan Verma, Ananthram Swami, and Santiago Segarra. Fully distributed online training of graph neural networks in networked systems. In *IEEE International Conference on Machine Learning for Communication and Networking*, 2025.

[279] Noa Onoszko, Gustav Karlsson, Olof Mogren, and Edvin Listo Zec. Decentralized federated learning of deep neural networks on non-iid data. In *International Workshop on Federated Learning for User Privacy and Data Confidentiality in Conjunction with ICML*, 2021.

[280] Róbert Ormándi, István Hegedűs, and Márk Jelasity. Gossip learning with linear models on fully distributed data. *Concurrency and Computation: Practice and Experience*, 25(4):556–571, 2013.

[281] Qiying Pan and Yifei Zhu. Fedwalk: Communication efficient federated unsupervised node embedding with differential privacy. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2022.

[282] Qiying Pan, Yifei Zhu, and Lingyang Chu. Lumos: Heterogeneity-aware federated graph learning over decentralized devices. In *IEEE International Conference on Data Engineering*, 2023.

[283] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2009.

[284] Theodore Papamarkou, Tolga Birdal, Michael Bronstein, Gunnar Carlsson, Justin Curry, Yue Gao, Mustafa Hajij, Roland Kwitt, Pietro Liò, Paolo Di Lorenzo, Vasileios Maroulas, Nina Miolane, Farzana Nasrin, Karthikeyan Natesan Ramamurthy, Bastian Rieck, Simone Scardapane, Michael T. Schaub, Petar Veličković, Bei Wang, Yusu Wang, Guo-Wei Wei, and Ghada Zamzmi. Position: Topological deep learning is the new frontier for relational learning. In *International Conference on Machine Learning*, 2024.

[285] Pál András Papp and Roger Wattenhofer. An introduction to graph neural networks from a distributed computing perspective. In *International Conference on Distributed Computing and Internet Technology*, 2022.

[286] Jong-Ik Park and Carlee Joe-Wong. Federated learning with flexible architectures. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2024.

[287] Jong-Ik Park and Carlee Joe-Wong. Fedtlu: Federated learning with targeted layer updates. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2025.

[288] Matthias Paulik, Matt Seigel, Henry Mason, Dominic Telaar, Joris Kluivers, Rogier van Dalen, Chi Wai Lau, Luke Carlson, Filip Granqvist, Chris Vandevelde, Sudeep Agarwal, Julien Freudiger, Andrew Byde, Abhishek Bhowmick, Gaurav Kapoor, Si Beaumont, Áine Cahill, Dominic Hughes, Omid Javidbakht, Fei Dong, Rehan Rishi, and Stanley Hung. Federated evaluation and tuning for on-device personalization: System design & applications. Technical report, Apple, 2021.

[289] Xinjun Pei, Xiaoheng Deng, Shengwei Tian, Jianqing Liu, and Kaiping Xue. Privacy-enhanced graph neural network for decentralized local graphs. *IEEE Transactions on Information Forensics and Security*, 19:1614–1629, 2023.

[290] Hao Peng, Haoran Li, Yangqiu Song, Vincent Zheng, and Jianxin Li. Differentially private federated knowledge graphs embedding. In *ACM International Conference on Information & Knowledge Management*, 2021.

[291] Jingshu Peng, Zhao Chen, Yingxia Shao, Yanyan Shen, Lei Chen, and Jiannong Cao. Sancus: Staleness-aware communication-avoiding full-graph decentralized training in large-scale graph neural networks. *Proceedings of the VLDB Endowment*, 15(9):1937–1950, 2022.

[292] Liang Peng, Nan Wang, Nicha Dvornek, Xiaofeng Zhu, and Xiaoxiao Li. Fedni: Federated graph learning with network inpainting for population-based disease prediction. *IEEE Transactions on Medical Imaging*, 42(7):2032–2043, 2022.

[293] Le Trieu Phong, Yoshinori Aono, Takuya Hayashi, Lihua Wang, and Shiho Moriai. Privacy-preserving deep learning via additively homomorphic encryption. *IEEE Transactions on Information Forensics and Security*, 13(5):1333–1345, 2017.

[294] Vasyl Pihur, Aleksandra Korolova, Frederick Liu, Subhash Sankuratripati, Moti Yung, Dachuan Huang, and Ruogu Zeng. Differentially-private "draw and discard" machine learning: Training distributed model from enormous crowds. In *International Symposium on Cyber Security, Cryptology, and Machine Learning*, 2022.

[295] Benny Pinkas, Thomas Schneider, and Michael Zohner. Scalable private set intersection based on ot extension. *ACM Transactions on Privacy and Security*, 21(2):1–35, 2018.

[296] Oleg Platonov, Denis Kuznedelev, Artem Babenko, and Liudmila Prokhorenkova. Characterizing graph datasets for node classification: Homophily-heterophily dichotomy and beyond. *Advances in Neural Information Processing Systems*, 36:523–548, 2023.

[297] Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In *International Conference on Learning Representations*, 2023.

[298] Daniel Pérez Herrera, Zheng Chen, and Erik G. Larsson. Faster convergence with less communication: Broadcast-based subgraph sampling for decentralized learning over wireless networks. *IEEE Open Journal of the Communications Society*, 2025.

[299] Chendi Qian, Gaurav Rattan, Floris Geerts, Mathias Niepert, and Christopher Morris. Ordered subgraph aggregation networks. *Advances in Neural Information Processing Systems*, 35:21030–21045, 2022.

[300] Yeqing Qiu, Chenyu Huang, Jianzong Wang, Zhangcheng Huang, and Jing Xiao. A privacy-preserving subgraph-level federated graph neural network via differential privacy. In *International Conference on Knowledge Science, Engineering and Management*, 2022.

[301] Liang Qu, Ningzhi Tang, Ruiqi Zheng, Quoc Viet Hung Nguyen, Zi Huang, Yuhui Shi, and Hongzhi Yin. Semi-decentralized federated ego graph learning for recommendation. In *ACM Web Conference*, 2023.

[302] Joaquin Quiñonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset Shift in Machine Learning*. MIT Press, 2022.

[303] S Sundhar Ram, A Nedić, and Venugopal V Veeravalli. Asynchronous gossip algorithms for stochastic optimization. In *IEEE Conference on Decision and Control*, 2009.

[304] S Sundhar Ram, A Nedić, and Venugopal V Veeravalli. Incremental stochastic subgradient algorithms for convex optimization. *SIAM Journal on Optimization*, 20(2):691–717, 2009.

[305] Swaroop Ramaswamy, Om Thakkar, Rajiv Mathews, Galen Andrew, H. Brendan McMahan, and Françoise Beaufays. Training production language models without memorizing user data. Technical report, Google, 2020.

[306] Saeed Razavikia, José Mairton Barros da Silva, and Carlo Fischione. ChannelComp: A general method for computation by communications. *IEEE Transactions on Communications*, 6:1497–1511, 2024.

[307] Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H. Brendan McMahan. Adaptive federated optimization. In *International Conference on Learning Representations*, 2021.

[308] Jinke Ren, Yinghui He, Dingzhu Wen, Guanding Yu, Kaibin Huang, and Dongning Guo. Scheduling for cellular federated edge learning with importance and channel awareness. *IEEE Transactions on Wireless Communications*, 19(11):7690–7703, 2020.

[309] Zhenghang Ren, Liu Yang, and Kai Chen. Improving availability of vertical federated learning: Relaxing inference on non-overlapping data. *ACM Transactions on Intelligent Systems and Technology*, 13(4):1–20, 2022.

[310] Mónica Ribero, Haris Vikalo, and Gustavo De Veciana. Federated learning under intermittent client availability and time-varying communication constraints. *IEEE Journal of Selected Topics in Signal Processing*, 17(1):98–111, 2022.

[311] Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 400–407, 1951.

[312] Angelo Rodio, Francescomaria Faticanti, Othmane Marfoq, Giovanni Neglia, and Emilio Leonardi. Federated learning under heterogeneous and correlated client availability. In *IEEE Infocom Conference on Computer Communications*, 2023.

[313] Frank Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386, 1958.

[314] Yichen Ruan, Xiaoxi Zhang, Shu-Che Liang, and Carlee Joe-Wong. Towards flexible device participation in federated learning. In *International Conference on Artificial Intelligence and Statistics*, 2021.

[315] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *Nature*, 323(6088):533–536, 1986.

[316] Fahad Sabah, Yuwen Chen, Zhen Yang, Muhammad Azam, Nadeem Ahmad, and Raheem Sarwar. Model optimization techniques in personalized federated learning: A survey. *Expert Systems with Applications*, 243:122874, 2024.

[317] R Olfati Saber and Richard M Murray. Consensus protocols for networks of dynamic agents. In *IEEE American Control Conference*, 2003.

[318] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Transactions on Neural Networks and Learning Systems*, 32(8):3710–3722, 2020.

[319] Simone Scardapane, Indro Spinelli, and Paolo Di Lorenzo. Distributed training of graph convolutional networks. *IEEE Transactions on Signal and Information Processing over Networks*, 7:87–100, 2020.

[320] Aviv Shamsian, Aviv Navon, Ethan Fetaya, and Gal Chechik. Personalized federated learning using hypernetworks. In *International Conference on Machine Learning*, 2021.

[321] Yingxia Shao, Hongzheng Li, Xizhi Gu, Hongbo Yin, Yawen Li, Xupeng Miao, Wentao Zhang, Bin Cui, and Lei Chen. Distributed graph neural network training: A survey. *ACM Computing Surveys*, 56(8):1–39, 2024.

[322] Wei Shen, Mang Ye, Wei Yu, and Pong C Yuen. Build yourself before collaboration: Vertical federated learning with limited aligned samples. *IEEE Transactions on Mobile Computing*, 24(7):6503–6516, 2025.

[323] Xiao Shen, Quanyu Dai, Fu-lai Chung, Wei Lu, and Kup-Sze Choi. Adversarial deep network embedding for cross-network node classification. In *AAAI Conference on Artificial Intelligence*, 2020.

[324] Zitong Shi, Guancheng Wan, Wenke Huang, Guibin Zhang, He Li, Carl Yang, and Mang Ye. Eagles: Towards effective, efficient, and economical federated graph learning via unified sparsification. In *International Conference on Machine Learning*, 2025.

[325] 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*, 69: 500–514, 2021.

[326] Neta Shoham, Tomer Avidor, Aviv Keren, Nadav Israel, Daniel Benditkis, Liron Mor-Yosef, and Itai Zeitak. Overcoming forgetting in federated learning on non-iid data. In *NEURIPS Workshop on Federated Learning for Data Privacy and Confidentiality*, 2019.

[327] Ilia Shumailov, Daniel Ramage, Sarah Meiklejohn, Peter Kairouz, Florian Hartmann, Borja Balle, and Eugene Bagdasarian. Trusted machine learning models unlock private inference for problems currently infeasible with cryptography. Technical report, Google, 2025.

[328] Adam Smith, Abhradeep Thakurta, and Jalaj Upadhyay. Is interaction necessary for distributed private learning? In *IEEE Symposium on Security and Privacy*, 2017.

[329] Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S Talwalkar. Federated multi-task learning. *Advances in Neural Information Processing Systems*, 30, 2017.

[330] John Snow. On the mode of communication of cholera. *Edinburgh medical journal*, 1855.

[331] Olga Solodova, Nick Richardson, Deniz Oktay, and Ryan P Adams. Graph neural networks gone hogwild. In *The International Conference on Learning Representations*, 2025.

[332] Shuang Song, Kamalika Chaudhuri, and Anand D Sarwate. Stochastic gradient descent with differentially private updates. In *IEEE Global Conference on Signal and Information Processing*, 2013.

[333] Joshua Southern, Yam Eitan, Guy Bar-Shalom, Michael M. Bronstein, Haggai Maron, and Fabrizio Frasca. Balancing efficiency and expressiveness: Subgraph GNNs with walk-based centrality. In *International Conference on Machine Learning*, 2025.

[334] Sebastian U. Stich. Local SGD converges fast and communicates little. In *International Conference on Learning Representations*, 2019.

[335] Zhenchao Sun, Hongzhi Yin, Hongxu Chen, Tong Chen, Lizhen Cui, and Fan Yang. Disease prediction via graph neural networks. *IEEE Journal of Biomedical and Health Informatics*, 25(3):818–826, 2020.

[336] Zhenyu Sun, Ziyang Zhang, Zheng Xu, Gauri Joshi, Pranay Sharma, and Ermin Wei. Debiasing federated learning with correlated client participation. In *International Conference on Learning Representations*, 2025.

[337] Canh T Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with Moreau envelopes. *Advances in Neural Information Processing Systems*, 33:21394–21405, 2020.

[338] Kunal Talwar, Shan Wang, Audra McMillan, Vojta Jina, Vitaly Feldman, Pansy Bansal, Bailey Basile, Aine Cahill, Yi Sheng Chan, Mike Chatzidakis, Junye Chen, Oliver Chick, Mona Chitnis, Suman Ganta, Yusuf Goren, Filip Granqvist, Kristine Guo, Frederic Jacobs, Omid Javidbakht, Albert Liu, Richard Low, Dan Mascenik, Steve Myers, David Park, Wonhee Park, Gianni Parsa, Tommy Pauly, Christian Priebe, Rehan Rishi, Guy Rothblum, Michael Scaria, Linmao Song, Congzheng Song, Karl Tarbe, Sebastian Vogt, Luke Winstrom, and Shundong Zhou. Samplable anonymous aggregation for private federated data analysis. In *ACM SIGSAC Conference on Computer and Communications Security*, 2024.

[339] Alysa Ziying Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

[340] Zihan Tan, Guancheng Wan, Wenke Huang, He Li, Guibin Zhang, Carl Yang, and Mang Ye. Fedspa: Generalizable federated graph learning under homophily heterogeneity. In *Computer Vision and Pattern Recognition Conference*, 2025.

[341] Yijun Tian, Shichao Pei, Xiangliang Zhang, Chuxu Zhang, and Nitesh V Chawla. Knowledge distillation on graphs: A survey. *ACM Computing Surveys*, 57(8):1–16, 2025.

[342] Jake Topping, Francesco Di Giovanni, Benjamin Paul Chamberlain, Xiaowen Dong, and Michael M. Bronstein. Understanding over-squashing and bottlenecks on graphs via curvature. In *International Conference on Learning Representations*, 2022.

[343] John Tsitsiklis, Dimitri Bertsekas, and Michael Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.

[344] Saeed Vahidian, Mahdi Morafah, and Bill Lin. Personalized federated learning by structured and unstructured pruning under data heterogeneity. In *IEEE International Conference on Distributed Computing Systems Workshops*, 2021.

[345] Leslie G Valiant. A theory of the learnable. *Communications of the ACM*, 27(11):1134–1142, 1984.

[346] Paul Vanhaesebrouck, Aurélien Bellet, and Marc Tommasi. Decentralized collaborative learning of personalized models over networks. In *Nternational Conference on Artificial Intelligence and Statistics*, 2017.

[347] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018.

[348] Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. In *International Conference on Learning Representations*, 2019.

[349] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. Split learning for health: Distributed deep learning without sharing raw patient data. In *ICLR Workshop on AI for Social Good*, 2019.

[350] Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226, 1995.

[351] Binghui Wang, Ang Li, Meng Pang, Hai Li, and Yiran Chen. Graphfl: A federated learning framework for semi-supervised node classification on graphs. In *IEEE International Conference on Data Mining*, 2022.

[352] Hao Wang, Zakhary Kaplan, Di Niu, and Baochun Li. Optimizing federated learning on non-iid data with reinforcement learning. In *IEEE Infocom Conference on Computer Communications*, 2020.

[353] Jialei Wang, Mladen Kolar, and Nathan Srerbo. Distributed multi-task learning. In *International Conference on Artificial Intelligence and Statistics*, 2016.

[354] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H. Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. In *Advances in Neural Information Processing Systems*, 2020.

[355] Kaibin Wang, Qiang He, Feifei Chen, Chunyang Chen, Faliang Huang, Hai Jin, and Yun Yang. Flexifed: Personalized federated learning for edge clients with heterogeneous model architectures. In *ACM Web Conference*, 2023.

[356] Ning Wang, Xiaokui Xiao, Yin Yang, Jun Zhao, Siu Cheung Hui, Hyejin Shin, Junbum Shin, and Ge Yu. Collecting and analyzing multidimensional data with local differential privacy. In *IEEE International Conference on Data Engineering*, 2019.

[357] Shilong Wang, Jianchun Liu, Hongli Xu, Chenxia Tang, Qianpiao Ma, and Liusheng Huang. Towards communication-efficient decentralized federated graph learning over non-iid data. *arXiv:2509.08409*, 2025.

[358] Shiqiang Wang and Mingyue Ji. A unified analysis of federated learning with arbitrary client participation. *Advances in Neural Information Processing Systems*, 35:19124–19137, 2022.

[359] Shiqiang Wang and Mingyue Ji. A lightweight method for tackling unknown participation statistics in federated averaging. In *International Conference on Learning Representations*, 2024.

[360] Shuai Wang, Yik-Chung Wu, Minghua Xia, Rui Wang, and H. Vincent Poor. Machine intelligence at the edge with learning centric power allocation. *IEEE Transactions on Wireless Communications*, 21 (10):7926–7941, 2020.

[361] Xinghan Wang, Anusha Lalitha, Tara Javidi, and Farinaz Koushanfar. Peer-to-peer variational federated learning over arbitrary graphs. *IEEE Journal on Selected Areas in Information Theory*, 30(2): 129–150, 2022.

[362] Yiwei Wang, Wei Wang, Yuxuan Liang, Yujun Cai, and Bryan Hooi. Mixup for node and graph classification. In *Web Conference*, 2021.

[363] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics*, 38(5):1–12, 2019.

[364] Yuzhuo Wang, Hongzhi Wang, Xianglong Liu, and Yu Yan. GFedKG: GNN-based federated embedding model for knowledge graph completion. *Knowledge-Based Systems*, 301:112290, 2024.

[365] Zehong Wang, Zheyuan Liu, Tianyi Ma, Jiazheng Li, Zheyuan Zhang, Xingbo Fu, Yiyang Li, Zhengqing Yuan, Wei Song, Yijun Ma, Qingkai Zeng, Xiusi Chen, Jianan Zhao, Jundong Li, Meng Jiang, Pietro Lio, Nitesh Chawla, Chuxu Zhang, and Yanfang Ye. Graph foundation models: A comprehensive survey. *arXiv:2505.15116*, 2025.

[366] Zheng Wang, Xiaoliang Fan, Jianzhong Qi, Haibing Jin, Peizhen Yang, Siqi Shen, and Cheng Wang. Fedgs: Federated graph-based sampling with arbitrary client availability. In *AAAI Conference on Artificial Intelligence*, 2023.

[367] Zhiyang Wang, Mark Eisen, and Alejandro Ribeiro. Learning decentralized wireless resource allocations with graph neural networks. *IEEE Transactions on Signal Processing*, 70:1850–1863, 2022.

[368] Ermin Wei and Asuman Ozdaglar. Distributed alternating direction method of multipliers. In *IEEE Conference on Decision and Control*, 2012.

[369] Jinliang Wei, Wei Dai, Aurick Qiao, Qirong Ho, Henggang Cui, Gregory R Ganger, Phillip B Gibbons, Garth A Gibson, and Eric P Xing. Managed communication and consistency for fast data-parallel iterative analytics. In *ACM Symposium on Cloud Computing*, 2015.

[370] Wanli Wen, Zihan Chen, Howard H Yang, Wenchao Xia, and Tony QS Quek. Joint scheduling and resource allocation for hierarchical federated edge learning. *IEEE Transactions on Wireless Communications*, 21(8):5857–5872, 2022.

[371] Mariel Werner, Lie He, Michael Jordan, Martin Jaggi, and Sai Praneeth Karimireddy. Provably personalized and robust federated learning. *Transactions on Machine Learning Research*, 2024.

[372] Chuhan Wu, Fangzhao Wu, Yang Cao, Yongfeng Huang, and Xing Xie. Fedgnn: Federated graph neural network for privacy-preserving recommendation. *arXiv:2102.04925*, 2021.

[373] Chuhan Wu, Fangzhao Wu, Lingjuan Lyu, Tao Qi, Yongfeng Huang, and Xing Xie. A federated graph neural network framework for privacy-preserving personalization. *Nature Communications*, 13(1):3091, 2022.

[374] Fan Wu, Yunhui Long, Ce Zhang, and Bo Li. Linkteller: Recovering private edges from graph neural networks via influence analysis. In *IEEE Symposium on Security and Privacy*, 2022.

[375] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph convolutional networks. In *International Conference on Machine Learning*, 2019.

[376] Qiong Wu, Xu Chen, Zhi Zhou, and Junshan Zhang. Fedhome: Cloud-edge based personalized federated learning for in-home health monitoring. *IEEE Transactions on Mobile Computing*, 21(8):2818–2832, 2020.

[377] Qiong Wu, Xu Chen, Tao Ouyang, Zhi Zhou, Xiaoxi Zhang, Shusen Yang, and Junshan Zhang. Hiflash: Communication-efficient hierarchical federated learning with adaptive staleness control and heterogeneity-aware client-edge association. *IEEE Transactions on Parallel and Distributed Systems*, 34(5):1560–1579, 2023.

[378] Ying-Xin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. Discovering invariant rationales for graph neural networks. *arXiv:2201.12872*, 2022.

[379] Zhenqin Wu, Bharath Ramsundar, Evan N Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S Pappu, Karl Leswing, and Vijay Pande. Moleculenet: A benchmark for molecular machine learning. *Chemical Science*, 9(2):513–530, 2018.

[380] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1):4–24, 2020.

[381] Ming Xiang, Stratis Ioannidis, Edmund Yeh, Carlee Joe-Wong, and Lili Su. Efficient federated learning against heterogeneous and non-stationary client unavailability. *Advances in Neural Information Processing Systems*, 37:104281–104328, 2024.

[382] Cong Xie, Oluwasanmi Koyejo, and Indranil Gupta. Asynchronous federated optimization. In *Annual Workshop on Optimization for Machine Learning*, 2020.

[383] Han Xie, Jing Ma, Li Xiong, and Carl Yang. Federated graph classification over non-iid graphs. *Advances in Neural Information Processing Systems*, 34:18839–18852, 2021.

[384] Pengtao Xie, Jin Kyu Kim, Yi Zhou, Qirong Ho, Abhimanu Kumar, Yaoliang Yu, and Eric Xing. Distributed machine learning via sufficient factor broadcasting. *arXiv:1511.08486*, 2015.

[385] Eric P Xing, Qirong Ho, Pengtao Xie, and Dai Wei. Strategies and principles of distributed machine learning on big data. *Engineering*, 2(2):179–195, 2016.

[386] Zheng Xu, Yanxiang Zhang, Galen Andrew, Christopher A Choquette-Choo, Peter Kairouz, H Brendan McMahan, Jesse Rosenstock, and Yuanbo Zhang. Federated learning of gboard language models with differential privacy. Technical report, Google, 2023.

[387] Bo Yan, Yang Cao, Haoyu Wang, Wenchuan Yang, Junping Du, and Chuan Shi. Federated heterogeneous graph neural network for privacy-preserving recommendation. In *ACM Web Conference*, 2024.

[388] Howard H Yang, Zuozhu Liu, Tony QS Quek, and H Vincent Poor. Scheduling policies for federated learning in wireless networks. *IEEE Transactions on Communications*, 68(1):317–333, 2019.

[389] Miao Yang, Ximin Wang, Hongbin Zhu, Haifeng Wang, and Hua Qian. Federated learning with class imbalance reduction. In *IEEE European Signal Processing Conference*, 2021.

[390] Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 10(2):1–19, 2019.

[391] Timothy Yang, Galen Andrew, Hubert Eichner, Haicheng Sun, Wei Li, Nicholas Kong, Daniel Ramage, and Françoise Beaufays. Applied federated learning: Improving Google keyboard query suggestions. Technical report, Google, 2018.

[392] Yitao Yang, Xiucai Ye, and Tetsuya Sakurai. Multi-view federated learning with data collaboration. In *International Conference on Machine Learning*, 2022.

[393] Xin Yao and Lifeng Sun. Continual local training for better initialization of federated models. In *IEEE International Conference on Image Processing*, 2020.

[394] Yuhang Yao, Weizhao Jin, Srivatsan Ravi, and Carlee Joe-Wong. Fedgcn: Convergence-communication tradeoffs in federated training of graph convolutional networks. *Advances in Neural Information Processing Systems*, 36:79748–79760, 2023.

[395] Gilad Yehudai, Ethan Fetaya, Eli Meirom, Gal Chechik, and Haggai Maron. From local structures to size generalization in graph neural networks. In *International Conference on Machine Learning*, 2021.

[396] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems*, 33: 5812–5823, 2020.

[397] Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel restarted SGD with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *AAAI Conference on Artificial Intelligence*, 2019.

[398] Hui Yu, Yu Zhang, Ligang He, Donghao He, Qikun Li, Jin Zhao, Xiaofei Liao, Hai Jin, Lin Gu, and Haikun Liu. CDA-GNN: A chain-driven accelerator for efficient asynchronous graph neural network. In *ACM/IEEE Design Automation Conference*, 2024.

[399] Edvin Listo Zec, Ebba Ekblom, Martin Willbo, Olof Mogren, and Sarunas Girdzijauskas. Decentralized adaptive clustering of deep nets is beneficial for client collaboration. In *International Workshop on Trustworthy Federated Learning*, 2022.

[400] Edvin Listo Zec, Johan Östman, Olof Mogren, and Daniel Gillblad. Efficient node selection in private personalized decentralized learning. In *Northern Lights Deep Learning Conference*, 2024.

[401] Liekang Zeng, Chongyu Yang, Peng Huang, Zhi Zhou, Shuai Yu, and Xu Chen. GNN at the edge: Cost-efficient graph neural network processing over distributed edge servers. *IEEE Journal on Selected Areas in Communications*, 41(3):720–739, 2022.

[402] Chen Zhang, Yu Xie, Hang Bai, Bin Yu, Weihong Li, and Yuan Gao. A survey on federated learning. *Knowledge-Based Systems*, 216:106775, 2021.

[403] Chenhan Zhang, Shuyu Zhang, JQ James, and Shui Yu. FASTGNN: A topological information protected federated learning approach for traffic speed forecasting. *IEEE Transactions on Industrial Informatics*, 17(12):8464–8474, 2021.

[404] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. Mixup: Beyond empirical risk minimization. In *International Conference on Learning Representations*, 2018.

[405] Huanding Zhang, Tao Shen, Fei Wu, Mingyang Yin, Hongxia Yang, and Chao Wu. Federated graph learning–a position paper. *arXiv:2105.11099*, 2021.

[406] Kai Zhang, Yu Wang, Hongyi Wang, Lifu Huang, Carl Yang, and Lichao Sun. Efficient federated learning on knowledge graphs via privacy-preserving relation embedding aggregation. In *ACL Workshop on Federated Learning for Natural Language Processing*, 2022.

[407] Ke Zhang, Carl Yang, Xiaoxiao Li, Lichao Sun, and Siu Ming Yiu. Subgraph federated learning with missing neighbor generation. *Advances in Neural Information Processing Systems*, 34:6671–6682, 2021.

[408] Kexin Zhang, Shuhan Liu, Song Wang, Weili Shi, Chen Chen, Pan Li, Sheng Li, Jundong Li, and Kaize Ding. A survey of deep graph learning under distribution shifts: From graph out-of-distribution generalization to adaptation. *arXiv:2410.19265*, 2024.

[409] Taolin Zhang, Chengyuan Mai, Yaomin Chang, Chuan Chen, Lin Shu, and Zibin Zheng. Fedego: Privacy-preserving personalized federated graph learning with ego-graphs. *ACM Transactions on Knowledge Discovery from Data*, 18(2):1–27, 2023.

[410] Yuchen Zhang, Martin J Wainwright, and John C Duchi. Communication-efficient algorithms for statistical optimization. *Advances in Neural Information Processing Systems*, 25, 2012.

[411] Zezhong Zhang, Guangxu Zhu, Rui Wang, Vincent K. N. Lau, and Kaibin Huang. Turning channel noise into an accelerator for over-the-air principal component analysis. *IEEE Transactions on Wireless Communications*, 2022.

[412] Zhenxiao Zhang, Zhidong Gao, Yuanxiong Guo, and Yanmin Gong. Scalable and low-latency federated learning with cooperative mobile edge networking. *IEEE Transactions on Mobile Computing*, 23(1): 812–822, 2022.

[413] Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv:1806.00582*, 2018.

[414] Xiaolin Zheng, Zhongyu Wang, Chaochao Chen, Jiashu Qian, and Yao Yang. Decentralized graph neural network for privacy-preserving recommendation. In *ACM International Conference on Information and Knowledge Management*, 2023.

[415] Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. *Advances in Neural Information Processing Systems*, 16, 2003.

[416] Lifeng Zhou, Vishnu D Sharma, Qingbiao Li, Amanda Prorok, Alejandro Ribeiro, Pratap Tokekar, and Vijay Kumar. Graph neural networks for decentralized multi-robot submodular action selection. *arXiv:2105.08601*, 2021.

[417] Shenglong Zhou and Geoffrey Ye Li. Federated learning via inexact ADMM. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):9699–9708, 2023.

[418] Xiaokang Zhou, Xiaozhou Ye, I Kevin, Kai Wang, Wei Liang, Nirmal Kumar C Nair, Shohei Shimizu, Zheng Yan, and Qun Jin. Hierarchical federated learning with social context clustering-based participant selection for internet of medical things applications. *IEEE Transactions on Computational Social Systems*, 10(4):1742–1751, 2023.

[419] Zhou Zhou, Youliang Tian, and Changgen Peng. Privacy-preserving federated learning framework with general aggregation and multiparty entity matching. *Wireless Communications and Mobile Computing*, 2021(1):6692061, 2021.

[420] Hangyu Zhu, Jinjin Xu, Shiqing Liu, and Yaochu Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 465:371–390, 2021.

[421] Jiong Zhu, Yujun Yan, Lingxiao Zhao, Mark Heimann, Leman Akoglu, and Danai Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in Neural Information Processing Systems*, 33:7793–7804, 2020.

[422] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. In *Advances in Neural Information Processing Systems*, 2019.

[423] Xiaojin Zhu and Zoubin Ghahramani. Learning from labeled and unlabeled data with label propagation. Technical report, Carnegie Mellon University, 2002.

[424] Xiaojin Jerry Zhu. Semi-supervised learning literature survey. *CS Technical Reports University of Wisconsin-Madision*, 2005.

[425] Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In *International Conference on Machine Learning*, 2021.

[426] Martin Zinkevich, John Langford, and Alex Smola. Slow learners are fast. *Advances in Neural Information Processing Systems*, 22, 2009.

[427] Martin Zinkevich, Markus Weimer, Lihong Li, and Alex Smola. Parallelized stochastic gradient descent. *Advances in Neural Information Processing Systems*, 23, 2010.

[428] Daniel Zügner, Oliver Borchert, Amir Akbarnejad, and Stephan Günnemann. Adversarial attacks on graph neural networks: Perturbations and their patterns. *ACM Transactions on Knowledge Discovery from Data*, 14(5):1–31, 2020.

A Notations and Glossary

Glossary

In this work, we use several terms whose definitions are either subjective or inconsistently used in the literature; we clarify them here for consistency and clarity. These definitions are limited to the present work.

Agent An autonomous entity with local data, computational capabilities, and the ability to communicate with other agents.

Server A central coordinating unit that supervises a group of agents, possessing computational capabilities and the ability to communicate with agents.

Distributed A property of a system in which some or all components (e.g., data, computation, or control) are spread across multiple agents.

Data parallelism A property of a distributed system where data collected centrally is distributed across multiple agents.

Model parallelism A property of a distributed system where parts of the model maintained centrally are distributed across multiple agents.

Collaborative (system) A distributed system of agents (either federated or decentralized) where agents cooperate to achieve shared objectives, though their willingness to collaborate may be limited by local constraints or incentives such as but not limited to local objective, computational and communication capabilities, data availability, or privacy requirements.

Federated (system) A distributed system of agents that are federated under the supervision of a central server coordinates.

Decentralized (system) A distributed system of agents where all agents operate without a central server.

Local Describes a property or object that pertains to a single agent (e.g., local data, local model).

Global Describes a property or object shared across or computed from all or many agents. It does not imply that each agent has a full view of this property (e.g., global graph).

Topology The discrete structure describing the communication or computational relationships between agents, typically modeled as a graph (see Hajij et al. (152) for a formal equivalence between topology and graph).

Sample A single data point drawn from a dataset.

Feature An observed attribute or measurable property of a sample.

Label The target property or variable of a sample that we aim to infer from its features.

Training The process where models learn patterns from data by optimizing parameters to minimize prediction errors on the training dataset.

Inference The process by which the model predicts the label of a sample based on its features using learned parameters.

Learning The comprehensive framework encompasses various methodologies, algorithms, and theoretical foundations for training models to perform inference tasks effectively.

Euclidean Pertaining to data where each sample is represented as a feature vector in $\mathbb{R}^n$ equipped with the standard Euclidean inner product.

Non-Euclidean Pertaining to data with inherent relatonal structure that cannot be faithfully represented in $\mathbb{R}^n$ with the Euclidean inner product. Graphs are the most commonly used example of discrete non-Euclidean structures.

Notations

Agents and Data	
$\mathcal{D}_k$, $\mathcal{D}$	Local and global data distributions
$\mathcal{I}_k$, $\mathcal{I}$, i	Local and global index sets of samples, and sample index
$\mathcal{K}$, K, k, k^*, s, t	Set of agents, number of agents, agent index, active agent, source and target agent
n_k, n	Number of samples in local and global datasets
Ω	Similarity matrix between agents
Π, π_i	State of a Markov chain and its i-th element
S_k, S	Local and global datasets
$W \in \mathbb{R}^{K \times K}$	Communication topology matrix between agents
Features, Labels, Representations	
d_x, d_y, d_z	Dimensions of feature vector, label space, and embedding space
x, $\mathcal{X}$	Feature sample and feature domain
x^i, x_k^i	Feature of the i-th sample, global and local
x_k, $\mathcal{X}_k$	Local fraction of feature x in local domain $\mathcal{X}_k$
y, $\mathcal{Y}$	Label and label domain
y^i, $\hat{y}^i$	Label of the i-th sample and its inferred value
z, $\mathcal{Z}$	Representation of sample (x, y) and representation domain
z_k, $\mathcal{Z}_k$	Local fraction of global representation z in local domain $\mathcal{Z}_k$
Graphs and Topology	
A	Adjacency matrix of $\mathcal{G}$
$\mathcal{E}_k$, $\mathcal{E}$	Local and global sets of edges
$\mathcal{E}_{cut,k}$	Local set of cut edges
$G = (A, X)$	Graph data sample in matrix form
G_k, G_k^i	Local subgraph and i-th local graph sample
G^v, G_k^v	Global and local L-hop neighborhoods of node v
$\mathcal{G}_k$, $\mathcal{G}$	Local and global graph structures (homogeneous)
$\tilde{G}_k$, $\tilde{\mathcal{G}}_k$	Local subgraph data and structure (from G_k and $\mathcal{G}_k$), respectively, expanded to L-hop neighborhoods
$\mathcal{H}_{\mathcal{G}}$	Homogeneous projection of heterogeneous graph
N_k, N	Number of vertices in local and global vertex sets
$\mathcal{O}_k$, $\mathcal{O}$	Local and global sets of object types
$\mathcal{R}_k$, $\mathcal{R}$	Local and global sets of relation types
ρ_k, ρ	Local and global relation type mappings
$\mathcal{T}^v$	Computational tree for target node v
τ_k, τ	Local and global node type mappings
$\mathcal{V}_k$, $\mathcal{V}$, v	Local and global sets of vertices (nodes), and vertex index
$\mathcal{V}_{f,k}$, $\mathcal{V}_{p,k}$, $\mathcal{V}_{cut,k}$	Sets of foreign nodes, private nodes, and vertex cut
Node-Level Quantities	
$h_v^{(l)}$, $h_{v,k}^{(l)}$	Global and local embeddings of node v at layer l
X	Feature matrix of the graph
x_v, $x_{v,k}$	Global and local features of node v
z_v, $z_{v,k}$	Global and local node representation of node v
Models and Learning	
E, e	Number of local iterations and local iteration index
F_k, $\mathcal{F}_k$	Local model and its function space
F, $\mathcal{F}$	Global model and its function space
F_θ, $F_{\theta,\psi}$	Model F parametrized by θ and by (θ, ψ)
$f_{Mes}^{(l)}$, $f_{Agg}^{(l)}$, $f_{Up}^{(l)}$	Message, aggregation, and update functions at layer l
f_ψ	Predictor parametrized by ψ
$\mathcal{H}$	Hypernetwork
ℓ	Pairwise loss
L, l	Number of layers and layer index
$\mathcal{L}_k$, $\mathcal{L}$	Local and global population losses
$\hat{\mathcal{L}}_k$, $\hat{\mathcal{L}}$	Local and global empirical losses
ϕ_k, ϕ	Local and global encoders
$\phi_{k,\theta}$, ϕ_θ	Local and global encoders parametrized by θ
ϕ_{node}, ϕ_{edge}, ϕ_{graph}	Node, edge, and graph encoder functions
ψ, Ψ	Predictor parameters and parameter space
T, t	Number of global iterations and global iteration index
θ, Θ	Model parameters and parameter space

Table 3: Notations used in this paper, organized by comprehensive categories and alphabetical order.