
SpreadGNN: Serverless Multi-task Federated Learning for Graph Neural Networks

Chaoyang He^{*1} Emir Ceyani^{*1} Keshav Balasubramanian^{*1} Murali Annavaram¹ Salman Avestimehr¹

Abstract

Graph Neural Networks (GNNs) are the first choice methods for graph machine learning problems thanks to their ability to learn state-of-the-art level representations from graph-structured data. However, centralizing a massive amount of real-world graph data for GNN training is prohibitive due to user-side privacy concerns, regulation restrictions, and commercial competition. Federated Learning is the de-facto standard for collaborative training of machine learning models over many distributed edge devices without the need for centralization. Nevertheless, training graph neural networks in a federated setting is vaguely defined and brings statistical and systems challenges. This work proposes SpreadGNN, a novel multi-task federated training framework capable of operating in the presence of partial labels and absence of a central server for the first time in the literature. SpreadGNN extends federated multi-task learning to realistic serverless settings for GNNs, and utilizes a novel optimization algorithm with a convergence guarantee, *Decentralized Periodic Averaging SGD (DPA-SGD)*, to solve decentralized multi-task learning problems. We empirically demonstrate the efficacy of our framework on a variety of non-I.I.D. distributed graph-level molecular property prediction datasets with partial labels. Our results show that SpreadGNN outperforms GNN models trained over a central server-dependent federated learning system, even in constrained topologies.

1. Introduction

Graph Neural Networks (GNNs) (Hamilton et al., 2017) are expressive models that can distill structural knowledge into highly representative embeddings. While graphs are the representation of choice in domains such as social networks (Ma et al., 2019; Bian et al., 2020), knowledge graphs for recommendation systems (Chen et al., 2020), in this work we focus on molecular graphs that are the core of drug discovery, molecular property prediction (Gilmer et al., 2017; Kearnes et al., 2016) and virtual screening (Wallach et al., 2015; Zheng et al., 2019). Molecular graphs differ from their more well known counterparts such as social network graphs. First, each molecule is a graph representation of the basic atoms and bonds that constitute the molecule and hence the size of the graph is small. Second, even though each graph may be small there are numerous molecules that are being developed continuously for varied use cases. Hence, what they lack in size they make up for in structural heterogeneity. Third, molecules can be labeled along multiple orthogonal dimensions. Since each graph has multiple labels the learning itself can be characterized as multi-task learning. For instance, whether a molecule has potentially harmful interaction with a diabetes drug, or whether that molecule can turn toxic under certain conditions are distinct labels. Molecular property analysis and labeling requires wet-lab experiments, which is time-consuming and resource-costly. As a consequence, many entities may only have partially labeled molecules even if they know the graph structure. Finally, molecules are coveted inventions and hence entities often possess proprietary graph representation that cannot be shared with other institutions for competitive and regulatory reasons. But training collectively over a private set of molecular graphs can have immense societal benefits such as accelerated drug discovery.

Federated Learning (FL) is a distributed learning paradigm that addresses this data isolation problem via collaborative training. In this paradigm, training is an act of collaboration between multiple clients (such as research institutions) without requiring centralized local data while providing a certain degree of user-level privacy (McMahan et al., 2017; Kairouz et al., 2019). However there are still challenges and shortcomings to training GNNs in a federated setting. As

^{*}Equal contribution ¹Viterbi School of Engineering, University of Southern California, Los Angeles, CA. Correspondence to: Chaoyang He <chaoyang.he@usc.edu>, Emir Ceyani <ceyani@usc.edu>.

This work was presented at the International Workshop on Federated Learning for User Privacy and Data Confidentiality in Conjunction with ICML 2021 (FL-ICML'21). This workshop does not have official proceedings and this paper is non-archival. Copyright 2021 by the author(s).

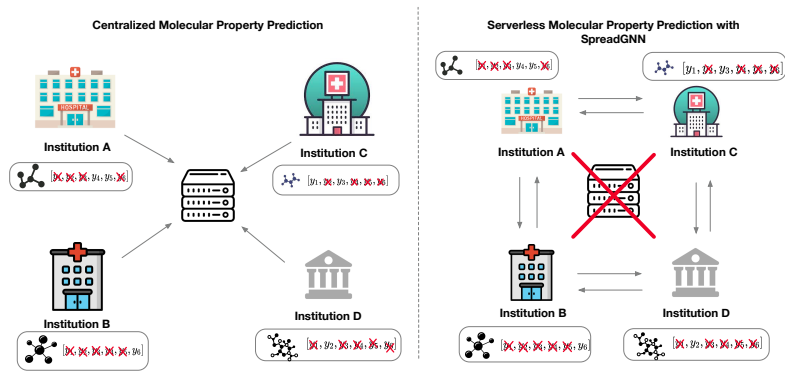


Figure 1: Serverless Multi-task Federated Learning for Graph Neural Networks.

shown in (He et al., 2021), federated GNNs perform poorly in a non-iid setting. This setting (Figure 1) is the typical case in molecular graphs since each owner may have different molecules and even when they have the same molecular graph each owner may have an incomplete set of labels for each molecule. The left half of Figure 1 shows a simpler case where all the clients can communicate through a central server. But in practice the presence of a central server is not feasible when multiple competing entities may want to collaboratively learn. The challenges are further compounded by the lack of a central server as shown in the right half of the Figure 1. Thus, it remains an open problem to design a federated learning framework for molecular GNNs, for a realistic setting, in which clients only have partial labels and one in which there is no reliance on a central server. This is the problem we seek to address in this work.

We propose a multitask federated learning framework called SpreadGNN that operates in the presence of partial labels and absence of a central server as shown in Figure 1. We use the word task and class interchangeably, with each label being composed of multiple tasks. First, we present a multitask learning (MTL) formulation to learn from partial labels. Second, in our MTL formulation, we utilize decentralized periodic averaging stochastic gradient descent (DPA-SGD) to solve the serverless MTL optimization problem, and also provide a theoretical guarantee on the convergence properties for DPA-SGD, which further verifies the rationality of our design.

We evaluate SpreadGNN on graph level molecular property prediction and regression tasks. We synthesize non-I.I.D. and partially labeled datasets by using curated data from the MoleculeNet (Wu et al., 2018) machine learning benchmark. With extensive experiments and analysis, we find that SpreadGNN can achieve even better performance than FedAvg (McMahan et al., 2016b), not only when all clients can communicate with each other, but also when clients are constrained to communicate with a subset of other clients.

We plan on publishing the source code of SpreadGNN as well as related datasets for future exploration.

2. SpreadGNN Framework

2.1. Preliminary: Federated Graph Neural Networks for Graph Level Learning

We seek to learn *graph level* representations in a federated learning setting over decentralized graph datasets located in edge servers which cannot be centralized for training due to **privacy and regulation restrictions** (He et al., 2021). For instance, compounds in molecular trials (Rong et al., 2020b) may not be shared across entities because of intellectual property or regulatory concerns. Under this setting, we assume that there are K clients in the FL network, and the k^{th} client has its own dataset $\mathcal{D}^{(k)} := \left\{ \left(G_i^{(k)}, y_i^{(k)} \right) \right\}_{i=1}^{N^{(k)}}$, where $G_i^{(k)} = (\mathcal{V}_i^{(k)}, \mathcal{E}_i^{(k)})$ is the i^{th} graph sample in $\mathcal{D}^{(k)}$ with node & edge feature sets $\mathbf{X}^{(k)} = \left\{ \mathbf{x}_m^{(k)} \right\}_{m \in \mathcal{V}_i^{(k)}}$ and $\mathbf{Z}^{(k)} = \left\{ e_{m,n}^{(k)} \right\}_{m,n \in \mathcal{V}_i^{(k)}}$, $\mathbf{y}_i^{(k)}$ is the corresponding multi-class label of $G_i^{(k)}$, $N^{(k)}$ is the sample number in dataset $\mathcal{D}^{(k)}$, and $N = \sum_{k=1}^K N^{(k)}$.

Each client owns a GNN with a readout, to learn graph-level representations. We call this model of a GNN followed by a readout function, a *graph classifier*. Multiple clients are interested in collaborating to improve their GNN models without necessarily revealing their graph datasets. In this work, we build our theory upon the Message Passing Neural Network (MPNN) framework (Gilmer et al., 2017; Rong et al., 2020c) as most spatial GNN models (Kipf and Welling, 2016; Veličković et al., 2018; Hamilton et al., 2017) can be unified into this framework. The forward pass of an MPNN has two phases: a message-passing phase (Eq (1)) and an update phase (Eq (2)). For each client, we define the graph classifier with an L-layer GNN followed by a readout




Figure 2: Federated Graph MultiTask Learning Framework (FedGMTL).

function as follows:

$$\mathbf{m}_i^{(k,\ell+1)} = \text{AGG} \left(\left\{ M_{\theta}^{(k,\ell+1)} \left(\mathbf{h}_i^{(k,\ell)}, \mathbf{h}_j^{(k,\ell)}, \mathbf{e}_{i,j} \right) \mid j \in \mathcal{N}_i \right\} \right), \ell=0, \dots, L-1 \quad (1)$$

$$\mathbf{h}_i^{(k,\ell+1)} = U_{\Psi}^{(k,\ell+1)} \left(\mathbf{h}_i^{(k,\ell)}, \mathbf{m}_i^{(k,\ell+1)} \right), \ell=0, \dots, L-1 \quad (2)$$

$$\hat{\mathbf{y}}_i^{(k)} = \mathbf{R}_{\Phi_{\text{pool}}, \Phi_{\text{task}}} \left(\left\{ \mathbf{h}_j^{(k,L)} \mid j \in \mathcal{V}_i^{(k)} \right\} \right) \quad (3)$$

where $\mathbf{h}_i^{(k,0)} = \mathbf{x}_i^{(k)}$ is the k^{th} client’s node features, ℓ is the layer index, AGG is the aggregation function (e.g., in the GCN model (Kipf and Welling, 2016), the aggregation function is a simple SUM operation), and \mathcal{N}_i is the neighborhood set of node i . In Eq. (1), $M_{\theta}^{(k,\ell+1)}(\cdot)$ is the message generation function which takes the hidden state of current node \mathbf{h}_i , the hidden state of the neighbor node \mathbf{h}_j and the edge features $\mathbf{e}_{i,j}$ as inputs to gather and transform neighbors’ messages. In other words, M_{θ} combines a vertex’s hidden state with the edge and vertex data from its neighbors to generate a new message. $U_{\Psi}^{(k,\ell+1)}(\cdot)$ is the state update function that updates the model using the aggregated feature $\mathbf{m}_i^{(k,\ell+1)}$ as in Eq. (2). After propagating through L GNN layers, the final module of the graph classifier is a readout function $\mathbf{R}_{\Phi_{\text{pool}}, \Phi_{\text{task}}}(\cdot)$ which allows clients to predict a label for the graph, given node embeddings that are learned from Eq.(2). In general the readout is composed of two neural networks: the pooling function parameterized by Φ_{pool} ; and a task classifier parameterized by Φ_{task} . The role of the pooling function is to learn a single graph embedding given node embedding from Eq (2). The task classifier then uses the graph level embedding to predict a label for the graph.

To formulate GNN-based FL, using the model definition above, we define $\mathbf{W} = \{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task}}\}$ as the overall learnable weights. Note that \mathbf{W} is independent of graph structure as both the GNN and Readout parameters make no assumptions about the input graph. Thus, one can learn \mathbf{W} using a FL based approach. The overall FL task can be formulated as a distributed optimization problem as:

$$\min_{\mathbf{W}} F(\mathbf{W}) \stackrel{\text{def}}{=} \min_{\mathbf{W}} \sum_{k=1}^K \frac{N^{(k)}}{N} \cdot f^{(k)}(\mathbf{W}) \quad (4)$$

where $f^{(k)}(\mathbf{W}) = \frac{1}{N^{(k)}} \sum_{i=1}^{N^{(k)}} \mathcal{L}(\hat{\mathbf{y}}_i^{(k)}, \mathbf{y}_i^{(k)})$ is the k^{th} client’s local objective function that measures the local empirical risk over the heterogeneous graph dataset $\mathcal{D}^{(k)}$. \mathcal{L} is the loss function of the global graph classifier.

With such a formulation, it might seem like an optimization problem tailored for a FedAvg based optimizers (McMahan et al., 2016b; He et al., 2020a; Reddi et al., 2020). Unfortunately, in molecular graph settings this is not the case for the following reasons. (a) In our setting, clients belong to a decentralized, serverless topology. There is no central server that can average model parameters from the clients. (b) Clients in our setting possess incomplete labels, hence the dimensions of Φ_{task} can be different on different clients. For instance, a client may only have a partial toxicity labels for a molecule, while another client may have a molecule’s interaction property with another drug compound. Even with such incomplete information from each client, our learning task is interested in classifying each molecule across multiple label categories. Federated Multi-task Learning (FMTL) (Smith et al., 2017; Caldas et al., 2018) is a popular framework designed to deal with such issues. However, current approaches do not generalize to non-convex deep models such as graph neural networks. To combat these issues, we first introduce a centralized FMTL framework for graph neural networks in Section 2.2. But reliance on a central server may not be feasible in molecular graphs owned by multiple competing entities. Hence, we enhance this centralized FMTL to a serverless scenario in section 2.3, which we call SpreadGNN.

2.2. Federated Multi-Task Learning with Graph Neural Networks

Under the regularized MTL paradigm (Evgeniou and Pontil, 2004), we define the centralized federated graph MTL problem (FedGMTL) as follows:

$$\begin{aligned} \min_{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task}}} & \sum_{k=1}^K \frac{1}{N_k} \sum_{i=1}^{N_k} \mathcal{L}(\hat{\mathbf{y}}_i^{(k)}, \mathbf{y}_i^{(k)}) + \mathcal{R}(\mathbf{W}^{(k)}, \Omega), \\ \text{s.t.} & \quad \Omega \geq 0 \quad \text{and} \quad \text{Tr}(\Omega) = 1. \end{aligned} \quad (5)$$

where

$$\mathcal{R}(\mathbf{W}^{(k)}, \Omega) = \frac{1}{2} \lambda_1 \text{Tr}(\Phi_{\text{task}} \Omega^{-1} \Phi_{\text{task}}^T) + \frac{1}{2} \sum_{\chi \in \{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task}}\}} \lambda_{\chi} \|\chi\|_F^2 \quad (6)$$

is the bi-convex regularizer introduced in (Zhang and Yeung, 2012). The first term of the Eq. (5) models the summation of different empirical loss of each client, which is what Eq. (4) exactly tries to address. The second term serves as a task-relationship regularizer with $\Omega \in \mathbb{R}^{\mathcal{S} \times \mathcal{S}}$ being the covariance matrix for \mathcal{S} different tasks constraining the task weights $\Phi_{\text{task}} = [\Phi_{\text{task},1}, \dots, \Phi_{\text{task},\mathcal{S}}] \in \mathbb{R}^{d \times \mathcal{S}}$ through matrix trace $\text{Tr}(\Phi_{\text{task}} \Omega^{-1} \Phi_{\text{task}}^T)$. Recall that each client in our setting may only have a partial set of tasks in the labels of its training dataset, but still needs to make predictions for tasks it does not have in its labels during test time. This regularizer helps clients relate its own tasks to tasks in other clients. Intuitively, it determines how closely two tasks i and j are related. The closer $\Phi_{\text{task},i}$ and $\Phi_{\text{task},j}$ are, the larger $\Omega_{i,j}$ will be. If Ω is an identity matrix, then each node is independent to each other. But as our results show, there is often a strong correlation between different molecular properties. This compels us to use a federated learning model.

Figure 2-a depicts the FedGMTL framework where clients' graph classifier weights are using an FL server. While the above formulation enhances the FMTL with a constrained regularizer that can be used for GNN learning, we still need to solve the final challenge, which is to remove the reliance on a central server to perform the computations in Eq. (5). Therefore, we propose a Decentralized Graph Multi-Task Learning framework, SpreadGNN, and utilize our novel optimizer, *Decentralized Periodic Averaging SGD (DPA-SGD)* to extend the FMTL framework to our setting. The decentralized framework is shown in the Figure 2-c. Note that each client's learning task remains the same but the aggregation process differs in SpreadGNN.

2.3. SpreadGNN: Serverless Federated MTL for GNNs

To transition to the serverless setting, we introduce a novel optimization method called Decentralized Periodic Averaging SGD (DPA-SGD). The main idea of DPA-SGD is that each client applies SGD locally, and synchronizes all parameters with only its neighbors during a communication round that occurs every τ iterations. A decentralized system in which all clients are not necessarily connected to all other

clients also makes it impossible to maintain one single task covariance matrix Ω . Thus we propose using distinct covariance matrices Ω_k in each client that are efficiently updated using the exchange mechanism illustrated in Figure 3. We formalize this idea as follows. Consider one particular client m having task weights $\Phi_{\text{task},m} \in \mathbb{R}^{d \times \mathcal{S}_m}$ where \mathcal{S}_m is the number of tasks that client m has. Recall that the regularizer in Eq. (6) is how various clients learn to predict tasks that are not within its \mathcal{S}_m .

In the decentralized setting, we emphasize that clients can collectively learn the exhaustive set of tasks, even when clients may not have access to some of the classes in each multi-class label. That is and $\cup_i \mathcal{S}_i = \mathcal{S}$. Then, the new non-convex objective function can be defined as:

$$\begin{aligned} \min_{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task}}} & \sum_{k=1}^K \frac{1}{N_k} \left[\sum_{i=1}^{N_k} l(\hat{\mathbf{y}}_i^{(k)}(\mathbf{X}_k, \mathbf{Z}_k; \mathbf{W}_k), \mathbf{y}_i^k) + \frac{1}{2} \lambda_1 \text{Tr}(\Phi_{\text{task},\mathcal{M}_k} \Omega_k^{-1} \Phi_{\text{task},\mathcal{M}_k}^T) \right] \\ & + \frac{1}{2} \sum_{\chi \in \{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task}}\}} \lambda_{\chi} \|\chi\|_F^2, \\ \text{s.t.} & \quad \Omega_k \geq 0 \quad \text{and} \quad \text{Tr}(\Omega_k) = 1, \quad k = 1, 2, \dots, K. \end{aligned} \quad (7)$$

where $\mathbf{W}_k = \{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task},k}\}$ is the set of all learnable weights for client k , $\mathcal{M}_k = k \cup \mathcal{N}_k$ is the neighbor set for client k including itself. This gives rise to: $\Phi_{\text{task},\mathcal{M}_k} = [\Phi_{\text{task},1} \parallel \Phi_{\text{task},2} \parallel \dots \parallel \Phi_{\text{task},|\mathcal{M}_k|}] \in \mathbb{R}^{d \times |\mathcal{S}_{\mathcal{M}_k}|}$ which is the task weight matrix for client k and its neighbors and \parallel represents the row-wise union operation. The matrix $\Omega_k \in \mathbb{R}^{|\mathcal{S}_{\mathcal{M}_k}| \times |\mathcal{S}_{\mathcal{M}_k}|}$ represents the correlation amongst all the available tasks for the set \mathcal{M}_k .

To solve this non-convex problem, we apply the alternating optimization method presented in (Zhang and Yeung, 2012), where \mathbf{W}_k and Ω_k are updated in an alternative fashion.

Optimizing \mathbf{W}_k : For simplicity, we define $\underline{\Omega} = \{\Omega_1, \Omega_2, \dots, \Omega_K\}$ to represent the set of correlation matrices for all clients. Fixing $\underline{\Omega}$, we can use SGD to update \mathbf{W}_k jointly. Let $L = \sum_{k=1}^K \frac{1}{N_k} \sum_{i=1}^{N_k} l(\hat{\mathbf{y}}_i^{(k)}(\mathbf{X}_i^k, \mathbf{Z}_i^k; \mathbf{W}_k), \mathbf{y}_i^k)$. Then, our problem can then be reformulated as:

$$\begin{aligned} G(\mathbf{W}_k | \underline{\Omega}) = L + & \sum_{k=1}^K \frac{1}{N_k} \sum_{i=1}^{N_k} \left[\frac{1}{2} \lambda_1 \text{Tr}(\Phi_{\text{task},\mathcal{M}_k} \Omega_k^{-1} \Phi_{\text{task},\mathcal{M}_k}^T) \right] \\ & + \frac{1}{2} \sum_{\chi \in \{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task}}\}} \lambda_{\chi} \|\chi\|_F^2. \end{aligned} \quad (8)$$

where the summation in (8) is amongst all the nodes connected to node k . Then the gradient formulations for each node are:

$$\frac{\partial G(\mathbf{W}_k | \underline{\Omega})}{\partial \Phi_{\text{task},\mathcal{M}_k}} = \frac{\partial L}{\partial \Phi_{\text{task},\mathcal{M}_k}} + \lambda_1 \sum_{i=1}^{|\mathcal{M}_k|} \frac{1}{N_i} \Phi_{\text{task},\mathcal{M}_k} \Omega_i^{-1} + \lambda_2 \Phi_{\text{task},\mathcal{M}_k} \quad (9)$$

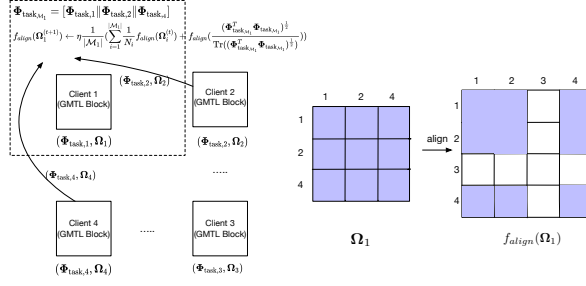


Figure 3: Decentralized Multi-Task Learning Correlation Matrix Exchanging Algorithm

$$\frac{\partial G(\mathbf{W}_k \| \underline{\Omega})}{\partial \chi} = \frac{\partial L}{\partial \chi} + \lambda_{\chi} \chi, \quad \forall \chi \in \{\theta, \Psi, \Phi_{pool}\} \quad (10)$$

Optimizing $\Omega_k \in \underline{\Omega}$: In (Zhang and Yeung, 2012), an analytical solution for Ω is equal to $\Omega^* = \frac{(\Phi_{task}^T \Phi_{task})^{\frac{1}{2}}}{\text{Tr}((\Phi_{task}^T \Phi_{task})^{\frac{1}{2}})}$. However, this solution is only applicable for the centralized case. This is because the missing central node forbids averaging parameters globally. So here we propose a novel way to update each $\Omega_k \in \underline{\Omega}$:

$$f_{align}(\Omega_k^{(t+1)}) \leftarrow \eta \frac{1}{|\mathcal{M}_k|} \left(\sum_{i=1}^{|\mathcal{M}_k|} \frac{1}{N_i} f_{align}(\Omega_i^{(t)}) + f_{align} \left(\frac{(\Phi_{task, \mathcal{M}_k}^T \Phi_{task, \mathcal{M}_k})^{\frac{1}{2}}}{\text{Tr}((\Phi_{task, \mathcal{M}_k}^T \Phi_{task, \mathcal{M}_k})^{\frac{1}{2}})} \right) \right) \quad (11)$$

The first averaging term can incorporate the nearby nodes correlation into its own. It should be noticed that each Ω_i may have a different dimension (different number of neighbors), so this averaging algorithm is based on the node-wise alignment in the global Ω . The f_{align} function is illustrated in Figure 3. Here, f_{align} operates on each client and the second term captures the new correlation between its neighbors as shown in Figure 3. We refer readers to the convergence analysis of DPA-SGD and the Algorithm 1 for a better view of SpreadGNN in the Appendix.

3. Experiments

3.1. Setup

Implementation. All experiments are conducted on a single GPU server equipped with 2 Nvidia Geforce GTX 1080Ti GPUs and an AMD Ryzen 7 3700X 8-Core Processor. Our models are built on top of the FedML framework (He et al., 2020b) and PyTorch (Paszke et al., 2019).

Multi-task Dataset. We use molecular datasets from the MoleculeNet (Wu et al., 2018) machine learning benchmark in our evaluation described in Table 2. The label for each molecule graph is a vector in which each element denotes a property of the molecule.

Non-I.I.D. Partition for Quantity and Label Skew. We

introduce non-I.I.D.ness in two additive ways. The first is a non-I.I.D. split of the training data based on quantity. Here we use a Dirichlet distribution parameterized by α to split the training data between clients. Specifically, the number of training samples present in each client is non-I.I.D. The second source of non-I.I.D.ness is a label masking scheme designed to represent the scenario in which different clients may possess partial labels as shown in Figure 1. More specifically, we mask out a subset of classes in each label on every client. In our experiments, the sets of unmasked classes across all clients are mutually exclusive and collectively exhaustive. This setting simulates a worst case scenario where no two clients share the same task. However our framework is just as applicable when there is label overlap. Such masking, introduces a class imbalance between the clients making the label distribution non-I.I.D. as well.

Network Topology. We first evaluate our framework in a complete topology in which all clients are connected to all other clients to measure the efficacy of our proposed regularizer. We then perform ablation studies on the number of neighbors of each client to stress test our framework in the more constrained setting¹.

3.2. Results

Our results summarized in Table 1 demonstrate that SpreadGNN (third column) outperforms a centralized federated learning system that uses FedAvg (first column) when all clients can communicate with all other clients. This shows that by using the combination of the task regularizer in equation 6 and the DPA-SGD algorithm, we can eliminate the dependence on a central server and enable clients to learn more effectively in the presence of missing molecular properties in their respective labels. Additionally, the results also show that our framework is agnostic to the type of GNN model being used in the molecular property prediction task since both GraphSAGE and GAT benefit from our framework. Our framework also works in the case a trusted central server is available (second column). The

¹A more comprehensive hyperparameter list for network topology and models can be found in the Appendix D.

presence of a trusted central server improves accuracy in a few scenarios. However, SpreadGNN provides competing performance in a more realistic setting where central server is not feasible.

3.2.1. ABLATION STUDIES

In this section, we provide some insights on how hyperparameters affect the results. More detailed analysis can be found in the Appendix E.

Task-Relationship Regularizer. Figure 4 illustrates the effect of λ_1 on both regression as well as classification tasks. Interestingly, regression is much more robust to variation in λ_1 while classification demands more careful tuning of λ_1 to achieve optimal model performance. This implies that the different properties in the regression task are more independent than the properties in the classification task.

Period. The communication period τ is another important hyperparameter in our framework. In general, our experiments suggest that a lower period is better, but this is not always the case. Figures 6 & 7 in the Appendix illustrate the effect of communication period on SIDER and Tox21 datasets. As we increase the communication period τ more, model performance decreases. However, selecting $\tau = 5$ can sometimes be better than averaging & exchanging each round. This indicates that, tuning τ is important for while controlling the tradeoff between the performance and the running time.

Network Topology. The network topology dictates how many neighbors each client can communicate with in a communication round. While Table 1 shows that SpreadGNN outperforms FedAvg in a complete topology, Figure 5 shows that our framework performs outperforms FedAvg even when clients are constrained to communicate with fewer neighbors. We can also see that it's not just $n_{neighbors}$ that matters, the topology in which clients are connected does too. When $n_{neighbors} = 2$, a ring topology outperforms a random topology, as a ring guarantees a path from any client to any other client. Thus, learning is shared indirectly between all clients. The same is not true in a random topology.

4. Conclusion

In this work, we propose SpreadGNN, a framework to train graph neural networks in a serverless federated system. We motivate our framework through a realistic setting, in which clients involved in molecular machine learning research cannot share data with each other due to privacy and competition. Moreover, we further recognize that clients may only possess partial labels of their data. Through experimental analysis, we show for the first time, that training graph neural networks in a federated manner does not re-

quire a centralized topology and that our framework can address the non-iidness in dataset size and label distribution across clients. SpreadGNN can outperform a central server depend baseline even when clients can only communicate with a few neighbors. To support our empirical results, we also provide a convergence analysis of the *DPA-SGD* optimization algorithm used by SpreadGNN.

References

- Tox21 challenge. <https://tripod.nih.gov/tox21/challenge/>, 2017.
- Amr Ahmed, Abhimanyu Das, and Alexander J. Smola. Scalable hierarchical multitask learning algorithms for conversion optimization in display advertising. In *Proceedings of the 7th ACM international conference on Web search and data mining*, pages 153–162. ACM, 2014.
- Inci M. Baytas, Ming Yan, Anil K. Jain, and Jiayu Zhou. Asynchronous multi-task learning. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pages 11–20. IEEE, 2016.
- Tian Bian, Xi Xiao, Tingyang Xu, Peilin Zhao, Wenbing Huang, Yu Rong, and Junzhou Huang. Rumor detection on social media with bi-directional graph convolutional networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 549–556, 2020.
- Léon Bottou, Frank E. Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- Sebastian Caldas, Virginia Smith, and Ameet Talwalkar. Federated Kernelized Multi-Task Learning. *The Conference on Systems and Machine Learning*, page 3, 2018.
- Chaochao Chen, Jamie Cui, Guanfeng Liu, Jia Wu, and Li Wang. Survey and open problems in privacy preserving knowledge graph: Merging, query, representation, completion and applications. *arXiv preprint arXiv:2011.10180*, 2020.
- Connor W Coley, Regina Barzilay, William H Green, Tommi S Jaakkola, and Klavs F Jensen. Convolutional embedding of attributed molecular graphs for physical property prediction. *Journal of chemical information and modeling*, 57(8):1757–1772, 2017.
- Henggang Cui, James Cipar, Qirong Ho, Jin Kyu Kim, Seunghak Lee, Abhimanu Kumar, Jinliang Wei, Wei Dai, Gregory R. Ganger, and Phillip B. Gibbons. Exploiting Bounded Staleness to Speed Up Big Data Analytics. In *USENIX Annual Technical Conference*, pages 37–48, 2014.

	GraphSAGE(Hamilton et al., 2017)			GAT(Veličković et al., 2018)		
	FedAvg	FedGMTL	SpreadGNN	FedAvg	FedGMTL	SpreadGNN
SIDER (4 clients)	0.582	0.629	0.5873	0.5857	0.61	0.6034
Tox21 (8 clients)	0.5548	0.6664	0.585	0.6035	0.6594	0.6056
MUV (8 clients)	0.6578	0.6856	0.703	0.709	0.6899	0.713
QM8 (8 clients)	0.02982	0.03624	0.02824	0.0392	0.0488	0.0315

Table 1: Results on the molecular property prediction task. SpreadGNN uses a complete topology: all clients communicated with each other. Communication period $\tau = 1$. We use Adam (Kingma and Ba, 2015) as the client optimizer in all of our experiments.



Figure 4: Effect of Task-Relationship Regularizer on Learning

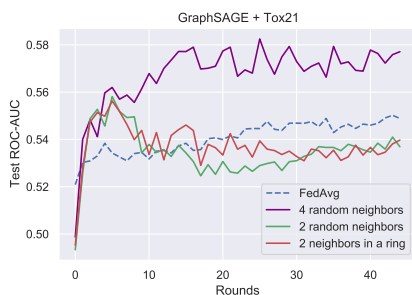


Figure 5: Effect of Topology on Graphsage + Tox21

Wei Dai, Yi Zhou, Nanqing Dong, Hao Zhang, and Eric P. Xing. Toward Understanding the Impact of Staleness in Distributed Machine Learning. *arXiv:1810.03264 [cs, stat]*, October 2018. URL <http://arxiv.org/abs/1810.03264>. arXiv: 1810.03264.

Jeffrey Dean, Greg S Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Quoc V Le, Mark Z Mao, Marc Aurelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, and Andrew Y Ng. Large Scale Distributed Deep Networks. page 11.

Sanghamitra Dutta, Gauri Joshi, Soumyadip Ghosh, Parijat Dube, and Priya Nagpurkar. Slow and Stale Gradients Can Win the Race: Error-Runtime Trade-offs in Distributed SGD. *arXiv preprint arXiv:1803.01113*, 2018.

David Duvenaud, Dougal Maclaurin, Jorge Aguilera-

Iparraquirre, Rafael Gómez-Bombarelli, Timothy Hirzel, Alán Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. *arXiv preprint arXiv:1509.09292*, 2015.

Theodoros Evgeniou and Massimiliano Pontil. Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 109–117, 2004.

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*, pages 1263–1272. PMLR, 2017.

Suyog Gupta, Wei Zhang, and Fei Wang. Model accuracy and runtime tradeoff in distributed deep learning: A systematic study. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pages 171–180. IEEE, 2016.

Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring network structure, dynamics, and function using networkx. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.

William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. *CoRR*, abs/1706.02216, 2017. URL <http://arxiv.org/abs/1706.02216>.

- 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, 2020a.
- Chaoyang He, Songze Li, Jinhyun So, Mi Zhang, Hongyi Wang, Xiaoyang Wang, Praneeth Vepakomma, Abhishek Singh, Hang Qiu, Li Shen, Peilin Zhao, Yan Kang, Yang Liu, Ramesh Raskar, Qiang Yang, Murali Annavaram, and Salman Avestimehr. Fedml: A research library and benchmark for federated machine learning. *arXiv preprint arXiv:2007.13518*, 2020b.
- Chaoyang He, Keshav Balasubramanian, Emir Ceyani, Yu Rong, Peilin Zhao, Junzhou Huang, Murali Annavaram, and Salman Avestimehr. Fedgraphnn: A federated learning system and benchmark for graph neural networks. *CoRR*, abs/2104.07145, 2021. URL <https://arxiv.org/abs/2104.07145>.
- Martin Jaggi, Virginia Smith, Martin Takáč, Jonathan Terhorst, Sanjay Krishnan, Thomas Hofmann, and Michael I. Jordan. Communication-efficient distributed dual coordinate ascent. In *Advances in neural information processing systems*, pages 3068–3076, 2014.
- Meng Jiang, Taeho Jung, Ryan Karl, and Tong Zhao. Federated dynamic gnn with secure aggregation. *arXiv preprint arXiv:2009.07351*, 2020.
- Zhanhong Jiang, Aditya Balu, Chinmay Hegde, and Soumik Sarkar. Collaborative deep learning in fixed topology networks. In *Advances in Neural Information Processing Systems*, pages 5904–5914, 2017.
- Peter H. Jin, Qiaochu Yuan, Forrest Iandola, and Kurt Keutzer. How to scale distributed deep learning? *arXiv preprint arXiv:1611.04581*, 2016.
- Xin Jin, Ping Luo, Fuzhen Zhuang, Jia He, and Qing He. Collaborating between local and global learning for distributed online multiple tasks. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*, pages 113–122. ACM, 2015.
- Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Keith Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *arXiv preprint arXiv:1912.04977*, 2019.
- Steven Kearnes, Kevin McCloskey, Marc Berndl, Vijay Pande, and Patrick Riley. Molecular graph convolutions: moving beyond fingerprints. *Journal of computer-aided molecular design*, 30(8):595–608, 2016.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.
- Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *CoRR*, abs/1609.02907, 2016. URL <http://arxiv.org/abs/1609.02907>.
- Johannes Klicpera, Janek Groß, and Stephan Günnemann. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*, 2020.
- Jakub Konečný, Brendan McMahan, and Daniel Ramage. Federated Optimization: Distributed Optimization Beyond the Datacenter. *arXiv:1511.03575 [cs, math]*, November 2015. URL <http://arxiv.org/abs/1511.03575>. arXiv: 1511.03575.
- Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated Learning: Strategies for Improving Communication Efficiency. *arXiv:1610.05492 [cs]*, October 2016. URL <http://arxiv.org/abs/1610.05492>. arXiv: 1610.05492.
- Michael Kuhn, Ivica Letunic, Lars Juhl Jensen, and Peer Bork. The sider database of drugs and side effects. *Nucleic acids research*, 44(D1):D1075–D1079, 2016.
- Greg Landrum. Rdkit: Open-source cheminformatics, 2006. URL <http://www.rdkit.org>.
- Mu Li, David G. Andersen, Jun Woo Park, Alexander J. Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J. Shekita, and Bor-Yiing Su. Scaling Distributed Machine Learning with the Parameter Server. In *OSDI*, volume 14, pages 583–598, 2014.
- Qimai Li, Zhichao Han, and Xiao-Ming Wu. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can Decentralized Algorithms Outperform Centralized Algorithms? A Case Study for Decentralized Parallel Stochastic Gradient Descent. *arXiv:1705.09056 [cs, math, stat]*, May 2017. URL <http://arxiv.org/abs/1705.09056>. arXiv: 1705.09056.
- Tao Lin, Sebastian U. Stich, and Martin Jaggi. Don't Use Large Mini-Batches, Use Local SGD. *arXiv:1808.07217 [cs, stat]*, August 2018. URL <http://arxiv.org/abs/1808.07217>. arXiv: 1808.07217.
- Sulin Liu, Sinno Jialin Pan, and Qirong Ho. Distributed multi-task relationship learning. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 937–946. ACM, 2017.

- Chenxin Ma, Virginia Smith, Martin Jaggi, Michael I. Jordan, Peter Richtárik, and Martin Takáč. Adding vs. Averaging in Distributed Primal-Dual Optimization. *arXiv:1502.03508 [cs]*, February 2015. URL <http://arxiv.org/abs/1502.03508>. arXiv:1502.03508.
- Jing Ma, Wei Gao, and Kam-Fai Wong. Detect rumors on twitter by promoting information campaigns with generative adversarial learning. In *The World Wide Web Conference*, pages 3049–3055, 2019.
- David Mateos-Núñez, Jorge Cortés, and Jorge Cortes. Distributed optimization for multi-task learning via nuclear-norm approximation. In *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, volume 48, pages 64–69, 2015.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017.
- H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data. *arXiv:1602.05629 [cs]*, February 2016a. URL <http://arxiv.org/abs/1602.05629>. arXiv:1602.05629.
- H. Brendan McMahan, Eider Moore, Daniel Ramage, and Blaise Agüera y Arcas. Federated learning of deep networks using model averaging. *CoRR*, abs/1602.05629, 2016b. URL <http://arxiv.org/abs/1602.05629>.
- Guangxu Mei, Ziyu Guo, Shijun Liu, and Li Pan. Sgnn: A graph neural network based federated learning approach by hiding structure. In *2019 IEEE International Conference on Big Data (Big Data)*, pages 2560–2568. IEEE, 2019.
- Chuizheng Meng, Sirisha Rambhatla, and Yan Liu. Cross-node federated graph neural network for spatio-temporal data modeling, 2021. URL https://openreview.net/forum?id=HWX5j6Bv_ih.
- Ioannis Mitliagkas, Ce Zhang, Stefan Hadjis, and Christopher Ré. Asynchrony begets momentum, with an application to deep learning. In *Communication, Control, and Computing (Allerton), 2016 54th Annual Allerton Conference on*, pages 997–1004. IEEE, 2016.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. *CoRR*, abs/1912.01703, 2019. URL <http://arxiv.org/abs/1912.01703>.
- Raghunathan Ramakrishnan, Mia Hartmann, Enrico Tapavicza, and O. Anatole von Lilienfeld. Electronic spectra from tddft and machine learning in chemical space. *The Journal of Chemical Physics*, 143(8):084111, Aug 2015. ISSN 1089-7690. doi: 10.1063/1.4928757. URL <http://dx.doi.org/10.1063/1.4928757>.
- Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in neural information processing systems*, pages 693–701, 2011.
- Sashank Reddi, Zachary Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and H Brendan McMahan. Adaptive federated optimization. *arXiv preprint arXiv:2003.00295*, 2020.
- David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- Sebastian G Rohrer and Knut Baumann. Maximum unbiased validation (muv) data sets for virtual screening based on pubchem bioactivity data. *Journal of chemical information and modeling*, 49(2):169–184, 2009.
- Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-supervised graph transformer on large-scale molecular data, 2020a.
- Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-supervised graph transformer on large-scale molecular data. *Advances in Neural Information Processing Systems*, 33, 2020b.
- Yu Rong, Tingyang Xu, Junzhou Huang, Wenbing Huang, Hong Cheng, Yao Ma, Yiqi Wang, Tyler Derr, Lingfei Wu, and Tengfei Ma. Deep graph learning: Foundations, advances and applications. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery; Data Mining, KDD '20*, page 3555–3556, New York, NY, USA, 2020c. Association for Computing Machinery. ISBN 9781450379984. doi: 10.1145/3394486.3406474. URL <https://doi.org/10.1145/3394486.3406474>.
- Kristof T Schütt, Farhad Arbabzadah, Stefan Chmiela, Klaus R Müller, and Alexandre Tkatchenko. Quantum-chemical insights from deep tensor neural networks. *Nature communications*, 8(1):1–8, 2017a.

- Kristof T Schütt, Pieter-Jan Kindermans, Huziel E Saucedo, Stefan Chmiela, Alexandre Tkatchenko, and Klaus-Robert Müller. Schnet: A continuous-filter convolutional neural network for modeling quantum interactions. *arXiv preprint arXiv:1706.08566*, 2017b.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. *Journal of Machine Learning Research*, 14(Feb):567–599, 2013.
- Virginia Smith, Simone Forte, Chenxin Ma, Martin Takac, Michael I. Jordan, and Martin Jaggi. CoCoA: A General Framework for Communication-Efficient Distributed Optimization. *arXiv:1611.02189 [cs]*, November 2016. URL <http://arxiv.org/abs/1611.02189>. arXiv:1611.02189.
- Virginia Smith, Chao-Kai Chiang, Maziar Sanjabi, and Ameet S. Talwalkar. Federated multi-task learning. In *Advances in Neural Information Processing Systems*, pages 4424–4434, 2017.
- Sebastian U. Stich. Local SGD Converges Fast and Communicates Little. September 2018. URL <https://openreview.net/forum?id=S1g2JnRcFX>.
- Toyotaro Suzumura, Yi Zhou, Natahalie Baracaldo, Guangnan Ye, Keith Houck, Ryo Kawahara, Ali Anwar, Lucia Larise Stavarache, Yuji Watanabe, Pablo Loyola, et al. Towards federated graph learning for collaborative financial crimes detection. *arXiv preprint arXiv:1909.12946*, 2019.
- Hanlin Tang, Shaoduo Gan, Ce Zhang, Tong Zhang, and Ji Liu. Communication Compression for Decentralized Training. *arXiv:1803.06443 [cs, stat]*, March 2018. URL <http://arxiv.org/abs/1803.06443>. arXiv:1803.06443.
- Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks, 2018.
- Izhar Wallach, Michael Dzamba, and Abraham Heifets. Atomnet: a deep convolutional neural network for bioactivity prediction in structure-based drug discovery. *arXiv preprint arXiv:1510.02855*, 2015.
- Binghui Wang, Ang Li, Hai Li, and Yiran Chen. Graphfl: A federated learning framework for semi-supervised node classification on graphs. *arXiv preprint arXiv:2012.04187*, 2020.
- Jialei Wang, Mladen Kolar, and Nathan Srebro. Distributed multi-task learning. In *Artificial Intelligence and Statistics*, pages 751–760, 2016.
- Jianyu Wang and Gauri Joshi. Cooperative SGD: A unified Framework for the Design and Analysis of Communication-Efficient SGD Algorithms. August 2018. URL <https://arxiv.org/abs/1808.07576v2>.
- 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.
- Kevin Yang, Kyle Swanson, Wengong Jin, Connor Coley, Philipp Eiden, Hua Gao, Angel Guzman-Perez, Timothy Hopper, Brian Kelley, Miriam Mathea, et al. Analyzing learned molecular representations for property prediction. *Journal of chemical information and modeling*, 59(8):3370–3388, 2019.
- Tianbao Yang. Trading computation for communication: Distributed stochastic dual coordinate ascent. In *Advances in Neural Information Processing Systems*, pages 629–637, 2013.
- Tianbao Yang, Shenghuo Zhu, Rong Jin, and Yuanqing Lin. Analysis of distributed stochastic dual coordinate ascent. *arXiv preprint arXiv:1312.1031*, 2013.
- Hao Yu, Sen Yang, and Shenghuo Zhu. Parallel Restarted SGD with Faster Convergence and Less Communication: Demystifying Why Model Averaging Works for Deep Learning. *arXiv:1807.06629 [cs, math]*, July 2018. URL <http://arxiv.org/abs/1807.06629>. arXiv:1807.06629.
- Yu Zhang and Dit-Yan Yeung. A convex formulation for learning task relationships in multi-task learning. *arXiv preprint arXiv:1203.3536*, 2012.
- Liangzhen Zheng, Jingrong Fan, and Yuguang Mu. Onionnet: a multiple-layer intermolecular-contact-based convolutional neural network for protein–ligand binding affinity prediction. *ACS omega*, 4(14):15956–15965, 2019.
- Jun Zhou, Chaochao Chen, Longfei Zheng, Xiaolin Zheng, Bingzhe Wu, Ziqi Liu, and Li Wang. Privacy-preserving graph neural network for node classification. *arXiv preprint arXiv:2005.11903*, 2020.

Appendix

This section includes additional information that a reader might find useful. Apart from the proof of Theorem , we include the algorithm sketch for SpreadGNN using DPA-SGD, a list of related works, a more detailed description of the datasets we used, hyperparameter configurations used in our experiments and additional ablation studies on communication period and network topology.

A. Algorithm Sketch

Algorithm 1 SpreadGNN : Serverless Multi-task Federated Learning for Graph Neural Networks

Require: initial parameters for each node $\mathbf{W}_k^{(t=0)} = \{\theta^{(t=0)}, \Psi^{(t=0)}, \Phi_{\text{pool}}^{(t=0)}, \Phi_{\text{task},k}^{(t=0)}\}$ and $\underline{\Omega}^{(t=0)} = (\Omega_1^{(t=0)}, \Omega_2^{(t=0)}, \dots, \Omega_K^{(t=0)})$; learning rate η ; maximum number of global iterations T , maximum number of client epochs E ; communication period τ .

- 1: **for** all nodes: $k = 1, 2, \dots, K$ **in parallel do**
- 2: **for** $t = 1$ to T **do do**
- 3: **for** $k = 1$ to E (epoch loop) **do do**
- 4: **for** $m \in MB$ (mini-batch loop) **do do**
- 5: Read a minibatch m
- 6: Calculate gradient: $g(\mathbf{W}_k^{t,m}) = \partial G(\mathbf{W}_k^{t,m} | \underline{\Omega}^{t,m}) / \partial \mathbf{W}_k^{t,m}$
- 7: Update the local k^{th} optimization variables:
 $\mathbf{W}_k^{(t+1,m)} \leftarrow \mathbf{W}_k^{(t,m)} - \eta g(\mathbf{W}_k^{(t,m)})$
 $\Omega_k^{(t+1,m)} \leftarrow (\Phi_{\text{task},\mathcal{M}_k}^T \Phi_{\text{task},\mathcal{M}_k})^{\frac{1}{2}} / \text{Tr}((\Phi_{\text{task},\mathcal{M}_k}^T \Phi_{\text{task},\mathcal{M}_k})^{\frac{1}{2}})$
- 8: **end for**
- 9: **end for**
- 10: **if** $t \bmod \tau = 0$ **then**
- 11: Perform aggregation and alignment over neighbors for node k :
 $\mathbf{W}_k^{(t+1)} \leftarrow (\sum_{j=1}^{|\mathcal{M}_k|} \frac{1}{N_j} \mathbf{W}_k^{(t)}) / |\mathcal{M}_k|$,
 $f_{align}(\Omega_k^{t+1}) \leftarrow \eta (\sum_{j=\mathcal{M}_k \setminus k} \frac{1}{N_j} f_{align}(\Omega_j^{(t)}) + f_{align}(\frac{(\Phi_{\text{task},\mathcal{M}_k}^T \Phi_{\text{task},\mathcal{M}_k})^{\frac{1}{2}}}{\text{Tr}((\Phi_{\text{task},\mathcal{M}_k}^T \Phi_{\text{task},\mathcal{M}_k})^{\frac{1}{2}})})) / |\mathcal{M}_k|$
- 12: **end if**
- 13: **end for**
- 14: **end for**

B. Related Works

Molecular Representation Learning. (Rogers and Hahn, 2010) encode the neighbors of atoms in the molecule into a fix-length vector to obtain vectors-space representations. To improve the expressive power of chemical fingerprints, (Duvenaud et al., 2015; Coley et al., 2017) use CNNs to learn rich molecule embeddings for downstream tasks like property prediction. (Kearnes et al., 2016; Schütt et al., 2017a;b) explore the graph convolutional network to encode molecular graphs into neural fingerprints. To better capture the interactions among atoms, (Gilmer et al., 2017) proposes to use a message passing framework and (Yang et al., 2019; Klicpera et al., 2020) extend this framework to model bond interactions.

FL. Early examples of research into federated learning include (Konečný et al., 2015; 2016; McMahan et al., 2016a). To address both statistical and system challenges, (Smith et al., 2017) and (Caldas et al., 2018) propose a multi-task learning framework for federated learning and its related optimization algorithm, which extends early works in distributed machine learning (Shalev-Shwartz and Zhang, 2013; Yang, 2013; Yang et al., 2013; Jaggi et al., 2014; Ma et al., 2015; Smith et al., 2016). The main limitation of (Smith et al., 2017) and (Caldas et al., 2018), however, is that strong duality is only guaranteed when the objective function is convex, which can not be generalized to the non-convex setting.(Ahmed et al., 2014; Jin et al., 2015; Mateos-Núñez et al., 2015; Wang et al., 2016; Baytas et al., 2016; Liu et al., 2017) extends federated multi-task learning to the distributed multi-task learnings setting, but not only limitation remains same , but also nodes performing same amount of work is prohibitive in FL.

Federated Graph Neural Networks. (Suzumura et al., 2019) and (Mei et al., 2019) use computed graph statistics for information exchange and aggregation to avoid node information leakage. (Jiang et al., 2020; Zhou et al., 2020) utilize cryptographic approaches to incorporate into GNN learning. (Wang et al., 2020) propose a hybrid of federated and meta learning to solve the semi-supervised graph node classification problem in decentralized social network datasets. (Meng et al., 2021) uses an edge-cloud partitioned GNN model for spatio-temporal traffic forecasting tasks over sensor networks. The previous works do not consider graph learning in a decentralized setting.

Stochastic Gradient Descent Optimization. In large scale distributed machine learning problems learning, synchronized mini-batch SGD is a well-known method to address the communication bottleneck by increasing the computation to communication ratio (Dean et al.; Li et al., 2014; Cui et al., 2014). It is shown that FedAvg (Konečný et al., 2016) is a special case of local SGD which allow nodes to perform local updates and infrequent synchronization between them to communicate less while converging fast (Stich, 2018; Wang and Joshi, 2018; Yu et al., 2018; Lin et al., 2018). Decentralized SGD, another approach to reducing communication, was successfully applied to deep learning [Jin et al., 2016; Jiang et al., 2017; Lian et al., 2017]. Asynchronous SGD is a potential method that can alleviate synchronization delays in distributed learning [Recht et al., 2011; Cui et al., 2014; Gupta et al., 2016; Mitliagkas et al., 2016; Dutta et al., 2018], but existing asynchronous SGD doesn’t fit for federated learning because the staleness problem is particularly severe due to the reason of heterogeneity in the federated setting (Dai et al., 2018).

C. Dataset Details

Table 2: Dataset summary used in our experiments.

Dataset	# Molecules	Avg # Nodes	Avg # Edges	# Tasks	Task Type	Evaluation Metric
SIDER	1427	33.64	35.36	27	Classification	ROC-AUC
Tox21	7831	18.51	25.94	12	Classification	ROC-AUC
MUV	93087	24.23	76.80	17	Classification	ROC-AUC
QM8	21786	7.77	23.95	12	Regression	MAE

Table 2 summarizes the necessary information of benchmark datasets (Wu et al., 2018). The details of each dataset are listed below:

- **SIDER** (Kuhn et al., 2016), or Side Effect Resource, the dataset consists of marketed drugs with their adverse drug reactions.
- **Tox21** (tox, 2017) is a dataset which records the toxicity of compounds.
- **MUV** (Rohrer and Baumann, 2009) is a subset of PubChem BioAssay processed via refined nearest neighbor analysis. Contains 17 tasks for around 90 thousand compounds and is specifically designed for validation of virtual screening techniques.
- **QM8** (Ramakrishnan et al., 2015) is composed from a recent study on modeling quantum mechanical calculations of electronic spectra and excited state energy of small molecules.

C.1. Feature Extraction Procedure for Molecules

The feature extraction is in two steps:

1. Atom-level feature extraction and Molecule object construction using RDKit (Landrum, 2006).
2. Constructing graphs from molecule objects using NetworkX (Hagberg et al., 2008).

Atom features, shown in Table 3, are the atom features we used. It’s exactly the same set of features as used in (Rong et al., 2020a).

Table 3: Atom features

Features	Size	Description
atom type	100	Representation of atom (e.g., C, N, O), by its atomic number
formal charge	5	An integer electronic charge assigned to atom
number of bonds	6	Number of bonds the atom is involved in
chirality	5	Number of bonded hydrogen atoms
number of H	5	Number of bonded hydrogen atoms
atomic mass	1	Mass of the atom, divided by 100
aromaticity	1	Whether this atom is part of an aromatic system
hybridization	5	SP, SP2, SP3, SP3D, or SP3D2

D. Model Hyperparameters

D.1. Model Architecture

As explained in section 2.1 our model is made up of a GNN and Readout. The GNNs we use are GAT (Veličković et al., 2018) and GraphSAGE (Hamilton et al., 2017). Each accepts input node features $X_v \in \mathcal{R}^{|V_M| \times d_{input}}$ and outputs node embeddings $h_v \in \mathcal{R}^{|V_M| \times d_{node}}$, $v \in V_M$. Where V_M is the set of atoms in molecule M . Given the output node embeddings from the GNN the Readout function we use is defined as follows:

$$R_{\Phi_{pool}, \Phi_{task}}(h_v, X_v) = \text{MEAN}(\text{ReLU}(\Phi_{task}(\text{ReLU}(\Phi_{pool}(X_v || h_v))))))$$

where $||$ represents the row wise concatenation operation. $\Phi_{pool} \in \mathcal{R}^{(d_{node} + d_{input}) \times d_{pool}}$ and $\Phi_{task} \in \mathcal{R}^{d_{pool} \times d_{out}}$ are learnable transformation matrices. d_{out} represents the number of classes/tasks present in the classification label. The MEAN operation here is a column wise mean. Note that while our general description of the readout in section 2.1 does not include the input features as part of the input, we find that including the input features leads to better generalization.

D.2. Hyperparameter Configurations

For each task, we utilize grid search to find the best results. Table 4 lists all the hyper-parameters ranges used in our experiments. All hyper-parameter tuning is run on a single GPU. The best hyperparameters for each dataset and model are listed in Table 5. The batch-size is kept 1. This pertains to processing a single molecule at a time. The number of GNN layers were fixed to 2 because having too many GNN layers result in over-smoothing phenomenon as shown in (Li et al., 2018). For all experiments, we used Adam optimizer.

Table 4: Hyperparameter Range for Experiments

hyper-parameter	Description	Range
Learning rate	Rate of speed at which the model learns.	[0.00015, 0.001, 0.0015, 0.0025, 0.015, 0.15]
Dropout rate	Dropout ratio	[0.3, 0.5, 0.6]
Node embedding dimension (d_{node})	Dimensionality of the node embedding	64
Hidden layer dimension	GNN hidden layer dimensionality	64
Readout embedding dimension (d_{pool})	Readout Hidden Layer Dimensionality	64
Graph embedding dimension (d_{out})	Dimensionality of the final graph embedding	64
Attention heads	Number of attention heads required for GAT	1-7
Alpha	LeakyRELU parameter used in GAT model	0.2
Rounds	Number of federating learning rounds	150
Epoch	Epoch of clients	1
Number of clients	Number of users in a federated learning round	4-10
Communication Period	Exchanging Period between clients	1

Table 5: Hyperparameters used in our experiments. For SpreadGNN we use a communication period $\tau = 1$ and a complete topology (all clients connected to all other clients) in all experiments.

	Parameters	GraphSAGE			GAT		
		FedAvg	FedGMTL	SpreadGNN	FedAvg	FedGMTL	SpreadGNN
SIDER	ROC-AUC Score	0.582	0.629	0.5873	0.5857	0.61	0.603
	Partition alpha	0.2	0.2	0.2	0.2	0.2	0.2
	Learning rate	0.0015	0.0015	0.0015	0.0015	0.0015	0.0015
	Dropout rate	0.3	0.3	0.3	0.3	0.3	0.3
	Node embedding dimension	64	64	64	64	64	64
	Hidden layer dimension	64	64	64	64	64	64
	Readout embedding dimension	64	64	64	64	64	64
	Graph embedding dimension	64	64	64	64	64	64
	Attention Heads	NA	NA	NA	2	2	2
	Leaky ReLU alpha	NA	NA	NA	0.2	0.2	0.2
	Number of Clients	4	4	4	4	4	4
Task Regularizer	NA	0.001	0.001	NA	0.001	0.001	
Tox21	ROC-AUC Score	0.5548	0.6644	0.585	0.6035	0.6594	0.6056
	Partition alpha	0.1	0.1	0.1	0.1	0.1	0.1
	Learning rate	0.0015	0.0015	0.0015	0.0015	0.0015	0.0015
	Dropout rate	0.3	0.3	0.3	0.3	0.3	0.3
	Node embedding dimension	64	64	64	64	64	64
	Hidden layer dimension	64	64	64	64	64	64
	Readout embedding dimension	64	64	64	64	64	64
	Graph embedding dimension	64	64	64	64	64	64
	Attention Heads	NA	NA	NA	2	2	2
	Leaky ReLU alpha	NA	NA	NA	0.2	0.2	0.2
	Number of Clients	8	8	8	8	8	8
Task Regularizer	NA	0.001	0.001	NA	0.001	0.001	
MUV	ROC-AUC Score	0.6578	0.6856	0.703	0.709	0.6899	0.713
	Partition alpha	0.3	0.3	0.3	0.3	0.3	0.3
	Learning rate	0.001	0.001	0.001	0.0025	0.0025	0.0025
	Dropout rate	0.3	0.3	0.3	0.3	0.3	0.3
	Node embedding dimension	64	64	64	64	64	64
	Hidden layer dimension	64	64	64	64	64	64
	Readout embedding dimension	64	64	64	64	64	64
	Graph embedding dimension	64	64	64	64	64	64
	Attention Heads	NA	NA	NA	2	2	2
	Leaky ReLU alpha	NA	NA	NA	0.2	0.2	0.2
	Number of Clients	8	8	8	8	8	8
Task Regularizer	NA	0.001	0.001	NA	0.002	0.002	
QM8	RMSE Score	0.02982	0.03624	0.02824	0.0392	0.0488	0.0333
	Partition alpha	0.5	0.5	0.5	0.5	0.5	0.5
	Learning rate	0.0015	0.0015	0.0015	0.0015	0.0015	0.0015
	Dropout rate	0.3	0.3	0.3	0.3	0.3	0.3
	Node embedding dimension	64	64	64	64	64	64
	Hidden layer dimension	64	64	64	64	64	64
	Readout embedding dimension	64	64	64	64	64	64
	Graph embedding dimension	64	64	64	64	64	64
	Attention Heads	NA	NA	NA	2	2	2
	Leaky ReLU alpha	NA	NA	NA	0.2	0.2	0.2
	Number of Clients	8	8	8	8	8	8
Task Regularizer	NA	0.3	0.3	NA	0.3	0.3	

E. Detailed Ablation Studies

E.1. Effect of Communication Period τ




Figure 6: Effect of Communication Period τ on GraphSAGE Model




Figure 7: Effect of Communication Period τ on GAT Model

E.2. Effect of Serverless Network Topology

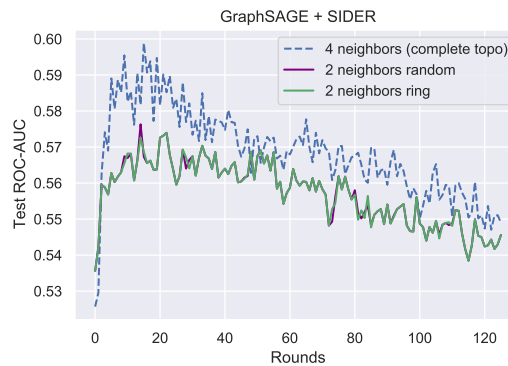


Figure 8: Effect of Topology on Graphsage + Sider using a 4 client network

Figure 8 illustrates the effect of varying topologies on SpreadGNN on the Sider dataset when using Graphsage as the GNN. The qualitative behavior is similar to Figure 5, in that when each client is connected to more neighbors, the local model on each client is more robust. However, when the total number of clients involved in the network is smaller, the effect of

topology is understated and the total number of neighbors matters more. Recall that in an 8 client network, when each client was restricted to being connected to only 2 neighbors, random connections performed worse than a ring topology, meaning that the topology mattered as much as mere number of neighbors. However, in the case of the 4 client network in Figure 8 there is a minimal difference between a 2 neighbor random configuration and a 2 neighbor ring configuration.

F. Proof for Convergence of DPA-SGD

In our analysis, we assume that following properties hold (Bottou et al., 2018):

- The objective function $F(\cdot)$ is **L-Lipschitz**.
- $F(\cdot)$ is **lower bounded** by F_{inf} such that $F(\cdot) \geq F_{inf}$.
- The full gradient of the objective function $F(\cdot)$ is approximated by stochastic gradient g on a mini-batch ϵ_i with **an unbiased estimate**: $\mathbb{E}_{\epsilon_k}[g(\mathbf{x})] = \nabla F(\mathbf{x})$.
- The variance of stochastic gradient $g(\mathbf{x})$ given a mini-batch ϵ_k is **upper bounded** and $\text{Var}_{\epsilon_k}(g(\mathbf{x})) \leq \beta \|\nabla F(\mathbf{x})\|^2 + \sigma^2$, $\exists \beta, \sigma^2 \geq 0$, $\forall k$.

In order to get a more clear view of our algorithm, we reformulate the loss function on each worker as follows:

$$f_k \left(\mathbf{\Gamma}, \Phi_{\text{task}\mathcal{M}_k}, \mathbf{\Omega}^{(k)}; \xi_i^{(k)} \right) := \mathcal{L} \left(\hat{\mathbf{y}}_i^{(k)}(\mathbf{X}_{i,k}, \mathbf{Z}_{i,k}; \mathbf{W}_k), \mathbf{y}_i^{(k)} \right) + \frac{1}{2} \lambda_1 \text{Tr}(\Phi_{\text{task}\mathcal{M}_k} \mathbf{\Omega}_k^{-1} \Phi_{\text{task}\mathcal{M}_k}) \\ + \frac{1}{2} \sum_{\chi \in \{\theta, \Psi, \Phi_{\text{pool}}, \Phi_{\text{task}}\}} \lambda_\chi \|\chi\|_F^2,$$

where we include all the parameters that is different on each worker to be $\delta^{(k)} := \{\Phi_{\text{task}\mathcal{M}_k}\}$, $\mathbf{\Gamma}$ denotes the shared parameters, and $\xi_i^{(k)}$ is the random variable that denotes the data samples $(\mathbf{X}_{i,k}, \mathbf{Z}_{i,k}, \mathbf{y}_i^{(k)})$.

Therefore, the original objective function can be cast into the following form:

$$F \left(\mathbf{\Gamma}, \delta^{(1:K)}, \mathbf{\Omega}^{(1:K)} \right) = \frac{1}{n} \sum_{i=1}^K \mathbb{E}_{\xi_i^{(k)}} f_k \left(\mathbf{\Gamma}, \delta^{(k)}, \mathbf{\Omega}^{(k)}; \xi_i^{(k)} \right).$$

Notice that the updating rule for $\mathbf{\Gamma}$, $\delta^{(1:K)}$, and $\mathbf{\Omega}^{(1:K)}$ are different in that:

$$\mathbf{\Gamma}_{t+1} = \mathbf{\Gamma}_t - \eta \frac{\partial F}{\partial \mathbf{\Gamma}_t} \mathbf{M}, \\ \delta_{t+1}^{(k)} = \delta_t^{(k)} - \eta \frac{\partial F}{\partial \delta_t^{(k)}}, \\ \mathbf{\Omega}_{t+1}^{(k)} = \text{argmin}_{\mathbf{\Omega}^{(k)}} F \left(\mathbf{\Gamma}, \delta_t^{(1:K)}, \mathbf{\Omega}_t^{(1:K)} \right).$$

From the update rule of $\mathbf{\Omega}_t^{(1:K)}$, we know that

$$\mathbb{E}_t F \left(\bar{\mathbf{\Gamma}}_{t+1}, \delta_{t+1}^{(1:K)}, \mathbf{\Omega}_{t+1}^{(1:K)} \right) - \mathbb{E}_t F \left(\bar{\mathbf{\Gamma}}_t, \delta_t^{(1:K)}, \mathbf{\Omega}_t^{(1:K)} \right) \\ \leq \mathbb{E}_t \left\langle \frac{\partial F}{\partial \bar{\mathbf{\Gamma}}_t}, \bar{\mathbf{\Gamma}}_{t+1} - \bar{\mathbf{\Gamma}}_t \right\rangle + \mathbb{E}_t \left\langle \frac{\partial F}{\partial \delta_t^{(1:K)}}, \delta_{t+1}^{(1:K)} - \delta_t^{(1:K)} \right\rangle + \frac{L}{2} \mathbb{E}_t \left(\|\bar{\mathbf{\Gamma}}_{t+1} - \bar{\mathbf{\Gamma}}_t\|_F^2 + \|\delta_{t+1}^{(1:K)} - \delta_t^{(1:K)}\|_F^2 \right) \\ \leq -\eta \left(1 - \frac{L\eta}{2} \right) \mathbb{E}_t \left\| \frac{\partial F}{\partial \delta_t^{(1:K)}} \right\|_F^2 - \frac{\eta}{2} \mathbb{E}_t \left\| \frac{\partial F}{\partial \bar{\mathbf{\Gamma}}_t} \right\|_F^2 - \frac{\eta(1-L\eta)}{2} \mathbb{E}_t \left\| \frac{\partial F}{\partial \mathbf{\Gamma}_t} \frac{1_n}{n} \right\|_F^2 + \frac{\eta}{2} \left\| \frac{\partial F}{\partial \bar{\mathbf{\Gamma}}_t} - \frac{\partial F}{\partial \mathbf{\Gamma}_t} \frac{1_n}{n} \right\|_F^2 \\ + \frac{L\eta^2(\sigma_\Gamma^2 + \sigma_\delta^2)}{2K}, \tag{12}$$

where the overline $\bar{*}$ denotes the expectation operation \mathbb{E} , σ_{Γ}^2 and σ_{δ}^2 are the variance bounds for the stochastic gradients of Γ_t and δ_t respectively. To estimate the upper bound for $\mathbb{E} \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} - \frac{\partial F}{\partial \Gamma_t} \frac{1_n}{n} \right\|_F^2$, we utilize convexity and Lipschitzness as follows:

$$\begin{aligned} \mathbb{E} \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} - \frac{\partial F}{\partial \Gamma_t} \frac{1_n}{n} \right\|_F^2 &= \frac{1}{n^2} \mathbb{E} \left\| \sum_{i=1}^n \left(\frac{\partial F_i}{\partial \bar{\Gamma}_t} - \frac{\partial F_i}{\partial \Gamma_t^{(i)}} \right) \right\|_F^2 \\ &\leq \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left\| \frac{\partial F_i}{\partial \bar{\Gamma}_t} - \frac{\partial F_i}{\partial \Gamma_t^{(i)}} \right\|_F^2 \\ &\leq \frac{L^2}{n} \mathbb{E} \sum_{i=1}^n \left\| \bar{\Gamma}_t - \Gamma_t^{(i)} \right\|_F^2. \end{aligned}$$

Therefore, we re-write the lower bound (12) as

$$-\eta \left(1 - \frac{L\eta}{2} \right) \mathbb{E}_t \left\| \frac{\partial F}{\partial \delta_t^{(1:K)}} \right\|_F^2 - \frac{\eta}{2} \mathbb{E}_t \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 - \frac{\eta(1-L\eta)}{2} \mathbb{E}_t \left\| \frac{\partial F}{\partial \Gamma_t} \frac{1_n}{n} \right\|_F^2 + \frac{\eta L^2}{2n} \mathbb{E}_t \sum_{i=1}^n \left\| \bar{\Gamma}_t - \Gamma_t^{(i)} \right\|_F^2 + \frac{L\eta^2(\sigma_{\Gamma}^2 + \sigma_{\delta}^2)}{2K}.$$

Summing the inequality above for all time-steps $t = 0, \dots, T$, we get

$$\begin{aligned} -\eta \left(1 - \frac{L\eta}{2} \right) \sum_{t=0}^T \mathbb{E}_t \left\| \frac{\partial F}{\partial \delta_t^{(1:K)}} \right\|_F^2 - \frac{\eta}{2} \sum_{t=0}^T \mathbb{E}_t \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 + \frac{\eta(1-L\eta)}{2} \sum_{t=0}^T \mathbb{E}_t \left\| \frac{\partial F}{\partial \Gamma_t} \frac{1_n}{n} \right\|_F^2 \\ + \frac{\eta L^2}{2n} \sum_{t=0}^T \mathbb{E}_t \sum_{i=1}^n \left\| \bar{\Gamma}_t - \Gamma_t^{(i)} \right\|_F^2 + \frac{L\eta(\sigma_{\Gamma}^2 + \sigma_{\delta}^2)T}{2K}. \end{aligned} \quad (13)$$

The main challenge however now becomes bounding the term with $\sum_{t=0}^T \mathbb{E}_t \sum_{i=1}^n \left\| \bar{\Gamma}_t - \Gamma_t^{(i)} \right\|_F^2$. Bounding it requires to derive another lower bound and using an available result. First re-write the $\mathbb{E} \left\| \frac{\partial f}{\partial \bar{\Gamma}_t} \right\|_F^2$ by utilizing Frobenius-norm properties:

$$\begin{aligned} \mathbb{E} \left\| \frac{\partial f}{\partial \bar{\Gamma}_t} \right\|_F^2 &= \mathbb{E} \left\| \left(\frac{\partial f}{\partial \Gamma_t} - \frac{\partial F}{\partial \Gamma_t} \right) + \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 \\ &= \mathbb{E} \left\| \frac{\partial f}{\partial \Gamma_t} - \frac{\partial F}{\partial \Gamma_t} \right\|_F^2 + \mathbb{E} \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 + 2\mathbb{E} \left\langle \frac{\partial f}{\partial \Gamma_t} - \frac{\partial F}{\partial \Gamma_t}, \frac{\partial F}{\partial \bar{\Gamma}_t} \right\rangle \\ &= \mathbb{E} \left\| \frac{\partial f}{\partial \Gamma_t} - \frac{\partial F}{\partial \Gamma_t} \right\|_F^2 + \mathbb{E} \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 \end{aligned}$$

Then bound each term as follows:

$$\begin{aligned} &\leq n\sigma_{\Gamma}^2 + \mathbb{E} \sum_{i=1}^n \left\| \left(\frac{\partial F_i}{\partial \Gamma_t^{(i)}} - \frac{\partial F_i}{\partial \bar{\Gamma}_t} \right) + \left(\frac{\partial F_i}{\partial \bar{\Gamma}_t} - \frac{\partial F}{\partial \bar{\Gamma}_t} \right) + \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 \\ &\leq n\sigma_{\Gamma}^2 + \sum_{i=1}^n \mathbb{E} \left\| \frac{\partial F_i}{\partial \Gamma_t^{(i)}} - \frac{\partial F_i}{\partial \bar{\Gamma}_t} \right\|_F^2 + \sum_{i=1}^n \mathbb{E} \left\| \frac{\partial F_i}{\partial \bar{\Gamma}_t} - \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 + \mathbb{E} \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2 \\ &\leq n\sigma_{\Gamma}^2 + L^2 \sum_{i=1}^n \mathbb{E} \left\| \bar{\Gamma}_t - \Gamma_t^{(i)} \right\|_F^2 + n\zeta^2 + n\mathbb{E} \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2, \end{aligned}$$

where first term is bounded by its stochastic gradient variance. To bound the second term $\mathbb{E} \left\| \frac{\partial F}{\partial \bar{\Gamma}_t} \right\|_F^2$ above, first bound the term with sum of the individual components' norm and then use add-subtract trick used with $\frac{\partial F}{\partial \bar{\Gamma}_t}$ and $\frac{\partial F_i}{\partial \bar{\Gamma}_t}$. Then, use the

facts $\sum_{i=1}^n \mathbb{E} \left\| \frac{\partial F_i}{\partial \bar{\mathbf{\Gamma}}_t} - \frac{\partial F_i}{\partial \mathbf{\Gamma}_t^{(i)}} \right\|_F^2 \leq L^2 \mathbb{E} \sum_{i=1}^n \|\bar{\mathbf{\Gamma}}_t - \mathbf{\Gamma}_t^{(i)}\|_F^2$ (an intermediate result derived above) and $\sum_{i=1}^n \mathbb{E} \left\| \frac{\partial F_i}{\partial \bar{\mathbf{\Gamma}}_t} - \frac{\partial F}{\partial \bar{\mathbf{\Gamma}}_t} \right\|_F^2 \leq n\zeta^2$, while ζ being the spectral gap of matrix \mathbf{M} . Then, from (Tang et al., 2018), we have

$$\sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left\| \bar{\mathbf{\Gamma}}_t - \mathbf{\Gamma}_t^{(i)} \right\|_F^2 \leq \frac{2}{(1-\zeta)^2} \sum_{t=1}^T \eta^2 \left\| \frac{\partial f}{\partial \bar{\mathbf{\Gamma}}_t} \right\|_F^2,$$

with the previously derived bound over $\sum_{t=1}^T \eta_F^2 \left\| \frac{\partial f}{\partial \bar{\mathbf{\Gamma}}_t} \right\|_F^2$ bound leading to another intermediate bound

$$\begin{aligned} \left(1 - \frac{2\eta^2 L^2}{(1-\zeta)^2}\right) \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left\| \bar{\mathbf{\Gamma}}_t - \mathbf{\Gamma}_t^{(i)} \right\|_F^2 &\leq \frac{2\eta^2}{(1-\zeta)^2} (n\sigma_{\bar{\Gamma}}^2 + n\zeta^2) + n \sum_{t=1}^T \mathbb{E} \left\| \frac{\partial F}{\partial \bar{\mathbf{\Gamma}}_t} \right\|_F^2 \\ \sum_{t=0}^T \sum_{i=1}^n \mathbb{E} \left\| \bar{\mathbf{\Gamma}}_t - \mathbf{\Gamma}_t^{(i)} \right\|_F^2 &\leq C_1 \eta^2 n (\sigma_{\bar{\Gamma}}^2 + \zeta^2) T + \sum_{t=1}^T \mathbb{E} \left\| \frac{\partial F}{\partial \bar{\mathbf{\Gamma}}_t} \right\|_F^2 \end{aligned} \quad (14)$$

where C_1 being a function of the spectral gap.

Finally, combining (13) with the intermediate bound (14) together first, bounding the Frobenius norms, and dividing both sides by T (to average in the end), we get the desired lower bound

$$\frac{2[F(\mathbf{x}_0) - F_{inf}]}{\eta T} + \frac{\eta L \sigma^2}{K} + \eta^2 L^2 \sigma^2 \left(\frac{1 + \zeta^2}{1 - \zeta^2} \tau - 1 \right)$$

where η is the learning rate that satisfies the given conditions, $\mathbf{x}_0 = \{\bar{\mathbf{\Gamma}}_0, \delta_0^{(1:K)}, \mathbf{\Omega}_0^{(1:K)}\}$ is the initial starting point and $\sigma^2 = \sigma_{\bar{\Gamma}}^2 + \sigma_{\delta}^2$ is the total variance bound over stochastic gradients of $\bar{\mathbf{\Gamma}}$ & δ \square

Our work differs from (Wang and Joshi, 2018) in that it does not provide adequate theoretical analysis and empirical evaluation for federated learning.