Federated Graph Machine Learning: A Survey of Concepts, Techniques, and Applications

Xingbo Fu, Binchi Zhang, Yushun Dong, Chen Chen, Jundong Li
University of Virginia
{xf3av, epb6gw, yd6eb, zrh6du, jundong}@virginia.edu

ABSTRACT

Graph machine learning has gained great attention in both academia and industry recently. Most of the graph machine learning models, such as Graph Neural Networks (GNNs), are trained over massive graph data. However, in many real-world scenarios, such as hospitalization prediction in healthcare systems, the graph data is usually stored at multiple data owners and cannot be directly accessed by any other parties due to privacy concerns and regulation restrictions. Federated Graph Machine Learning (FGML) is a promising solution to tackle this challenge by training graph machine learning models in a federated manner. In this survey, we conduct a comprehensive review of the literature in FGML. Specifically, we first provide a new taxonomy to divide the existing problems in FGML into two settings, namely, *FL with structured data* and *structured FL*. Then, we review the mainstream techniques in each setting and elaborate on how they address the challenges under FGML. In addition, we summarize the real-world applications of FGML from different domains and introduce open graph datasets and platforms adopted in FGML. Finally, we present several limitations in the existing studies with promising research directions in this field.

1. INTRODUCTION

In recent years, graphs have been widely used to represent complex data in a wide diversity of real-world domains, e.g., healthcare [112; 84], transportation [56; 136], bioinformatics [72; 137], and recommendation systems [13; 29]. Numerous graph machine learning techniques provide insights into understanding rich information hidden in graphs and show expressive performance in different tasks, such as node classification [139; 38] and link prediction [6; 21].

Although these graph machine learning techniques have made significant progress, most of them require a massive amount of graph data centrally stored on a single machine. However, with the emphasis on data security and user privacy [107], this requirement is often infeasible in the real world. Instead, graph data is usually distributed in multiple data owners (i.e., data isolation) and we are not able to collect the graph data in different places due to privacy concerns. For instance, a third-party company aims to train a graph machine learning model for a group of financial institutions to help them detect potential financial crimes and fraud cus-

tomers. Each financial institution owns its local dataset of customers, such as their demographics, as well as transaction records among them. The customers in each financial institution form a customer graph where edges represent the transaction records. Due to strict privacy policies and commercial competition, local customer data in each institution cannot be directly shared with the company or other institutions. Meanwhile, some institutions may have connections with others, which could be viewed as structural information among institutions. Generally, the major challenge for the company lies in training a graph machine learning model for financial crime detection based on the local customer graphs and structural information among institutions without directly accessing local customer data in each institution.

Federated Learning (FL) [76] is a distributed learning scheme which addresses the data isolation problem through collaborative training. It enables participants (i.e., clients) to jointly train a machine learning model without sharing their private data. Therefore, combining FL with graph machine learning becomes a promising solution to the aforementioned problem. In this paper, we term it Federated Graph Machine Learning (FGML). In general, FGML can be categorized as two settings with respect to the level of structural information. The first setting is *FL with structured data*. In FL with structured data, clients collaboratively train a graph machine learning model based on their graph data while keeping the graph data locally. The second setting is *structured FL*. In structured FL, there are structural information among the clients which forms a client-level graph. The client graph could be leveraged to design more effective federated optimization approaches.

While FGML provides a promising paradigm, the following challenges emerge and need to be addressed.

1. *Cross-client missing information.* A common scenario in FL with structured data is that each client owns a subgraph of the global graph and some nodes may have neighbors belonging to other clients. Due to privacy concerns, a node can only aggregate the features of its neighbors within the client but cannot access the features of those located on other clients, which leads to insufficient node representations [135; 88; 129; 11].

2. *Privacy leakage of graph structures.* In traditional FL, a client is not allowed to expose the features and labels of its data samples. In Fl with structured data, the privacy of structural information should also be considered. The structural information can be either directly exposed by sharing the adjacency matrix or

indirectly exposed by transmitting node embeddings [134; 66; 114; 92].

3. *Data heterogeneity across clients.* Unlike traditional FL where data heterogeneity comes from non-IID data samples [97; 46], graph data in FGML contains rich structural information [138; 50; 65; 51]. Meanwhile, divergent graph structures across clients can also affect performance of graph machine learning models.

4. *Parameter utilization strategies.* In structured FL, the client graph enables a client to obtain information from its neighbor clients. The effective strategies of fully utilizing neighbor information orchestrated by a central server or in a fully decentralized manner should be well designed in structured FL [55; 40; 78].

To tackle the above challenges, a great number of algorithms have been proposed in recent years. However, to the best of our knowledge, the existing surveys mainly focus on challenges and approaches in standard FL [58; 53; 123; 143; 100] yet only a few attempts have been made to survey specific problems and techniques in FGML [133; 63]. A position paper [133] provides a categorization of FGML but does not summarize main techniques in FGML. Another review paper [63] only covers a limited number of related papers in this topic and introduces the existing techniques very briefly. In this survey, we introduce the concepts of two problem settings in FGML. Then we review the current techniques under each setting and introduce real-world applications in FGML. We also summarize accessible graph datasets and platforms which can be used for applications of FGML. Finally, several promising future directions are presented. Our contributions in this paper can be summarized as follows.

- **Taxonomy of Techniques in FGML.** We propose a taxonomy of FGML based on different problem settings and summarize key challenges in each setting.

- **Comprehensive Technique Review.** We provide a comprehensive overview of the existing techniques in FGML. Compared with the existing reviews, we not only investigate a more extensive set of related work but also provide a more elaborate analysis of techniques instead of simply listing the steps of each method.

- **Real-World Applications.** We are the first to summarize real-world applications of FGML. We categorize the applications by their domains and introduce related works in each domain.

- **Datasets and Platforms.** We introduce the existing datasets and platforms in FGML, which facilitates developing algorithms and deploying applications in FGML for engineers and researchers.

- **Promising Future Directions.** We point out the limitations of the existing methods and provide promising future directions in FGML.

The rest of this paper is organized as follows. Section 2 briefly introduces definitions in graph machine learning as well as concepts and challenges of two settings in FGML. We review mainstream techniques in the two settings in Section 3 and Section 4, respectively. Section 5 further explores applications of FGML in the real world. Section 6 presents open graph datasets used in related FGML papers and two platforms for FGML. We also provide possible future directions in Section 7. Finally, Section 8 concludes this paper.

2. PROBLEM FORMULATION

In this section, we first present related definitions in graph machine learning and FL. Then we introduce the problem formulation of two different settings in FGML.

Notations. Throughout this paper, we use bold lowercase letters (e.g., $\mathbf{z}$) and bold uppercase letters (e.g., $\mathbf{A}$) to represent vectors and matrices, respectively. For any matrix, e.g., $\mathbf{A}$, we use $\mathbf{A}_i$ to denote its i-th row vector and $\mathbf{A}_{ij}$ to denote its (i, j)-th entry. The l_p norm of a vector $\mathbf{z}$ for $p \geq 1$ is denoted as $||\mathbf{z}||_p$. We use letters in calligraphy font (e.g., $\mathcal{V}$) to denote sets. $|\mathcal{V}|$ denotes the cardinality of set $\mathcal{V}$.

2.1 Graph Machine Learning

Definition 1. (Graphs) *A graph is $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the node set and $\mathcal{E}$ is the edge set. $v_i \in \mathcal{V}$ denotes a node and $e_{ij} = (v_i, v_j) \in \mathcal{E}$ denotes an edge between node v_i and node v_j.*

We use $\mathbf{A} \in \{0, 1\}^{n \times n}$ to represent the adjacency matrix of graph $\mathcal{G}$, where $n = |\mathcal{V}|$ is the total number of nodes. $\mathbf{A}_{ij} = 1$ implies that there exists an edge between node v_i and node v_j, otherwise $\mathbf{A}_{ij} = 0$. $\mathbf{D} \in \mathbb{R}^{n \times n}$ denotes the degree diagonal matrix where $\mathbf{D}_{ii} = \sum_j \mathbf{A}_{ij}$. The neighborhood of node v_i is defined as $N(v_i) = \{v_j \in \mathcal{V} | (v_i, v_j) \in \mathcal{E}\}$. For graph data with node features, we use $\mathbf{X} \in \mathbb{R}^{n \times d}$ to denote the node feature matrix where d is the number of node features.

Graphs can be categorized as homogeneous graphs (containing only one type of nodes and one type of edges) and heterogeneous graphs (whose nodes belong to more than one type of nodes and/or edges) according to the number of node types and edge types. The two typical heterogeneous graphs we mention in this paper are knowledge graphs (KGs) and user-item graphs.

Definition 2. (Knowledge Graphs) *A knowledge graph is a directed heterogeneous graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ where nodes are entities and edges are subject-property-object triple facts. Each edge of the form (head entity, relation, tail entity) (denoted as (h, r, t)) indicates a relationship r from a head entity h to a tail entity t.*

Definition 3. (User-Item Graphs) *A user-item graph is a heterogeneous graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Users and items serve as nodes and relations between users and items serve as edges. In some scenarios, relations also exist between users and between items.*

Definition 4. (Graph Machine Learning Models) *Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, a graph machine learning model f_ω parameterized by ω learns the node representations $\mathbf{H} \in \mathbb{R}^{n \times d_e}$ with respect to $\mathcal{G}$ for downstream tasks, where d_e is the dimension of node embeddings*

$$\mathbf{H} = f_\omega(\mathcal{G}). \tag{1}$$

For node classification tasks, we usually employ a softmax

|(a) FL with structured data | (b) Original FL | (c) Structured FL|

Figure 1: The framework comparison among original FL, FL with structured data and structured FL. (a) FL with structured data: each client owns graph data, i.e., a (sub)graph or multiple graphs. (b) Original FL: data samples on clients are typically Euclidean data and no links exist among edges. (c) Structured FL: clients are connected by links and form a client graph.

function to obtain the probability vector for each node based on its embedding, and then a loss function (e.g., cross entropy) is applied to measure the difference between predictions and the given node labels.

For graph classification tasks, a graph-level representation $\mathbf{h}_\mathcal{G}$ can be pooled from node representations

$$\mathbf{h}_\mathcal{G} = \mathrm{readout}(\mathbf{H}), \qquad (2)$$

where $\mathrm{readout}(\cdot)$ is a pooling function (e.g., mean pooling and sum pooling) which aggregates the embeddings of all nodes in the graph into a single embedding vector.

Without loss of generality, we mainly consider Graph Neural Networks (GNNs) (e.g., GCN [54], GAT [105], and Graph-Sage [37]) as graph machine learning models in this survey. In GNNs, each node v_i typically gathers the information from its neighbors $N(v_i)$ and aggregates them with its own information to update its representation $\mathbf{h}_i$. Mathematically, an L-layer GNN f_ω can be formulated as

$$\mathbf{h}_i^l = \sigma(\omega^l \cdot (\mathbf{h}_i^{l-1}, \mathrm{Agg}(\{\mathbf{h}_j^{l-1} | v_j \in N(v_i)\}))) \qquad (3)$$

for $l = 1, 2, \cdots, L$, where $\mathbf{h}_i^l$ is the representation of node v_i after the l-layer of f_ω and $\mathbf{h}_i^0 = \mathbf{X}_i$ is the raw feature of node v_i. ω^l is the learnable parameters in the l-layer of f_ω and $\mathrm{Agg}(\cdot)$ is the aggregation operation (e.g., mean pooling). σ is an activation function.

2.2 Federated Learning

FGML is a type of FL which involves structural information. Before introducing the concepts of two settings in FGML, we provide the definition of FL in this subsection.

Definition 5. (Federated Learning) *In standard FL, we consider a set of M clients $\mathcal{C} = \{c_k\}_{k=1}^M$. Each client c_k owns its private dataset $\mathcal{D}_k = \{(\boldsymbol{x}_i, y_i)_{i=1}^{N_k}\}$ sampled from its own data distribution, where $\boldsymbol{x}_i$ is the feature vector of i-th data sample and y_i is the corresponding label of the data sample. $N_k = |\mathcal{D}_k|$ is the number of data samples on client c_k and $N = \sum_{k=1}^M N_k$. Let l_k denote the loss function parameterized by ω on client c_k. The goal of FL is to optimize the overall objective function while keeping private datasets locally*

$$\min_\omega \sum_{k=1}^M \frac{N_k}{N} L_k(\omega) = \min_\omega \frac{1}{N} \sum_{k=1}^M \sum_{i=1}^{N_k} l_k(\mathbf{x}_i, y_i; \omega), \quad (4)$$

where L_k is the average loss over the local data on client c_k.

FedAvg [76] is a typical algorithm for federated optimization to obtain high model utility while preserving privacy. In FedAvg, only model parameters are transmitted between the central server and each client. Specifically, during each round t, the central server selects a subset of clients and sends them a copy of the current global model parameters ω^t for local training. Each selected client c_k updates the received copy ω_k^t by an optimizer such as stochastic gradient descent (SGD) for a variable number of iterations locally on its own dataset $\mathcal{D}_k$. Then the server collects updated model parameters ω_k^t from the selected clients and aggregates them to obtain a new global model ω^{t+1}. Finally, the server broadcasts the updated global model ω^{t+1} to clients for training in round $t + 1$.

It is worthwhile to note that GNN and FL both involve an *aggregation* operation. Aggregation in the context of GNN represents the operation that a node updates its representation by aggregating information from its neighbors. Aggregation in the context of FL represents the operation that the central server collects model parameters from clients and updates the global model parameters. Following the previous survey [63], we use *GNN aggregation* and *FL aggregation* in this survey to represent two aggregation operations in GNN and FL, respectively.

2.3 Federated Graph Machine Learning

Standard FL mainly deals with tasks on Euclidean data (e.g., image classification) and equally aggregates model parameters from clients. Different from standard FL, federated graph machine learning involves structural information in federated optimization. Based on the level of structural information, FGML can be categorized as two mainstreams.

Setting 1. (FL with Structured Data) In FL with structured data, clients possess private structured datasets (i.e., graphs) and jointly train a graph machine learning model orchestrated by the central server. Fig. 1(a) illustrates the framework of FL with structured data. Formally, each client c_k owns its private local data $\mathcal{D}_k = \{\mathcal{G}_1, \mathcal{G}_2, \cdots\}$, where each $\mathcal{G}_i = (\mathcal{V}_i, \mathcal{E}_i)$ is a graph with its node set $\mathcal{V}_i$ and edge set $\mathcal{E}_i$. The objective of FL with structured data is that each client collaboratively trains a graph machine learning model f_ω with other clients based on its local graph dataset

$\mathcal{D}_k$ while always keeping $\mathcal{D}_k$ locally. Note that each client in FL with structured data may have one single (sub)graph or multiple graphs. In general, clients train a graph machine learning model for graph-level tasks when each client c_k owns multiple graphs and N_k is the number of graphs on client c_k; on the contrary, when each client c_k owns one single graph $\mathcal{G}_k$ or a subgraph $\mathcal{G}_k$ of an entire graph, the graph machine learning model is for node-level tasks and $N_k = |\mathcal{V}_k|$ is the number of nodes in $\mathcal{G}_k$.

Setting 2. (Structured FL) In structured FL, relations exist among clients. When we take each client as a node, all the clients in structured FL will form a graph $\mathcal{G}^{\mathcal{C}} = \{\mathcal{V}^{\mathcal{C}}, \mathcal{E}^{\mathcal{C}}\}$ where $\mathcal{V}^{\mathcal{C}}$ is the client set and $\mathcal{E}^{\mathcal{C}}$ contains links between clients. Fig. 1(c) shows the framework of structured FL. Formally, given the client graph $\mathcal{G}^{\mathcal{C}} = \{\mathcal{V}^{\mathcal{C}}, \mathcal{E}^{\mathcal{C}}\}$, each client c_k collaboratively trains a machine learning model f_ω by interacting with its neighbors $N(c_k) = \{c_s | (c_k, c_s) \in \mathcal{E}^{\mathcal{C}}\}$. It is noteworthy that the datasets on clients do not have to be structured data.

In Section 3 and Section 4, we review the existing techniques in FL with structured data and structured FL and analyze how they solve the aforementioned challenges, respectively. We summarize these techniques in Table 1.

3. FL WITH STRUCTURED DATA

The goal of clients in FL with structured data is to jointly train a graph machine learning model based on their local graph datasets while preserving privacy. In this section, we review techniques in FL with structured data for improving model utility and tackling the aforementioned challenges. Fig. 2 illustrates the taxonomy of techniques in FL with structured data.

3.1 Cross-Client Information Reconstruction

When a graph is split into multiple subgraphs and each client owns a subgraph of the original graph, each node can only perform GNN aggregation on the information from a subset of its neighbors (i.e., those within the subgraph) but cannot obtain information from those located on other clients due to the privacy issue. The missing cross-client information leads to biased node embeddings on each client and therefore degrades the performance of graph machine learning models. The objective in this case is to reconstruct the important missing cross-client information for calculating node embeddings. The existing techniques can be categorized as intermediate result transmission and missing neighbor generation. The difference lies in whether the original global graph structure is known: the studies in intermediate result transmission assume that the central server is aware of the original graph structure, while those in missing neighbor generation do not.

3.1.1 Intermediate Result Transmission

When the central server is aware of the original graph structure (including missing cross-client links), it is able to collect intermediate results (e.g., node representations) in graph machine learning models from clients and compute node embeddings according to their complete neighbor lists including cross-client neighbors.

Considering an L-layer GCN model f_ω, the operation in the l-th layer of f_ω can be written as

$$\mathbf{H}^l = \sigma\left(\mathbf{L}\mathbf{H}^{l-1}\mathbf{W}^l\right), \tag{5}$$

(a) FL with structured data   (b) Cross-client information reconstruction

(c) Overlapping instance alignment   (d) Non-IID data adaptation
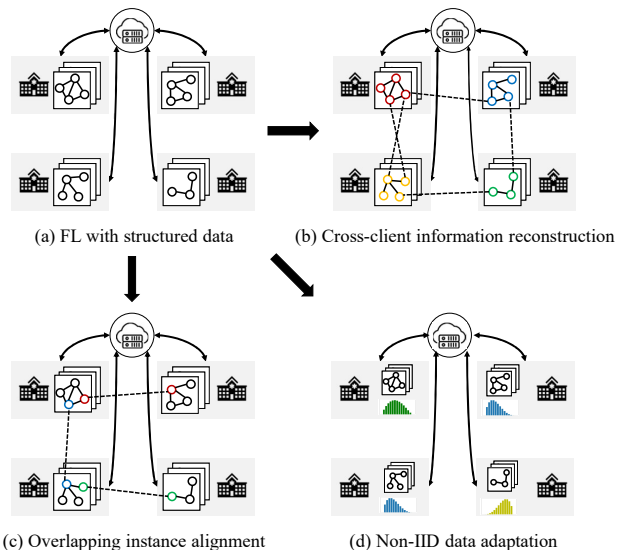
Figure 2: The taxonomy of techniques in FL with structured data. The techniques in (a) FL with structured data can be categorized as (b) cross-client information reconstruction which recovers missing links (e.g., dashed lines in (b)) between nodes from different clients, (c) overlapping instance alignment which aligns overlapping instances (e.g., nodes connected by dashed lines in (c)) among different clients, and (d) Non-IID data adaptation which tackles the non-IID characteristic of data across clients.

where $\mathbf{H}^l$ is the hidden representation after the l-th layer of f_ω and $\mathbf{H}^0 = \mathbf{X}$. $\mathbf{W}^l$ denotes the weight matrix of the l-th layer and $\mathbf{L} = (\mathbf{D}+\mathbf{I})^{-\frac{1}{2}}(\mathbf{A}+\mathbf{I})(\mathbf{D}+\mathbf{I})^{-\frac{1}{2}}$. In the federated setting, we rewrite the vanilla GCN model in a distributed manner, where the hidden matrix $\mathbf{H}^l_{(k)}$ after the l-th layer of f_ω on each client c_k can be computed by

$$\begin{aligned}
\mathbf{H}^l_{(k)} &= \sigma\left(\sum_{s=1}^{M} \mathbf{L}_{(ks)}\mathbf{H}^{l-1}_{(s)}\mathbf{W}^l_{(s)}\right) \\
&= \sigma\left(\mathbf{L}_{(kk)}\mathbf{H}^{l-1}_{(k)}\mathbf{W}^l_{(k)} + \sum_{s \neq k} \mathbf{L}_{(ks)}\mathbf{H}^{l-1}_{(s)}\mathbf{W}^l_{(s)}\right)
\end{aligned} \tag{6}$$

for $l = 1, 2, \cdots, L$, where $\mathbf{H}^0_{(k)} = \mathbf{X}_{(k)}$ are node features on client c_k. $\mathbf{W}^l_{(k)}$ denotes the local parameters in the l-th layer of the local graph machine learning model on client c_k. $\mathbf{L}_{(ks)}$ is a block of $\mathbf{L}$ corresponding to the rows of nodes in $\mathcal{V}_k$ and columns of nodes in $\mathcal{V}_s$.

An intuitive way of obtaining information from a node's neighbors located on other clients is to transmit node embeddings directly. To avoid exposing raw node features, each client performs GNN aggregation within its local subgraph when processing the first GCN layer [11; 64]

$$\mathbf{H}^1_{(k)} = \sigma(\mathbf{L}_{(kk)}\mathbf{X}_{(k)}\mathbf{W}^1_{(k)}). \tag{7}$$

Then, each client performs GNN aggregation for a node including the embeddings of its neighbors from other clients following Eq. (6) after the first GCN layer.

However, the block of degree diagonal matrix $\mathbf{D}_{(k)} \in \mathbb{R}^{N_k \times N_k}$ containing node degree information on client c_k is unknown for other clients due to privacy concerns and must be kept

Table 1: Summary of techniques in FGML.

Settings	Techniques	Approaches	Data on each client	Downstream Tasks
FL with Structured Data	Cross-Client Info. Recon.	FedGraph [11]	A subgraph	Node classification
	Cross-Client Info. Recon.	Glint [64]	A subgraph	Node classification
	Cross-Client Info. Recon.	PPSGCN [129]	A subgraph	Node classification
	Cross-Client Info. Recon.	FedSage+ [135]	A subgraph	Node classification
	Cross-Client Info. Recon.	FedNI [88]	A subgraph	Node classification
	Overlapping Ins. Align.	FedVGCN [81]	A graph	Node classification
	Overlapping Ins. Align.	VFGNN [142]	A graph	Node classification
	Overlapping Ins. Align.	FedSGC [17]	A graph	Node classification
	Overlapping Ins. Align.	SGNN [77]	A graph	Node classification
	Overlapping Ins. Align.	FedGL [10]	A subgraph	Node classification
	Overlapping Ins. Align.	FedE[15]	A KG	KG completion
	Overlapping Ins. Align.	FedR [134]	A KG	KG completion
	Overlapping Ins. Align.	FKGE [87]	A KG	KG completion
	Overlapping Ins. Align.	FedGNN [113]	A user-item graph	Rating prediction
	Overlapping Ins. Align.	FedPerGNN [114]	A user-item graph	Rating prediction
	Overlapping Ins. Align.	FeSoG [66]	A user-item graph	Rating prediction
	Non-IID data adaptation	FedAlign [61]	A subgraph	Node classification
	Non-IID data adaptation	FLIT+ [144]	Multiple graphs	Graph classification/regression
	Non-IID data adaptation	GraphFL [108]	A subgraph	Node classification
	Non-IID data adaptation	FML-ST [59]	A graph	Node regression
	Non-IID data adaptation	FedGCN [43]	Multiple graphs	Graph classification
	Non-IID data adaptation	ASFGNN [140]	A subgraph	Node classification
	Non-IID data adaptation	GCFL+ [117]	Multiple graphs	Graph classification
	Non-IID data adaptation	CTFL [132]	A subgraph	Node regression
Structured FL	Centralized Aggregation	SFL [12]	A node	Time series prediction
	Centralized Aggregation	BiG-Fed [120]	A node	Time series prediction
	Centralized Aggregation	CNFGNN [78]	A node	Time series prediction
	Fully Decentralized Trans.	D-FedGNN [86]	A node	Graph classification/regression
	Fully Decentralized Trans.	[55]	A node	Node regression
	Fully Decentralized Trans.	GFL [91]	A node	Node regression
	Fully Decentralized Trans.	FL DSGD/DSGT [69]	A node	Health record representation
	Fully Decentralized Trans.	FD DSGD/DSGT [70]	A node	Health record representation
	Fully Decentralized Trans.	[119]	A node	Image classification
	Fully Decentralized Trans.	SpreadGNN [40]	A node	Graph classification/regression
	Fully Decentralized Trans.	cPDS [5]	A node	Node classification
	Fully Decentralized Trans.	dFedU [24]	A node	Image classification
	Fully Decentralized Trans.	EF-HC [127]	A node	Image classification

locally during computation. There have been a series of corresponding solutions to this problem. For instance, PPS-GCN [129] reformulates each block $\mathbf{L}_{(ks)}$ by

$$\mathbf{L}_{(ks)} = (\mathbf{D}_{(k)} + \mathbf{I})^{-\frac{1}{2}} \cdot \tilde{\mathbf{L}}_{(ks)}, \qquad (8)$$

where $\tilde{\mathbf{L}}_{(ks)} = \mathbf{A}_{(ks)}(\mathbf{D}_{(k)} + \mathbf{I})^{-\frac{1}{2}}$. Therefore, Eq. (6) can be rewritten as

$$\begin{aligned} \mathbf{H}_{(k)}^l = \sigma(\mathbf{L}_{(kk)}\mathbf{H}_{(k)}^{l-1}\mathbf{W}_{(k)}^l \\ + (\mathbf{D}_{(k)} + \mathbf{I})^{-\frac{1}{2}} \sum_{s \neq k} \tilde{\mathbf{L}}_{(ks)}\mathbf{H}_{(s)}^{l-1}\mathbf{W}_{(s)}^l) \end{aligned} \qquad (9)$$

for $l = 1, 2, \cdots, L$, where $\tilde{\mathbf{L}}_{(ks)}\mathbf{H}_{(s)}^{l-1}\mathbf{W}_{(s)}^l$ can be calculated locally within client c_s and transmitted to the server. Consequently, we can learn node representations over clients without exchanging local graph structure information.

Although the raw data (i.e., node features and structural information) are preserved with intermediate result transmission, it requires the structural information of the original graph to compute node embeddings, which might be impractical in the real world. Furthermore, intermediate re-

sults are transmitted multiple times for each local update, which also brings significant communication costs.

3.1.2 Missing Neighbor Generation

When the original graph structure is unknown to the server, the techniques in intermediate result transmission fail since the cross-subgraph links carrying important information will never be captured by any client [135]. To tackle this issue, several approaches of missing neighbor generation have been proposed recently.

The intuition of missing neighbor generation is to design a missing neighbor generator to reconstruct the features of a node's cross-subgraph neighbors on other clients [135; 88]. Concretely, each client c_k first hides a subset of nodes and related edges in its local subgraph $\mathcal{G}_k$ following a specific strategy (e.g., Breadth-First Search) [88] to form an impaired subgraph $\bar{\mathcal{G}}_k$. Then each client trains $\mathcal{X}$ a predictor parameterized by θ_d for predicting the number $\tilde{n}_i$ of the masked neighbors of each node v_i in $\bar{\mathcal{G}}_k$ and an encoder (e.g., GCN) parameterized by θ_f for predicting the features $\tilde{\mathbf{X}}_{(i)}$

of the masked neighbors by minimizing the loss

$$L^n = \lambda^d L^d(\tilde{n}_i, n_i; \theta_d) + \lambda^f L^f(\{\tilde{\mathbf{X}}_{(i)}\}_{v_i \in \bar{\mathcal{G}}_k}, \{\mathbf{X}_{(i)}\}_{v_i \in \bar{\mathcal{G}}_k}; \theta_f), \quad (10)$$

where L^d and L^f are the loss functions for the predicted number of masked neighbors and their predicted features, respectively. $\mathbf{X}_{(i)}$ is the features of node v_i's masked neighbors. A cross-subgraph feature reconstruction term [135] is introduced to L^f, aiming to recover features of cross-subgraph missing neighbors by decreasing the distance between predicted node features and the closest node feature on other clients.

3.2 Overlapping Instance Alignment

In applications, an instance (e.g., a node in homogeneous graphs or an entity in KGs) could belong to two or more clients. Under this setting, the embeddings of an overlapping instance from different clients may come from different embedding spaces during collaborative training. To tackle this problem, the overlapping instance alignment technique is proposed [10; 142; 87]. The key idea is to learn global instance embeddings based on the local instance embeddings from clients. This technique can be applied to homogeneous graphs, KGs, and user-item graphs.

3.2.1 Homogeneous Graph-Based Alignment

The existing works about instance alignment in homogeneous graphs are mainly under vertical FL [123]. In generic vertical FL, clients have overlapping nodes but differ in the feature space. Unlike generic vertical FL, structural information in graph data is also taken into account in FGML. One common scenario is that a set of nodes are located on all clients but their features and relations are different across clients [81; 142]. The central server can collect local node embeddings from clients and align overlapping node embeddings. For instance, VFGNN [142] first computes local node embeddings $\mathbf{H}_{(k)}$ on each client c_k using a graph machine learning model. Then it combines the embeddings of each node v_i via a combination strategy

$$\mathbf{H} \leftarrow \text{COMBINE}(\{\mathbf{H}_{(k)}\}_{k=1}^M), \quad (11)$$

where $\text{COMBINE}(\cdot)$ denotes a combination operator (e.g., Concat, Mean and Regression).

Another scenario is that one client only contains structural information of graph data and other clients only contain node features [17]. In this scenario, graph machine learning models cannot be simply applied on each client since structural information and node features are located in different clients. The related techniques deal with this problem by protecting structural information and raw node features simultaneously during federated optimization. For example, SGNN [77] replaces the original adjacency matrix with a structural similarity matrix $\mathbf{A}^s$ where entry $\mathbf{A}_{ij}^s$ measures the structural similarity between node v_i and node v_j. SGNN computes $\mathbf{A}_{ij}^s$ by

$$\mathbf{A}_{ij}^s = \exp(-\text{dist}(\text{OD}(N(v_i)), \text{OD}(N(v_j)))), \quad (12)$$

where $\text{OD}(\cdot)$ returns a list of ordered degree given the input node list and $\text{dist}(\cdot, \cdot)$ is a distance function (e.g., dynamic time warping (DTW) [83]). Then SGNN embeds original features on the clients which only contain features using one-hot encoding and finally computes node embeddings with the structural similarity matrix. FedSGC [17] applies Ho-

momorphic Encryption (HE) [1] for secure transmission of the adjacency matrix and node features.

3.2.2 KG-Based Alignment

Suppose in a federated KG each client owns one KG and each KG may have overlapping entities which also exist on other clients. The key technique to improve the performance of KG embedding is to align the embeddings of overlapping entities across KGs [15; 134; 87]. More specifically, after each round t, the server collects each local embedding matrix from each client to update the global embedding matrix. Then the server distributes the global embeddings to corresponding clients for subsequent local training. As the first FL framework for KGs, FedE [15] enables the server to record all the unique entities from clients with an overall entity table. The server collects the entity embedding matrix $\mathbf{E}_{(k)}^t$ of each client c_k and aligns them by

$$\mathbf{E}^t = \left(\mathbb{1} \oslash \sum_{k=1}^M \mathbf{v}_k\right) \otimes \sum_{k=1}^M \mathbf{P}_{(k)} \mathbf{E}_{(k)}^t, \quad (13)$$

where $\mathbf{E}^t$ is the global entity embedding matrix, $\mathbb{1}$ denotes an all-one vector, $\oslash$ denotes element-wise division for vectors and $\otimes$ denotes element-wise multiplication with broadcasting. $\mathbf{P}_{(k)}$ denotes client c_k's permutation matrix that maps client c_k's entity matrix to the server's entity table. $\mathbf{v}_k$ denotes client c_k's existence vectors.

As the server maintains a complete table of entity embeddings, it can easily infer a relation embedding between two entities h and t by calculating

$$r' = \arg\max_r f(h, r, t), \quad (14)$$

where $f(\cdot)$ denotes a score function (e.g., TransE [4]) [134]. To tackle the privacy issue, FedR [134] was proposed based on relation embedding alignment.

Instead of aligning embeddings on the server, FKGE [87] enables entity alignment between clients. Inspired by PATE-GAN [52], FKGE involves a privacy-preserving adversarial translation (PPAT) network for adversarial learning. The PPAT network employs a generator as well as a student discriminator and multiple teacher discriminators. The generator first translates aligned entities' embeddings from $\mathcal{G}_k$ into synthesized embeddings and sents them to $\mathcal{G}_s$. The student and teacher discriminators distinguish between the synthesized embeddings and ground truth embeddings in $\mathcal{G}_s$ for each pair of KGs $(\mathcal{G}_k, \mathcal{G}_s)$ which have aligned entities $\mathcal{E}_k \cap \mathcal{E}_s$ and relations $\mathcal{R}_k \cap \mathcal{R}_s$. During the alignment, only synthesized embeddings and gradients are transmitted among clients and data privacy can be guaranteed [33].

3.2.3 User-Item Graph-Based Alignment

In a federated recommendation system, each user only has a first-order local user-item subgraph with its own item rating and its neighbors located on its device. A naive method is to align the embeddings of overlapping users and items directly. However, the server can easily infer a user's user-item links by recording the items with non-zero-gradient embeddings from this user because an item embedding gets updated on the user only when the item has the rating score from the user [66].

To tackle the privacy leakage, pseudo interacted item sampling and Local Differential Privacy (LDP) [30] techniques

are two common strategies [113; 114; 66]. Before sending gradients to the central server, each user u_k first samples some items that it has not interacted with (i.e., pseudo interacted items). Then it generates embedding gradients of the sampled items (e.g., using a Gaussian distribution) and combines them with the real embedding gradients. Finally, the user applies an LDP module to modify gradients by clipping and adding zero-mean Laplacian noise to gradients

$$g\prime_k = \text{clip}(g_k, \delta) + \text{Laplace}(0, \lambda), \qquad (15)$$

where g_k is the unified gradients of user u_k including model gradients and user/item embedding gradients, $\text{clip}(g_k, \delta)$ denotes limiting g_k with the threshold δ and λ is the strength of Laplacian noise.

3.3 Non-IID Data Adaptation

The data distribution on each client may diverge a lot both in node features and graph structures [117]. Such data heterogeneity may lead to severe model divergence in the federated setting and therefore degrade the performance of the global model. The intuition of mitigating the problem is either to train an effective global model or to train specialized models for each client. The existing techniques handling this problem can be categorized as single global model-based methods and personalized model-based methods.

3.3.1 Single Global Model-Based Methods

The goal of single global model-based methods is to train a global graph machine learning model over graph data from clients. The existing techniques tackle non-IID data across clients by designing loss functions and reweighting FL aggregations and interpolating local models.

Loss Function Designing. The intuition of loss function designing is to replace the original loss function, which is just for high model utility, with a new well-designed loss function that is also targeted at data heterogeneity. A common strategy is to add regularization terms into the local loss function. For instance, to deal with relational data (e.g., KGs) heterogeneity across clients, FedAlign [61] minimizes the average Optimal Transportation (OT) distance [106] between the basis matrices in basis decomposition [95] among clients. Mathematically, to train an L-layer graph machine learning model, the regularization term on client c_k can be rewritten as

$$L_k^r = \frac{\mu}{M} \sum_{k \neq s}^{M} \sum_{l=1}^{L} \text{OT}(\mathbf{V}_{(k)}^l, \mathbf{V}_{(s)}^l) + \lambda(||\nabla L_k(\omega)||_2 - 1)^2, \quad (16)$$

where $\mathbf{V}_{(k)}^l$ is the basis of l-th layer of the local graph machine learning model on client c_k and $\text{OT}(\cdot)$ computes OT distance. The second term is a weight penalty to make the objective function quasi-Lipschitz continuous. μ and λ are hyperparameters to adjust the contributions of each term.

In addition to regularization, another strategy in loss function designing is instance reweighting. For instance, FILT+ [144] pulls the local model closer to the global by minimizing the loss discrepancy between a local model and the global model. Specifically, FILT+ reweights instances on client c_k by putting more weights on samples with less confidence in the loss function

$$L_k^u = \sum_{i=1}^{N_k} (1 - \exp(-\Phi(\mathbf{x}_i, \omega, \omega_k)))^\gamma l_k(\hat{y}_i, y_i; \omega_k), \qquad (17)$$

where $\Phi(\cdot)$ is defined as

$$\Phi(\mathbf{x}_i, \omega, \omega_k) = \phi(\mathbf{x}_i, \omega_k) + \max(\phi(\mathbf{x}_i, \omega_k) - \phi(\mathbf{x}_i, \omega), 0). \quad (18)$$

Here γ is a hyperparameter and $\phi(\mathbf{x}_i, \omega)$ indicates the uncertainty of training sample $\mathbf{x}_i$ under the model ω. Generally, if the local model ω_k on client c_k is less confident about a sample $\mathbf{x}_i$ than the global model ω, this sample will obtain a higher weight in the objective function.

Inspired by the model-agnostic meta-learning (MAML) [28], some meta learning-based methods in FGML [108; 59] rewrite the loss function on each client c_k as

$$L_k^u = \sum_{i=1}^{N_k} l_k(\omega - \alpha \nabla l_k(\omega)), \qquad (19)$$

where α is a hyperparameter. Although the meta learning-based methods do not minimize the discrepancy between local models and the global model, they find an initial global model which can be easily adapted by clients after performing one or a few extra local updates.

FL Aggregation Reweighting. Apart from the loss function designing, reweighting local models during FL aggregation is also a solution to deal with Non-IID data. FedGCN [43] tries to reweight local model parameters via an attention mechanism. Considering an L-layer model f_ω, FedGCN assigns adaptive weights $\{\beta_k^{t,l}\}_{l=1}^L$ to the model parameter ω_k from each client c_k in round t for FL aggregation

$$\omega^{t+1,l} = \sum_{k=1}^{M} \beta_k^{t,l} \omega_k^{t,l} \qquad (20)$$

for $l = 1, 2, \cdots, L$. $\beta_k^{t,l}$ can be calculated through a softmax operation of score function $\alpha_k^{t,l}$

$$\alpha_k^{t,l} = \text{Attn}(\omega_k^{t,l}, \omega^{t,l}) = \mathbf{p}_k^l[\omega_k^{t,l}; \omega^{t,l}], \qquad (21)$$

where $\text{Attn}(\cdot)$ is the attention mechanism, $[\cdot; \cdot]$ indicates a concat operation, and $\mathbf{p}_k^l$ is a trainable vector. As a result, $\{\beta_k^{t,l}\}_{l=1}^L$ can dynamically measure the closeness between each local model ω_k and the global model ω.

Model Interpolation. The model interpolation technique developed based on parameter weighted average of the global and the local models. Specifically, the model ω_k^t on client c_k is a combination of its local model ω_k^{t-1} and the global model ω^t [140]

$$\omega_k^t = \alpha_k \omega_k^{t-1} + (1 - \alpha_k)\omega^t, \qquad (22)$$

where α_k is a mixing weight. The authors of [140] calculate it as Jensen–Shannon divergence [60] between local and global data distributions.

3.3.2 Personalized Model-Based Methods

Unlike training a single global graph machine learning model, the goal of learning personalized models is to train personalized graph machine learning models for each client. The resulting personalized models are tailored for specific clients and thus result in good performance. Formally, the objective function for training personalized graph machine learning models can be rewritten as

$$\min_{\omega_1, \omega_2, \cdots, \omega_M} \sum_{k=1}^{M} \frac{N_k}{N} L_k(\omega_k). \qquad (23)$$

(a) Structured FL

(b) Centralized aggregation

(c) Fully decentralized transmission

Figure 3: The taxonomy of techniques in structured FL. The techniques in (a) structured FL can be categorized as (b) centralized aggregation and (c) fully decentralized transmission. In (b) centralized aggregation, the server performs the FL aggregation and computes global client embeddings based on the client graph. In (c) fully decentralized transmission, parameters are transmitted among clients based on the client graph.

One common strategy for training personalized graph machine learning models is client clustering [94; 130; 31]. The intuition of client clustering is that the clients with similar data distribution can be clustered in a group and the clients in a group share the same model parameters. The basic idea of client clustering in FGML is to dynamically assign clients to multiple clusters based on their latest gradients of graph machine learning models [117; 132]. One problem in this idea is that the clustering result is significantly influenced by the latest gradients from clients, which are usually unstable during local training [117]. GCFL+ solves this problem by taking series of gradient norms into account for client clustering. Unlike collecting parameters in a cluster-wise manner in GCFL, CTFL [132] updates the global model based on representative clients of each cluster. A client's local model is updated as the average of the global model and the representative model of the cluster which includes the client.

4. STRUCTURED FL

In the real world, a client may have connections with others, such as road paths existing among traffic sensors. These connections usually contain rich information (e.g., the similarity of data distribution) among clients. Considering these connections, the clients can form a client graph. Structured FL takes the client graph $\mathcal{G}^{\mathcal{C}} = \{\mathcal{V}^{\mathcal{C}}, \mathcal{E}^{\mathcal{C}}\}$ into account and enables a client to obtain information from its neighbors. The key techniques in structured FL can be categorized as centralized aggregation and fully decentralized transmission. Fig. 3 illustrates the taxonomy of techniques in structured FL.

4.1 Centralized Aggregation

In structured FL with the central server, there exists structural information among clients. It is natural for the server to consider the structural information while updating parameters for each client. Generally, the server first collects parameters from clients as it does in standard FL. Then it updates parameters for each client through a graph machine learning model based on the client graph $\mathcal{G}^{\mathcal{C}}$ and finally sends the updated parameters back to clients.

Transmitting Model Parameters. Transmitting local model parameters to the central server as the input of graph machine learning models is a straightforward strategy. For instance, the central server in SFL [12] collects local model parameters $\{\omega_k^t\}_{k=1}^M$ from clients in round t and employs a GCN [54] to compute a graph-based local model parameters $\{\phi_k^{t+1}\}_{k=1}^M$ by

$$\{\phi_k^{t+1}\}_{k=1}^M \leftarrow \text{GCN}(\mathbf{A}^{\mathcal{C}}, \{\omega_k^t\}_{k=1}^M), \qquad (24)$$

where $\mathbf{A}^{\mathcal{C}}$ is the adjacency matrix of $\mathcal{G}^{\mathcal{C}}$ and $\text{GCN}(\cdot)$ represents the operations in GCN. The global model parameters ω^{t+1} are calculated by a $\text{readout}(\cdot)$ operation

$$\omega^{t+1} = \text{readout}(\{\phi_k^{t+1}\}_{k=1}^M). \qquad (25)$$

Each client c_k updates local mode parameters in round t by minimizing the local loss

$$\min l_k(\omega_k) + \lambda[R(\omega_k^t, \omega^t) + R(\omega_k^t, \phi_k^t)], \qquad (26)$$

where $R(\cdot)$ is a regularization term.

Similarly, the server in BiG-Fed [120] collects local model parameters $\{\omega_k^t\}_{k=1}^M$ as client representatives and computes a global client embedding $\mathbf{H}$ through a graph machine learning model based on $\{\omega_k^t\}_{k=1}^M$ and the client graph. Inspired by contrastive learning [16], Big-Fed optimizes the model by

$$\min \frac{1}{n} \sum_{k=1}^M \sum_{c_s \in N(c_k)} 1 - \cos(\mathbf{H}_k, \mathbf{H}_s)$$
$$+ \frac{1}{n} \sum_{k=1}^M \mathbb{E}_{c_s \sim P_k} \max(0, \cos(\mathbf{H}_k, \mathbf{H}_s)), \qquad (27)$$

where $\mathbf{H}_k$ is the global embedding of client c_k and P_k is a client sampling strategy. By minimizing the objective function in Eq. (27), each client will obtain a local model closer to its neighbors.

Transmitting Embeddings. Instead of gathering model parameters, another strategy for the server is to collect local embeddings from each client and compute global embeddings through graph machine learning models. For instance, CNFGNN [78] assumes that each client c_k represents a sensor with time series data $\mathbf{x}_k$. A client c_k first computes a temporal embedding $\mathbf{h}_k$ by a local encoder (e.g., GRU [18]) which models local temporal dynamics. The central server collects temporal embeddings $\{\mathbf{h}_k\}_{k=1}^M$ from clients and employs a graph machine learning model to compute the spatial embeddings $\{\mathbf{h}_k^{\mathcal{G}}\}_{k=1}^M$ of clients based on their temporal embeddings and client graph $\mathcal{G}^{\mathcal{C}}$. As a result, $\mathbf{h}_k^{\mathcal{G}}$ integrates information from client c_k's neighbors and involves spatial dynamics. $\mathbf{h}_k^{\mathcal{G}}$ is finally sent back to each client c_k as the input of a local decoder with $\mathbf{h}_k$ to make prediction.

4.2 Fully Decentralized Transmission

In the federated setting, one crucial bottleneck lies in high communication cost on the central server [40]. A feasible

solution to tackle this issue is to train a model in a fully decentralized fashion. Since the clients in structured FL form a client graph $\mathcal{G}^{\mathcal{C}} = \{\mathcal{V}^{\mathcal{C}}, \mathcal{E}^{\mathcal{C}}\}$, each client $c_k \in \mathcal{V}^{\mathcal{C}}$ can transmit parameters with its neighbors $N(c_k)$. Specifically, when a client c_k receives model parameters $\{\omega_s^t | c_s \in N(c_k)\}$ from its neighbors $N(c_k)$ [86; 55; 91], it can perform GNN aggregation to update its local model parameters ω_k^{t+1} by

$$\omega_k^{t+1} = \text{AGG}(\{\omega_s^t | c_s \in N(c_k)\}), \tag{28}$$

where $\text{AGG}(\cdot)$ is the aggregation function. An intuitive method [86; 91] following this strategy is to let each client c_k directly sum up the model parameters from its neighboring clients

$$\omega_k^{t+1} = \sum_{c_s \in N(c_k)} \mathbf{A}_{ks}^{\mathcal{C}} \cdot \omega_s^t. \tag{29}$$

However, this method results in losing local information of each client since the updated model is fully determined by neighboring information. This issue can be mitigated by involving local information during aggregation [69; 119; 40; 5; 24; 127]. For example, the authors of [69; 70] directly add local gradients during aggregation. Formally, the operation can be written as

$$\omega_k^{t+1} = \sum_{c_s \in N(c_k)} \mathbf{A}_{ks}^{\mathcal{C}} \cdot \omega_s^t - \alpha \nabla L_k(\omega_k^t). \tag{30}$$

Some methods [119] choose to retain local model parameters. Formally, these methods can be written as

$$\omega_k^{t+1} = \mathbf{A}_{kk}^{\mathcal{C}} \cdot \omega_k^t + \sum_{c_s \in N(c_k)} \mathbf{A}_{ks}^{\mathcal{C}} \cdot \omega_s^t - \alpha \nabla L_k(\omega_k^t). \tag{31}$$

5. APPLICATIONS

The number of applications of FGML is greatly increasing in various domains such as transportation, computer vision, recommendation systems, and healthcare. In this section, we elaborate some representative applications of FGML.

5.1 Transportation

Traffic prediction plays an important role in urban computing since it benefits reducing traffic congestion and improving transportation efficiency in smart cities [141]. The target of traffic prediction is to predict traffic speed or traffic flow of regions or road segments based on historical data collected by devices deployed in each region or road segment. Traffic data containing spatial-temporal information can be naturally represented as graphs and used as inputs of graph machine learning models for traffic flow prediction. In the cross-silo setting where a client typically represents an organization, each client constructs a local graph including devices on the client as nodes and edges are formulated by their physical properties (e.g., Euclidean distances). The clients jointly train a graph machine learning model for traffic flow prediction [67; 131; 132; 130]. In the cross-device setting where a client represents a device, the clients can form a graph based on their structure information and graph machine learning models are employed to calculate the embeddings of the clients. Due to the privacy issue, each client uploads its temporal embeddings to the central server instead of its raw data [78; 125].

Moreover, other applications of FGML in transportation systems, such as location representation [36; 116], routing planing [128; 124], and user mobility anomaly detection [103], are also attracting increasing attention.

5.2 Computer Vision

The existing applications of FGML in computer vision consist of image classification and object trajectory prediction [7; 9; 49]. The intuition is to construct graphs which incorporate semantic relationships among classes and objects on each client and embed them via graph machine learning models. For image classification, a graph is constructed on each client to represent classes (or domains) and the connections among them; then the clients jointly train a graph machine learning model to learn class embeddings [7; 9]. Typically, the existing techniques for image classification mainly employ CNN-based methods as the backbone model. Apart from the CNN-based backbone model (e.g., ResNet [42] and DenseNet [45]), domain and class-specific features and the GNN models are also transmitted for aggregation during federated optimization. For object trajectory prediction, the idea is to construct a series of dynamic graphs from videos. Each graph represents objects and their spatial relationships in a video frame. Besides aggregating information from nearby objects in the current graph, each object also aggregates dynamic information from those in the previous graph by a dynamic GNN. In a federated setting (e.g., a distributed surveillance system), the dynamic GNN is transmitted for collaborative training [49].

5.3 Recommendation Systems

The downstream task in graph-based federated recommendation systems is to produce high-quality item rating for each user. A graph machine learning model learns to predict unobserved item rating to a user by leveraging the embeddings of users and items in a user's local user-item subgraph with its own item rating. Aside from transmitting model parameters, user embeddings and item embeddings are also collected by the server, which leads to information leakage of item rating becoming a primary concern in federated recommendation systems [66; 113; 114]. This privacy issue is mitigated via adding the embeddings of pseudo interacted items and noise to gradients [113; 66].

5.4 Healthcare

Medical data such as medical images and disease symptoms are very sensitive and private, which results in medical datasets usually existing in isolated hospitals and medical institutes. The applications of FGML in healthcare are targeted at disease and hospitalization prediction based on medical images or health records stored in different hospitals and institutes. One key feature of healthcare applications is complex connections among medical images or health records because of patient interactions. The common strategy of modeling the connections is to construct a graph where each node represents an image or a medical record from a patient [104]. Due to the privacy issue, the graph is split into multiple subgraphs located in different hospitals and institutes. Several papers deal with cross-hospital links either by reconstructing cross-hospital connections [88] or by transmitting parameters with nearby hospitals [69; 70; 5]. Fed-CBT [2] learns connectional brain template representations by modeling them as graphs and fuses the graphs of one subject with a deep graph normalizer (DGN) [35]. STFL [68] takes each channel in polysomnography recordings as a node in a graph and trains a graph machine learning model for federated graph classification.

5.5 Other Applications

Apart from the aforementioned applications, FGML has manifold applications in other domains. A number of related papers have explored applications of FGML to various problems, such as human activity recognition [93], neural architecture search [109], packet routing [73], malware detection [20; 85], multi-armed bandits [145], drug discovery [74], and financial crime detection [99].

6. DATASETS AND PLATFORMS

In this section, we summarize the existing open graph datasets and platforms used for FGML research.

6.1 Datasets

We organize and introduce examples of real-world datasets in Table 2. These datasets are categorized by different application domains, such as citation networks, coauthor networks, social networks, molecules, proteins, KGs, recommendation systems, transportation, and healthcare. For each dataset, we provide basic statistical information including the number of graphs, the (average) number of nodes, the (average) number of edges. We also list the corresponding papers that use these datasets.

6.2 Platforms

Although a number of platforms [3; 8; 146] facilitate FL applications in multiple domains such as vision and language, only a few of them provide off-the-shelf supports for graph datasets and graph machine learning models. To the best of our knowledge, FedGraphNN [39] and FederatedScope-GNN (FS-G) [111] are the two platforms supporting tasks in FGML among the existing FL platforms.

FedGraphNN. FedGraphNN is an FL benchmark system for GNNs in FedML ecosystem [41]. It contains 36 graph datasets from 7 domains, such as molecules, proteins, KGs, recommendation systems, citation networks, and social networks. As for graph machine learning models, it supports a series of popular GNN-based models, such as GCN [54], GAT [105], GraphSage [37], and GIN [122], implemented via PyTorch Geometric [27].

FederatedScope-GNN. FS-G is built upon an event-driven FL framework FederatedScope [118] and it is compatible with various graph machine learning backends. For the ease of benchmarking approaches in FGML, FS-G incorporates two components specifically designed for FGML: Graph-DataZoo and GNNModelZoo. A GraphDataZoo integrates a rich collection of splitting mechanisms for dispersing a given standalone graph dataset into multiple clients. A GNNModelZoo integrates a number of algorithms including vanilla graph machine learning models (e.g., GCN [54], GAT [105], and GraphSage [37]) and some recent frameworks in FGML, such as FedSage+ [135], GCFL+ [117], and Fedgnn [113]. FS-G also supports federated hyper-parameter optimization and model personalization for FGML.

7. OPEN CHALLENGES AND FUTURE DIRECTIONS

In this section, we present some limitations in current studies and provide promising directions for future advances.

Data Heterogeneity of Graph Structures. Unlike the original FL where data distribution of features and labels are considered, the non-IID characteristic of graph structures is also a key challenge. Although a few papers analyze non-IID graph structures across clients [117], few of them address this issue completely. Despite some popular approaches designed for mitigating the non-IID characteristic in standard FL [31; 97; 46], the approaches in FGML should take graph structures into account.

Secure Aggregation of Instance Embeddings. As mentioned in Section 3, the central server may collect instance embeddings from clients for instance alignment in FGML, especially in KGs and user-item graphs. The existing studies mainly apply LDP techniques and pseudo instance sampling to alleviate privacy leakage due to embedding transmission [66; 113]. However, these algorithms can lead to performance degradation because of introducing noise. Thus, designing an effective yet secure aggregation scheme in FGML is still an open problem.

Communication Reduction Strategies in FGML. Communication is a critical bottleneck in the federated setting [76]. Large communication overhead makes it difficult to train an FL model. In FGML, we should consider both client-level and node-level relationships between different clients and it usually requires extra communication along these relations [40; 129]. Therefore, communication might be a more serious bottleneck in FGML. Further applied research on FGML should better consider the communication reduction strategies such as data compression [62] and local updating [58].

Fairness in FGML. Fairness is an important topic in FL. Without accessing the sensitive information (e.g., gender and race) in different clients, FL models might show distinct bias against some groups of data [26]. Furthermore, we want the model to have similar performance in each client in some FL scenarios [19; 57]. They are both requirements of fairness in FL. Considering structural information in FL, FGML provides many non-trivial challenges of fairness, e.g., how the structural information affects different fairness metrics in the federated training process. Novel fairness-aware FGML models are greatly expected.

Poisoning Attacks and Defenses in FGML. Recently, a few studies about poisoning attacks and defenses have been proposed [14; 121]. Apart from poisoning attacks on data features and model parameters in standard FL, poisoning attacks on graph structures can also affect collaborative training in FGML. Designing efficient attacks on graph structures in FGML and defending against such attacks could be a promising topic in security.

Benchmarks and Platforms. Compared with abundant benchmarks and platforms in standard FL [146; 3; 8; 118], applicable benchmarks and platforms in FGML are still at an infant stage. The current two platforms, FederatedScope-GNN [111] and FedGraphNN [39], lack either graph datasets or off-the-shelf FGML algorithms. In addition, splitting mechanisms in FGML are significantly different from those in standard FL, especially when a graph is split into multiple subgraphs across clients. Practical distributed graph data from the real world is needed for more practical graph partition.

8. CONCLUSIONS

A large number of powerful graph machine learning models have achieved remarkable success in different domains. How-

Table 2: Summary of graph datasets.

Category	Datasets	# Graphs	(Avg.) # Nodes	(Avg.) # Edges	Works Citation
Citation Networks	CORA [75]	1	2,708	5,429	[135; 11; 64; 81; 142; 108]
	CITESEER [32]	1	4,230	5,358	[135; 11; 81; 142; 108]
	PUBMED [96]	1	19,717	44,338	[135; 11; 64; 129; 81; 142]
	arXiv [44]	1	169,343	2,315,598	[142]
Coauthor Networks	CS [98]	1	18,333	81,894	[108; 135]
	Physics [98]	1	34,493	247,962	[64]
	Aminer [102]	1	3,923	9,023	[77]
Social Networks	KarateClub [126]	1	34	156	[120]
	REDDIT [37]	1	232,965	114,848,857	[11; 64; 129]
	REDDIT-BINARY [80]	2,000	429.63	497.75	[43]
	GITHUB [80]	12,725	113.79	234.64	[43]
	IMDB-BINARY [80]	1,000	19.77	96.53	[43; 117]
	IMDB-MULTI [80]	1,500	13	65.94	[43; 117]
	COLLAB [80]	5,000	74.49	2457.78	[43; 117]
	Amazon2M [44]	1	2,449,029	61,859,140	[108; 129]
Molecules	FreeSolv [115]	642	8.72	25.6	[144; 86]
	Lipophilicity [115]	4,200	27.04	86.04	[144; 86]
	ESOL [115]	1,128	13.29	40.65	[144; 86]
	MUV [115]	93,087	24.23	76.80	[40]
	QM8 [115]	21,786	7.77	23.95	[40]
	QM9 [115; 80]	133,885	8.8	27.6	[144]
	Tox21 [115; 80]	7,831	18.51	25.94	[144; 86; 40]
	SIDER [115]	1,427	33.64	35.36	[144; 86; 40]
	ClinTox [115]	1,478	26.13	27.86	[144; 86]
	BBBP [115]	2,039	24.05	25.94	[144; 86]
	BACE [115]	1,513	34.12	36.89	[144; 86]
	hERG [115]	10,572	29.39	94.09	[86]
	NCI1 [80]	4,110	29.87	32.3	[43; 117]
	MUTAG [80]	188	17.93	19.79	[117]
	BZR [80]	405	35.75	38.36	[117]
	COX2 [80]	467	41.22	43.45	[117]
	DHFR [80]	467	42.43	44.54	[117]
	PTC_MR [80]	344	14.29	14.69	[117]
	AIDS [80]	2,000	15.69	16.20	[117]
Proteins	DDI [80]	1,178	284.32	715.66	[43; 117]
	PROTEINS [80]	1,113	39.06	72.82	[43; 117]
	ENZYMES [80]	600	32.63	62.14	[43; 117]
Knowledge Graphs	FB15k-237 [22]	1	14,505	212,110	[15; 134]
	WN18RR [22]	1	40,559	71,839	[134]
	NELL-995 [110]	1	63,917	147,465	[15]
	DDB14 [110]	1	9,203	44,561	[134]
Recommendation Systems	Flixster [48]	1	6,000	26,173	[113; 114]
	Douban [71]	1	6,000	136,891	[113; 114]
	Yahoo [25]	1	6,000	5,335	[113; 114]
	MovieLens [79]	1	82,000	10,000,054	[113; 114]
	Ciao [101]	1	26,678	167,320	[66]
	Epinions [101]	1	35,079	225,579	[66]
	Filmtrust [34]	1	3,579	35,500	[66]
Transportation	Brazil [90]	1	131	1,038	[77]
	Europe [90]	1	399	5,995	[77]
	PeMSD4 [82]	1	307	-	[132; 12]
	PeMSD7 [82]	1	288	-	[132]
	PeMSD8 [82]	1	288	-	[12]
	PEMS-BAY [82]	1	325	-	[78; 12]
	METR-LA [47]	1	207	-	[78; 12]
Healthcare	ABIDE [23]	1	1,029	-	[88]
	ADNI [89]	1	911	-	[88]

ever, graph machine learning in a federated setting still faces a series of new challenges and therefore attracts massive attention from both researchers and practitioners. In this paper, we introduce the concepts of two problem settings in FGML. Then we review the current techniques under each setting in detail and introduce applications of FGML from different domains in the real world. We also summarize open graph datasets and platforms in FGML. In the end, some promising future directions are provided.

9. ACKNOWLEDGEMENTS

10. REFERENCES

[1] A. Acar, H. Aksu, A. S. Uluagac, and M. Conti. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys*, 2018.

[2] H. C. Bayram and I. Rekik. A federated multigraph integration approach for connectional brain template learning. In *ML-CDS*, 2021.

[3] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. Kiddon, J. Konečný, S. Mazzocchi, B. McMahan, et al. Towards federated learning at scale: System design. In *MLSys*, 2019.

[4] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In *NIPS*, 2013.

[5] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis, and W. Shi. Federated learning of predictive models from federated electronic health records. *International Journal of Medical Informatics*, 2018.

[6] L. Cai, J. Li, J. Wang, and S. Ji. Line graph neural networks for link prediction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.

[7] D. Caldarola, M. Mancini, F. Galasso, M. Ciccone, E. Rodola, and B. Caputo. Cluster-driven graph federated learning over multiple domains. In *CVPR Workshops*, 2021.

[8] S. Caldas, S. M. K. Duddu, P. Wu, T. Li, J. Konečný, H. B. McMahan, V. Smith, and A. Talwalkar. Leaf: A benchmark for federated settings. In *NeurIPS Workshops*, 2019.

[9] A. Chakravarty, A. Kar, R. Sethuraman, and D. Sheet. Federated learning for site aware chest radiograph screening. In *ISBI*, 2021.

[10] C. Chen, W. Hu, Z. Xu, and Z. Zheng. Fedgl: federated graph learning framework with global self-supervision. *arXiv preprint arXiv:2105.03170*, 2021.

[11] F. Chen, P. Li, T. Miyazaki, and C. Wu. Fedgraph: Federated graph learning with intelligent sampling. *IEEE Transactions on Parallel and Distributed Systems*, 2021.

[12] F. Chen, G. Long, Z. Wu, T. Zhou, and J. Jiang. Personalized federated learning with graph. *arXiv preprint arXiv:2203.00829*, 2022.

[13] H. Chen, L. Wang, Y. Lin, C.-C. M. Yeh, F. Wang, and H. Yang. Structured graph convolutional networks with stochastic masks for recommender systems. In *SIGIR*, 2021.

[14] J. Chen, G. Huang, H. Zheng, S. Yu, W. Jiang, and C. Cui. Graph-fraudster: Adversarial attacks on graph neural network-based vertical federated learning. *IEEE Transactions on Computational Social Systems*, 2022.

[15] M. Chen, W. Zhang, Z. Yuan, Y. Jia, and H. Chen. Fede: Embedding knowledge graphs in federated setting. In *IJCKG*, 2021.

[16] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, 2020.

[17] T.-H. Cheung, W. Dai, and S. Li. Fedsgc: Federated simple graph convolution for node classification. In *IJCAI Workshops*, 2021.

[18] K. Cho, B. van Merrienboer, C. Gucehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*, 2014.

[19] S. Cui, W. Pan, J. Liang, C. Zhang, and F. Wang. Addressing algorithmic disparity and performance inconsistency in federated learning. In *NeurIPS*, 2021.

[20] K. H. T. Dam, C.-H. B. Van Ouytsel, and A. Legay. Symbolic analysis meets federated learning to enhance malware identifier. *arXiv preprint arXiv:2204.14159*, 2022.

[21] D. Daza, M. Cochez, and P. Groth. Inductive entity representations from text via link prediction. In *WWW*, 2021.

[22] T. Dettmers, P. Minervini, P. Stenetorp, and S. Riedel. Convolutional 2d knowledge graph embeddings. In *AAAI*, 2018.

[23] A. Di Martino, C.-G. Yan, Q. Li, E. Denio, F. X. Castellanos, K. Alaerts, J. S. Anderson, M. Assaf, S. Y. Bookheimer, M. Dapretto, et al. The autism brain imaging data exchange: Towards a large-scale evaluation of the intrinsic brain architecture in autism. *Molecular psychiatry*, 2014.

[24] C. T. Dinh, T. T. Vu, N. H. Tran, M. N. Dao, and H. Zhang. A new look and convergence rate of federated multi-task learning with laplacian regularization. *arXiv preprint arXiv:2102.07148*, 2021.

[25] G. Dror, N. Koenigstein, Y. Koren, and M. Weimer. The yahoo! music dataset and kdd-cup'11. In *KDD Cup 2011*, 2012.

[26] Y. H. Ezzeldin, S. Yan, C. He, E. Ferrara, and S. Avestimehr. Fairfed: Enabling group fairness in federated learning. In *NeurIPS Workshops*, 2021.

[27] M. Fey and J. E. Lenssen. Fast graph representation learning with pytorch geometric. In *ICLR Workshops*, 2019.

[28] C. Finn, P. Abbeel, and S. Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.

[29] C. Gao, X. Wang, X. He, and Y. Li. Graph neural networks for recommender system. In *WSDM*, 2022.

[30] R. C. Geyer, T. Klein, and M. Nabi. Differentially private federated learning: A client level perspective. In *NIPS Workshops*, 2017.

[31] A. Ghosh, J. Chung, D. Yin, and K. Ramchandran. An efficient framework for clustered federated learning. In *NeurIPS*, 2020.

[32] C. L. Giles, K. D. Bollacker, and S. Lawrence. Citeseer: An automatic citation indexing system. In *DL*, 1998.

[33] Z. Guan, Y. Li, Z. Xue, Y. Liu, H. Gao, and Y. Shao. Federated graph neural network for cross-graph node classification. In *CCIS*, 2021.

[34] G. Guo, J. Zhang, and N. Yorke-Smith. A novel bayesian similarity measure for recommender systems. In *IJCAI*, 2013.

[35] M. B. Gurbuz and I. Rekik. Deep graph normalizer: A geometric deep learning approach for estimating connectional brain templates. In *MICCAI*, 2020.

[36] S. Gurukar, S. Parthasarathy, R. Ramnath, C. Calder, and S. Moosavi. Locationtrails: A federated approach to learning location embeddings. In *ASONAM*, 2021.

[37] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *NIPS*, 2017.

[38] M. Hang, J. Neville, and B. Ribeiro. A collective learning framework to boost gnn expressiveness for node classification. In *ICML*, 2021.

[39] C. He, K. Balasubramanian, E. Ceyani, C. Yang, H. Xie, L. Sun, L. He, L. Yang, P. S. Yu, Y. Rong, et al. Fedgraphnn: A federated learning system and benchmark for graph neural networks. In *ICLR Workshops*, 2021.

[40] C. He, E. Ceyani, K. Balasubramanian, M. Annavaram, and S. Avestimehr. Spreadgnn: Serverless multi-task federated learning for graph neural networks. In *AAAI*, 2022.

[41] C. He, S. Li, J. So, X. Zeng, M. Zhang, H. Wang, X. Wang, P. Vepakomma, A. Singh, H. Qiu, et al. Fedml: A research library and benchmark for federated machine learning. In *NeurIPS*, 2020.

[42] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[43] K. Hu, J. Wu, Y. Li, M. Lu, L. Weng, and M. Xia. Fedgcn: Federated learning-based graph convolutional networks for non-euclidean spatial data. *Mathematics*, 2022.

[44] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*, 2020.

[45] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.

[46] Y. Huang, L. Chu, Z. Zhou, L. Wang, J. Liu, J. Pei, and Y. Zhang. Personalized cross-silo federated learning on non-iid data. In *AAAI*, 2021.

[47] H. V. Jagadish, J. Gehrke, A. Labrinidis, Y. Papakonstantinou, J. M. Patel, R. Ramakrishnan, and C. Shahabi. Big data and its technical challenges. *Communications of the ACM*, 2014.

[48] M. Jamali and M. Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In *RecSys*, 2010.

[49] M. Jiang, T. Jung, R. Karl, and T. Zhao. Federated dynamic graph neural networks with secure aggregation for video-based distributed surveillance. *ACM Transactions on Intelligent Systems and Technology*, 2022.

[50] W. Jin, Y. Ma, X. Liu, X. Tang, S. Wang, and J. Tang. Graph structure learning for robust graph neural networks. In *KDD*, 2020.

[51] D. Johnson, H. Larochelle, and D. Tarlow. Learning graph structure with a finite-state automaton layer. In *NeurIPS*, 2020.

[52] J. Jordon, J. Yoon, and M. Van Der Schaar. Pategan: Generating synthetic data with differential privacy guarantees. In *ICLR*, 2018.

[53] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 2021.

[54] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017.

[55] A. Lalitha, O. C. Kilinc, T. Javidi, and F. Koushanfar. Peer-to-peer federated learning on graphs. *arXiv preprint arXiv:1901.11173*, 2019.

[56] M. Li and Z. Zhu. Spatial-temporal fusion graph neural networks for traffic flow forecasting. In *AAAI*, 2021.

[57] T. Li, S. Hu, A. Beirami, and V. Smith. Ditto: Fair and robust federated learning through personalization. In *ICML*, 2021.

[58] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 2020.

[59] W. Li and S. Wang. Federated meta-learning for spatial-temporal prediction. *Neural Computing and Applications*, 2022.

[60] J. Lin and S. Wong. A new directed divergence measure and its characterization. *International Journal of General System*, 1990.

[61] Y. Lin, C. Chen, C. Chen, and L. Wang. Improving federated relational data modeling via basis alignment and weight penalty. *arXiv preprint arXiv:2011.11369*, 2020.

[62] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally. Deep gradient compression: Reducing the communication bandwidth for distributed training. In *ICLR*, 2018.

[63] R. Liu and H. Yu. Federated graph neural networks: Overview, techniques and challenges. *arXiv preprint arXiv:2202.07256*, 2022.

[64] T. Liu, P. Li, and Y. Gu. Glint: Decentralized federated graph learning with traffic throttling and flow scheduling. In *IWQOS*, 2021.

[65] Y. Liu, Y. Zheng, D. Zhang, H. Chen, H. Peng, and S. Pan. Towards unsupervised deep graph structure learning. In *WWW*, 2022.

[66] Z. Liu, L. Yang, Z. Fan, H. Peng, and P. S. Yu. Federated social recommendation with graph neural network. *ACM Transactions on Intelligent Systems and Technology*, 2021.

[67] S. Lonare and R. Bhramaramba. Federated approach for privacy-preserving traffic prediction using graph convolutional network. *Journal of Shanghai Jiaotong University (Science)*, 2021.

[68] G. Lou, Y. Liu, T. Zhang, and X. Zheng. Stfl: A temporal-spatial federated learning framework for graph neural networks. In *AAAI Workshops*, 2021.

[69] S. Lu, Y. Zhang, and Y. Wang. Decentralized federated learning for electronic health records. In *CISS*, 2020.

[70] S. Lu, Y. Zhang, Y. Wang, and C. Mack. Learn electronic health records by fully decentralized federated learning. In *NeurIPS Workshops*, 2019.

[71] H. Ma, D. Zhou, C. Liu, M. R. Lyu, and I. King. Recommender systems with social regularization. In *WSDM*, 2011.

[72] O. Mahmood, E. Mansimov, R. Bonneau, and K. Cho. Masked graph modeling for molecule generation. *Nature communications*, 2021.

[73] X. Mai, Q. Fu, and Y. Chen. Packet routing with graph attention multi-agent reinforcement learning. In *GLOBECOM*, 2021.

[74] D. Manu, Y. Sheng, J. Yang, J. Deng, T. Geng, A. Li, C. Ding, W. Jiang, and L. Yang. Fl-disco: Federated generative adversarial network for graph-based molecule drug discovery. In *ICCAD*, 2021.

[75] A. K. McCallum, K. Nigam, J. Rennie, and K. Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 2000.

[76] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AISTATS*, 2017.

[77] G. Mei, Z. Guo, S. Liu, and L. Pan. Sgnn: A graph neural network based federated learning approach by hiding structure. In *BigData*, 2019.

[78] C. Meng, S. Rambhatla, and Y. Liu. Cross-node federated graph neural network for spatio-temporal data modeling. In *KDD*, 2021.

[79] B. N. Miller, I. Albert, S. K. Lam, J. A. Konstan, and J. Riedl. Movielens unplugged: Experiences with an occasionally connected recommender system. In *ACM IUI*, 2003.

[80] C. Morris, N. M. Kriege, F. Bause, K. Kersting, P. Mutzel, and M. Neumann. Tudataset: A collection of benchmark datasets for learning with graphs. In *ICML Workshops*, 2020.

[81] X. Ni, X. Xu, L. Lyu, C. Meng, and W. Wang. A vertical federated learning framework for graph convolutional network. *arXiv preprint arXiv:2106.11593*, 2021.

[82] C. D. of Transportation. California department of transportation. https://pems.dot.ca.gov/.

[83] N. L. Olsen, B. Markussen, and L. L. Raket. Simultaneous inference for misaligned multivariate functional data. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 2018.

[84] G. Panagopoulos, G. Nikolentzos, and M. Vazirgiannis. Transfer graph neural networks for pandemic forecasting. In *AAAI*, 2021.

[85] J. Payne and A. Kundu. Towards deep federated defenses against malware in cloud ecosystems. In *TPS-ISA*, 2019.

[86] Y. Pei, R. Mao, Y. Liu, C. Chen, S. Xu, F. Qiang, and B. E. Tech. Decentralized federated graph neural networks. In *IJCAI Workshops*, 2021.

[87] H. Peng, H. Li, Y. Song, V. Zheng, and J. Li. Differentially private federated knowledge graphs embedding. In *CIKM*, 2021.

[88] L. Peng, N. Wang, N. Dvornek, X. Zhu, and X. Li. Fedni: Federated graph learning with network inpainting for population-based disease prediction. *IEEE Transactions on Medical Imaging*, 2022.

[89] R. C. Petersen, P. Aisen, L. A. Beckett, M. Donohue, A. Gamst, D. J. Harvey, C. Jack, W. Jagust, L. Shaw, A. Toga, et al. Alzheimer's disease neuroimaging initiative (adni): clinical characterization. *Neurology*, 2010.

[90] L. F. Ribeiro, P. H. Saverese, and D. R. Figueiredo. struc2vec: Learning node representations from structural identity. In *KDD*, 2017.

[91] E. Rizk and A. H. Sayed. A graph federated architecture with privacy preserving learning. In *SPAWC*, 2021.

[92] S. Ru, B. Zhang, Y. Jie, C. Zhang, L. Wei, and C. Gu. Graph neural networks for privacy-preserving recommendation with secure hardware. In *NaNA*, 2021.

[93] A. Sarkar, T. Sen, and A. K. Roy. Grafehty: Graph neural network using federated learning for human activity recognition. In *ICMLA*, 2021.

[94] F. Sattler, K.-R. Müller, and W. Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.

[95] M. Schlichtkrull, T. N. Kipf, P. Bloem, R. v. d. Berg, I. Titov, and M. Welling. Modeling relational data with graph convolutional networks. In *ESWC*, 2018.

[96] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad. Collective classification in network data. *AI magazine*, 2008.

[97] A. Shamsian, A. Navon, E. Fetaya, and G. Chechik. Personalized federated learning using hypernetworks. In *ICML*, 2021.

[98] O. Shchur, M. Mumme, A. Bojchevski, and S. Gunnemann. Pitfalls of graph neural network evaluation. *arXiv preprint arXiv:1811.05868*, 2018.

[99] T. Suzumura, Y. Zhou, N. Baracaldo, G. Ye, K. Houck, R. Kawahara, A. Anwar, L. L. Stavarache, Y. Watanabe, P. Loyola, et al. Towards federated graph learning for collaborative financial crimes detection. In *NeurIPS Workshops*, 2019.

[100] A. Z. Tan, H. Yu, L. Cui, and Q. Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

[101] J. Tang, H. Gao, and H. Liu. mtrust: Discerning multifaceted trust in a connected world. In *WSDM*, 2012.

[102] J. Tang, J. Zhang, L. Yao, J. Li, L. Zhang, and Z. Su. Arnetminer: Extraction and mining of academic social networks. In *KDD*, 2008.

[103] Y. Tang, R. Jia, X. Zhou, Z. Li, H. Jin, and C. Zhao. Federated learning of user mobility anomaly based on graph attention networks. In *ICCC*, 2021.

[104] A. Thakur, P. Sharma, and D. A. Clifton. Dynamic neural graphs based federated reptile for semi-supervised multi-tasking in healthcare applications. *IEEE Journal of Biomedical and Health Informatics*, 2021.

[105] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. In *ICLR*, 2018.

[106] C. Villani. *Optimal transport: old and new*. 2009.

[107] P. Voigt and A. Von dem Bussche. The eu general data protection regulation (gdpr). *A Practical Guide, 1st Ed., Cham: Springer International Publishing*, 2017.

[108] B. Wang, A. Li, H. Li, and Y. Chen. Graphfl: A federated learning framework for semi-supervised node classification on graphs. *arXiv preprint arXiv:2012.04187*, 2020.

[109] C. Wang, B. Chen, G. Li, and H. Wang. Flagcns: Federated learning framework for automatic graph convolutional network search. *arXiv preprint arXiv:2104.04141*, 2021.

[110] H. Wang, H. Ren, and J. Leskovec. Relational message passing for knowledge graph completion. In *KDD*, 2021.

[111] Z. Wang, W. Kuang, Y. Xie, L. Yao, Y. Li, B. Ding, and J. Zhou. Federatedscope-gnn: Towards a unified, comprehensive and efficient package for federated graph learning. In *KDD*, 2022.

[112] Z. Wang, R. Wen, X. Chen, S. Cao, S.-L. Huang, B. Qian, and Y. Zheng. Online disease diagnosis with inductive heterogeneous graph convolutional networks. In *WWW*, 2021.

[113] C. Wu, F. Wu, Y. Cao, Y. Huang, and X. Xie. Fedgnn: Federated graph neural network for privacy-preserving recommendation. In *ICML Workshops*, 2021.

[114] C. Wu, F. Wu, L. Lyu, T. Qi, Y. Huang, and X. Xie. A federated graph neural network framework for privacy-preserving personalization. *Nature Communications*, 2022.

[115] Z. Wu, B. Ramsundar, E. N. Feinberg, J. Gomes, C. Geniesse, A. S. Pappu, K. Leswing, and V. Pande. Moleculenet: a benchmark for molecular machine learning. *Chemical science*, 2018.

[116] Z. Wu, X. Wu, and Y. Long. Multi-level federated graph learning and self-attention based personalized wi-fi indoor fingerprint localization. *IEEE Communications Letters*, 2022.

[117] H. Xie, J. Ma, L. Xiong, and C. Yang. Federated graph classification over non-iid graphs. In *NeurIPS*, 2021.

[118] Y. Xie, Z. Wang, D. Chen, D. Gao, L. Yao, W. Kuang, Y. Li, B. Ding, and J. Zhou. Federatedscope: A comprehensive and flexible federated learning platform via message passing. *arXiv preprint arXiv:2204.05011*, 2022.

[119] H. Xing, O. Simeone, and S. Bi. Decentralized federated learning via sgd over wireless d2d networks. In *SPAWC*, 2020.

[120] P. Xing, S. Lu, L. Wu, and H. Yu. Big-fed: Bilevel optimization enhanced graph-aided federated learning. In *ICML Workshops*, 2021.

[121] J. Xu, R. Wang, K. Liang, and S. Picek. More is better (mostly): On the backdoor attacks in federated graph neural networks. *arXiv preprint arXiv:2202.03195*, 2022.

[122] K. Xu, W. Hu, J. Leskovec, and S. Jegelka. How powerful are graph neural networks? In *ICLR*, 2019.

[123] Q. Yang, Y. Liu, T. Chen, and Y. Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology*, 2019.

[124] M. Ye, J. Zhang, Z. Guo, and H. J. Chao. Federated traffic engineering with supervised learning in multi-region networks. In *ICNP*, 2021.

[125] X. Yuan, J. Chen, J. Yang, N. Zhang, T. Yang, T. Han, and A. Taherkordi. Fedstn: Graph representation driven federated learning for edge computing enabled urban traffic flow prediction. *IEEE Transactions on Intelligent Transportation Systems*, 2022.

[126] W. W. Zachary. An information flow model for conflict and fission in small groups. *Journal of anthropological research*, 1977.

[127] S. Zehtabi, S. Hosseinalipour, and C. G. Brinton. Decentralized event-triggered federated learning with heterogeneous communication thresholds. *arXiv preprint arXiv:2204.03726*, 2022.

[128] T. Zeng, J. Guo, K. J. Kim, K. Parsons, P. Orlik, S. Di Cairano, and W. Saad. Multi-task federated learning for traffic prediction and its application to route planning. In *IV*, 2021.

[129] B. Zhang, M. Luo, S. Feng, Z. Liu, J. Zhou, and Q. Zheng. Ppsgcn: A privacy-preserving subgraph sampling based distributed gcn training method. *arXiv preprint arXiv:2110.12906*, 2021.

[130] C. Zhang, L. Cui, S. Yu, and J. James. A communication-efficient federated learning scheme for iot-based traffic forecasting. *IEEE Internet of Things Journal*, 2021.

[131] C. Zhang, S. Zhang, J. James, and S. Yu. Fastgnn: A topological information protected federated learning approach for traffic speed forecasting. *IEEE Transactions on Industrial Informatics*, 2021.

[132] C. Zhang, S. Zhang, S. Yu, and J. James. Graph-based traffic forecasting via communication-efficient federated learning. In *WCNC*, 2022.

[133] H. Zhang, T. Shen, F. Wu, M. Yin, H. Yang, and C. Wu. Federated graph learning–a position paper. *arXiv preprint arXiv:2105.11099*, 2021.

[134] K. Zhang, Y. Wang, H. Wang, L. Huang, C. Yang, and L. Sun. Efficient federated learning on knowledge graphs via privacy-preserving relation embedding aggregation. In *ACL Workshops*, 2022.

[135] K. Zhang, C. Yang, X. Li, L. Sun, and S. M. Yiu. Subgraph federated learning with missing neighbor generation. In *NeurIPS*, 2021.

[136] X. Zhang, C. Huang, Y. Xu, L. Xia, P. Dai, L. Bo, J. Zhang, and Y. Zheng. Traffic flow forecasting with spatial-temporal graph diffusion network. In *AAAI*, 2021.

[137] Z. Zhang, Q. Liu, H. Wang, C. Lu, and C.-K. Lee. Motif-based graph self-supervised learning for molecular property prediction. In *NeurIPS*, 2021.

[138] T. Zhao, Y. Liu, L. Neves, O. Woodford, M. Jiang, and N. Shah. Data augmentation for graph neural networks. In *AAAI*, 2021.

[139] T. Zhao, X. Zhang, and S. Wang. Graphsmote: Imbalanced node classification on graphs with graph neural networks. In *WSDM*, 2021.

[140] L. Zheng, J. Zhou, C. Chen, B. Wu, L. Wang, and B. Zhang. Asfgnn: Automated separated-federated graph neural network. *Peer-to-Peer Networking and Applications*, 2021.

[141] Y. Zheng, L. Capra, O. Wolfson, and H. Yang. Urban computing: Concepts, methodologies, and applications. *ACM Transactions on Intelligent Systems and Technology*, 2014.

[142] J. Zhou, C. Chen, L. Zheng, H. Wu, J. Wu, X. Zheng, B. Wu, Z. Liu, and L. Wang. Vertically federated graph neural network for privacy-preserving node classification. In *IJCAI*, 2022.

[143] H. Zhu, J. Xu, S. Liu, and Y. Jin. Federated learning on non-iid data: A survey. *Neurocomputing*, 2021.

[144] W. Zhu, A. White, and J. Luo. Federated learning of molecular properties in a heterogeneous setting. *arXiv preprint arXiv:2109.07258*, 2021.

[145] Z. Zhu, J. Zhu, J. Liu, and Y. Liu. Federated bandit: A gossiping approach. In *SIGMETRICS*, 2021.

[146] A. Ziller, A. Trask, A. Lopardo, B. Szymkow, B. Wagner, E. Bluemke, J.-M. Nounahon, J. Passerat-Palmbach, K. Prakash, N. Rose, et al. Pysyft: A library for easy federated learning. In *Federated Learning Systems*. 2021.