

Modeling Inter-Intra Heterogeneity for Graph Federated Learning

Wentao Yu¹, Shuo Chen^{2*}, Yongxin Tong³, Tianlong Gu⁴, Chen Gong^{5*}

¹School of Computer Science and Engineering, Nanjing University of Science and Technology, China

²Center for Advanced Intelligence Project, RIKEN, Japan

³School of Computer Science and Engineering, Beihang University, China

⁴Engineering Research Center of Trustworthy AI (Ministry of Education), Jinan University, China

⁵Department of Automation, Institute of Image Processing and Pattern Recognition, Shanghai Jiao Tong University, China
wentao.yu@njust.edu.cn, shuo.chen.ya@riken.jp, yxtong@buaa.edu.cn, gutianlong@jnu.edu.cn, goodgongchen@gmail.com

Abstract

Heterogeneity is a fundamental and challenging issue in federated learning, especially for the graph data due to the complex relationships among the graph nodes. To deal with the heterogeneity, lots of existing methods perform the weighted federation based on their calculated similarities between pairwise clients (*i.e.*, subgraphs). However, their inter-subgraph similarities estimated with the outputs of local models are less reliable, because the final outputs of local models may not comprehensively represent the real distribution of subgraph data. In addition, they ignore the critical intra-heterogeneity which usually exists within each subgraph itself. To address these issues, we propose a novel **F**ederated learning method by integrally modeling the **I**nter-**I**ntra **H**eterogeneity (FedIIH). For the inter-subgraph relationship, we propose a novel hierarchical variational model to infer the whole distribution of subgraph data in a multi-level form, so that we can accurately characterize the inter-subgraph similarities with the global perspective. For the intra-heterogeneity, we disentangle the subgraph into multiple latent factors and partition the model parameters into multiple parts, where each part corresponds to a single latent factor. Our FedIIH not only properly computes the distribution similarities between subgraphs, but also learns disentangled representations that are robust to irrelevant factors within subgraphs, so that it successfully considers the inter- and intra- heterogeneity simultaneously. Extensive experiments on six homophilic and five heterophilic graph datasets in both non-overlapping and overlapping settings demonstrate the effectiveness of our method when compared with nine state-of-the-art methods. Specifically, FedIIH averagely outperforms the second-best method by a large margin of 5.79% on all heterophilic datasets. Code is available at <https://github.com/blgpb/FedIIH>.

Introduction

Graphs are ubiquitous data structures in lots of important domains such as social media, transportation, and recommendation systems. In many real-world scenarios, a global graph is usually made up of multiple subgraphs that are distributed across devices, and subgraphs are only locally accessible due to privacy and regulatory concerns. Recently, Graph Federated Learning (GFL) has received increasing attention (He et al. 2022; Wang et al. 2023; Tan et al. 2023b;

Huang et al. 2023), where each client individually trains a local model based on the corresponding subgraph, and a central server adaptively aggregates the models from all clients.

Since subgraphs are different parts of a global graph, there inevitably exists heterogeneity among them, making it difficult to realize federated collaboration. Ignoring this heterogeneity will degrade the performance of several traditional Federated Learning (FL) methods (*e.g.*, FedAvg (McMahan et al. 2017)) which rely heavily on the assumption that all clients have similar data distributions. To deal with this deficiency, recently, a number of personalized FL methods (Li et al. 2020; Pillutla et al. 2022; Tan et al. 2023a; Baek et al. 2023; Zhang et al. 2023) have been proposed. For example, in (Baek et al. 2023; Li et al. 2023; Zhang et al. 2024), personalized GFL methods estimate the pairwise similarities between clients, and then they perform weighted federations based on the client similarities.

However, the performance of most existing personalized GFL methods is still limited due to their improper similarity calculations and ignored intra-heterogeneity. First, most existing methods (Baek et al. 2023; Li et al. 2023; Zhang et al. 2024) compute the inter-subgraph similarities based on the simplex outputs of local models. Since the outputs of local models cannot accurately reveal the whole distribution of subgraph data, the calculated similarities are hardly generalizable, leading to negative impacts on the weighted federation. Moreover, although most existing methods successfully consider inter-heterogeneity, they ignore the crucial intra-heterogeneity (*i.e.*, the heterogeneity of intra-subgraph), which commonly exists in real-world graphs. Intra-heterogeneity can be defined as different types of connections in the subgraph on each client. For example, a user in a social graph is connected to others for various different reasons, such as families, hobbies, studies, and work. Meanwhile, as confirmed by (Xie et al. 2021), inter-heterogeneity is defined as divergent distributions of both graph structures and node features among different clients.

To address the above-mentioned issues, we propose a novel **F**ederated learning method by integrally modeling the **I**nter-**I**ntra **H**eterogeneity (FedIIH), which is shown in the right panel of Fig. 1. On one hand, we propose a new Hierarchical Variational Graph AutoEncoder (HVGAE) by using the hierarchical Bayesian model (Gelman and Hill 2006; Tran, Ranganath, and Blei 2017), so as to build the poste-

*Corresponding authors: Chen Gong, Shuo Chen.




Figure 1: A framework comparison between existing methods and our FedIIIH.

rior dependencies between the latent factors of the local subgraphs and those of the global graph. Subsequently, we can successfully infer the subgraph data distribution and compute the similarities between pairwise clients (see Fig. 1c). On the other hand, we disentangle the local subgraph into multiple latent factors (see Fig. 1d). After that, the model parameters can be partitioned to correspond exactly to each latent factor, such that we can accordingly perform the separate federation based on our calculated similarities.

Our FedIIIH not only properly computes the distribution similarities between subgraphs, but also learns disentangled representations robust to irrelevant factors within each subgraph, effectively alleviating the inter-intra heterogeneity and significantly improving the model performances on different types of graphs. Main contributions of our work are:

- For the inter-heterogeneity, we propose a novel method HVGAE to characterize the inter-heterogeneity from a multi-level global perspective, so that we can infer the data distributions of local subgraphs and properly compute the distribution similarities between subgraphs.
- For the intra-heterogeneity, we disentangle the subgraph into several latent factors, such that the federations can be separately performed, and this is the first time in GFL that considers the intra-heterogeneity.
- Extensive experiments on eleven datasets demonstrate the effectiveness of our proposed FedIIIH, where our method averagely outperforms the second-best method by a large margin of 5.79% on all heterophilic graph data.

Related Work

Here we first describe the preliminaries, and then review the typical work related to this paper, including GFL and personalized FL.

Preliminaries

We focus on the task of node classification and aim to collaboratively train node classifiers with local subgraphs on distributed clients under the control of a server. For given M clients, each of them has a local subgraph $\mathcal{G}_m = (\mathcal{V}_m, \mathcal{E}_m)$, where \mathcal{V}_m represents the node set, and \mathcal{E}_m is the edge set ($m = 1, \dots, M$). The node feature matrix and adjacency matrix of \mathcal{G}_m is denoted as $\mathbf{X}_m \in \mathbb{R}^{n_m \times d}$ and $\mathbf{A}_m \in \mathbb{R}^{n_m \times n_m}$, respectively. Here n_m is the number of nodes in \mathcal{G}_m and d is the feature dimension. Due to the privacy constraints, \mathcal{G}_m on each client is inaccessible to the others.

Graph Federated Learning

Since each client owns only a part of the global graph, there inevitably exists heterogeneity. To deal with this issue, FED-PUB (Baek et al. 2023) estimates similarities between subgraphs based on the outputs of local models. Then, it performs a weighted averaging of the local models for each client based on the similarities. Similarly, FedGTA (Li et al. 2023) and FedGT (Zhang et al. 2024) compute the similarities based on the mixed moments of processed neighbor features and embedding vectors, respectively. Then, they both perform the weighted federation, where model parameters with high similarities are assigned with larger weights during the weighted federation. However, the outputs of local models (e.g., embedding vectors) cannot faithfully reveal the whole distribution of subgraph data, such that the computed similarities are improper, which may impact the weighted federation, thus decreasing the performances on the clients. Consequently, we infer the whole distribution of subgraph data in a multi-level global perspective, so as to properly compute the similarities.

Personalized Federated Learning

Heterogeneity is a fundamental and challenging problem in FL (Ye et al. 2023). To deal with the heterogeneity, personalized FL methods (Li et al. 2020; Tan et al. 2023a; Arivazhagan et al. 2019) have obtained increasing attention. Unlike FedAvg, which aims to train a global model collaboratively, personalized FL methods not only pursue commonalities among multiple clients, but also retain the personality of each client as much as possible. Personalized FL methods can be categorized as similarity-based methods (Baek et al. 2023; Li et al. 2023; Zhang et al. 2024), local customization-based methods (Li et al. 2020; Arivazhagan et al. 2019; T. Dinh, Tran, and Nguyen 2020), and meta-learning-based methods (Chen et al. 2018; Fallah, Mokhtari, and Ozdaglar 2020). For similarity-based methods, they compute the inter-subgraph similarities and then perform weighted federation. In contrast, as one of the local customization-based methods, FedProx (Li et al. 2020) customizes a personalized model for each client by adding a proximal term as a subproblem. Similarly, FedPer (Arivazhagan et al. 2019) only federates the weights of the backbone while training the personalized classification layer in each client. However, these methods only consider the inter-heterogeneity, while ignoring the intra-heterogeneity. Therefore, in this paper, we propose to

characterize both inter- and intra- heterogeneity.

Methodology

This section details our proposed FedIIIH. Specifically, we describe the modeling process of the intra-heterogeneity and inter-heterogeneity, respectively.

Modeling Intra-heterogeneity

Since there are diverse connecting relations among nodes in a real-world subgraph (Ma et al. 2019; Li et al. 2021; Guo et al. 2024), there is inevitably heterogeneity within the subgraph. To deal with the intra-heterogeneity, we aim to disentangle the subgraph into K latent factors, which are utilized to represent different relations within the subgraph. Note that most of the existing disentangled graph convolutional networks (Ma et al. 2019; Li et al. 2021; Guo et al. 2024) are applicable to learning disentangled representations in GFL. In this paper, the well-known Disentangled Graph Convolutional Network (DisenGCN) (Ma et al. 2019) is considered for its simplicity. Here we use the node u as an example to describe the disentangling process of DisenGCN. Given one node $i \in \{u\} \cup \{v | (u, v) \in \mathcal{G}_m\}$ in the subgraph \mathcal{G}_m , let $\mathbf{x}^i \in \mathbb{R}^d$ denote its node feature vector. This node feature \mathbf{x}^i is first projected to K subspaces, and its node representation in the k -th subspace can be represented by

$$\mathbf{z}^{i,k} = \frac{\sigma[(\mathbf{W}^k)^\top \mathbf{x}^i + \mathbf{b}^k]}{\|\sigma[(\mathbf{W}^k)^\top \mathbf{x}^i + \mathbf{b}^k]\|_2}, \quad (1)$$

where $\mathbf{W}^k \in \mathbb{R}^{d \times \frac{d_{\text{out}}}{K}}$ and $\mathbf{b}^k \in \mathbb{R}^{\frac{d_{\text{out}}}{K}}$ are learnable parameters. Here $\sigma[\cdot]$ denotes the nonlinear activation function, and d_{out} denotes the output dimension of node representations. After the projection operation via Eq. (1), $\mathbf{z}^{i,k} \in \mathbb{R}^{\frac{d_{\text{out}}}{K}}$ represents the aspect of node i that are related with the k -th latent factor. Second, DisenGCN proposes a neighborhood routing mechanism to identify the latent factor that causes the connection between node u and its neighbor node v , and accordingly extract features of v that are specific to that factor. Due to space limitations, details of the neighborhood routing mechanism are provided in the Appendix A. With this mechanism, the node representation in each subspace is aggregated independently, such that we can obtain the disentangled latent factors $\{\mathbf{c}^{u,1}, \mathbf{c}^{u,2}, \dots, \mathbf{c}^{u,K}\}$, where $\mathbf{c}^{u,k} \in \mathbb{R}^{\frac{d_{\text{out}}}{K}}$ represents the k -th aspect of node u . Note that there are no learnable parameters in the neighborhood routing mechanism. Finally, the disentangled node representation of u can be obtained by

$$\mathbf{h}^u = \text{Con}(\mathbf{c}^{u,1}, \mathbf{c}^{u,2}, \dots, \mathbf{c}^{u,K}), \quad (2)$$

where $\mathbf{h}^u \in \mathbb{R}^{d_{\text{out}}}$, and ‘Con’ denotes the concatenation operation performed along the column. Similarly, we can obtain all of the disentangled node representations in \mathcal{G}_m based on Eq. (2). After that, we use the matrix $\mathbf{H}_m \in \mathbb{R}^{n_m \times d_{\text{out}}}$ to denote the disentangled node representations in \mathcal{G}_m . According to Eq. (1), we can find that \mathbf{W}^k and \mathbf{b}^k correspond exactly to the k -th latent factor. Consequently, we propose to perform the separate federation with parameters (*i.e.*, \mathbf{W}^k and \mathbf{b}^k) specific to the k -th latent factor, which is described in the following section.




Figure 2: The graphical model of our proposed HVGAE, where $\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K$ denote the local latent factors on the m -th client, and $\alpha^1, \alpha^2, \dots, \alpha^K$ denote the global latent factors on the server.

Modeling Inter-heterogeneity

Since the local views of subgraphs may not be sufficient to estimate the inter-heterogeneity, we aim to model the inter-heterogeneity from a multi-level global perspective. Moreover, since the similarities between clients are essentially determined by the similarities of local subgraph data distributions, we seek to infer the data distribution of subgraphs and thereby compute the similarities between clients through divergences. Specifically, we propose the HVGAE to infer the subgraph data distribution of the m -th client based on the disentangled node representations (*i.e.*, \mathbf{H}_m). For ease of expression, we rewrite \mathbf{H}_m as $\mathbf{H}_m^1, \mathbf{H}_m^2, \dots, \mathbf{H}_m^K$, where $\mathbf{H}_m^k \in \mathbb{R}^{n_m \times \frac{d_{\text{out}}}{K}}$ denotes the disentangled node representations with the k -th latent factor. As $\mathbf{H}_m^1, \mathbf{H}_m^2, \dots, \mathbf{H}_m^K$ are deterministic results output by DisenGCN, we use $\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K$ to represent random latent variables.

Hierarchical model To build the posterior dependencies between the latent factors of the local subgraphs and those of the global graph, we assume that the latent factors of the subgraphs are governed by those of the global graph based on the theory of hierarchical Bayesian model (Gelman and Hill 2006; Tran, Ranganath, and Blei 2017). As shown in Fig. 2, $\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K$ on the m -th client are local latent factors, which are linked via the higher-level global latent factors $\alpha^1, \alpha^2, \dots, \alpha^K$ on the server, respectively. Note that $\alpha^1, \alpha^2, \dots, \alpha^K$ are shared across all clients, where $\alpha^k \in \mathbb{R}^{\frac{d_{\text{out}}}{K}}$.

Based on the graphical model in Fig. 2, the joint probability can be formulated as

$$p(\mathcal{G}_{1:M}, \tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \alpha^{1:K}) = \prod_{k=1}^K p(\alpha^k) \cdot \prod_{m=1}^M \prod_{k'=1}^K p(\mathcal{G}_m | \tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K) p(\tilde{\mathbf{H}}_m^k | \alpha^{k'}), \quad (3)$$

where $\mathcal{G}_{1:M}$ is the abbreviation for $(\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_M)$. $(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K)$ is the abbreviation for $(\tilde{\mathbf{H}}_1^1, \tilde{\mathbf{H}}_2^1, \dots, \tilde{\mathbf{H}}_M^1, \tilde{\mathbf{H}}_1^2, \tilde{\mathbf{H}}_2^2, \dots, \tilde{\mathbf{H}}_M^2, \tilde{\mathbf{H}}_1^K, \tilde{\mathbf{H}}_2^K, \dots, \tilde{\mathbf{H}}_M^K)$, $\alpha^{1:K}$ is the abbreviation for $(\alpha^1, \alpha^2, \dots, \alpha^K)$, and $p(\alpha^k)$ denotes the prior distribution of α^k . In Eq. (3), $p(\mathcal{G}_m | \tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K)$ denotes the conditional distribution, and $p(\tilde{\mathbf{H}}_m^k | \alpha^{k'})$ denotes the prior distribution of $\tilde{\mathbf{H}}_m^k$. Moreover, the true posterior distribution of both local and global latent factors can be formulated as $p(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \alpha^{1:K} | \mathcal{G}_{1:M})$. However,

this true posterior distribution is computationally intractable. Therefore, we attempt to approximate it with a tractable approximate posterior distribution $q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})$, which can be formulated as

$$\begin{aligned} & q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M}) \\ &= \prod_{k=1}^K q(\boldsymbol{\alpha}^k) \cdot \prod_{m=1}^M \prod_{k'=1}^K q(\tilde{\mathbf{H}}_m^{k'} | \mathcal{G}_m), \end{aligned} \quad (4)$$

where $q(\boldsymbol{\alpha}^k)$ denotes the marginal distribution, and $q(\tilde{\mathbf{H}}_m^{k'} | \mathcal{G}_m)$ represents the approximate posterior distribution of $\tilde{\mathbf{H}}_m^{k'}$. For ease of derivation, we can flexibly instantiate the above-mentioned distributions. The detailed instantiations are presented in the Appendix B.

Variational inference Since the approximate posterior distributions in Eq. (4) are still intractable, we propose to use the variational inference (Kingma and Welling 2013; Zhang et al. 2019) to infer them. According to the graphical model in Fig. 2, the Evidence Lower Bound (ELBO) can be derived as follows:

$$\begin{aligned} \mathcal{L}_{\text{ELBO}} &= \sum_{k=1}^K \log p(\tilde{\boldsymbol{\alpha}}^k) + \sum_{m=1}^M \left\{ \mathbb{E}_{q(\tilde{\mathbf{H}}_m | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m)] \right. \\ &\quad \left. - \sum_{k'=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^{k'} | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^{k'} | \tilde{\boldsymbol{\alpha}}^{k'})) \right\}, \end{aligned} \quad (5)$$

where $\log p(\tilde{\boldsymbol{\alpha}}^k)$ denotes the log prior probability of $\tilde{\boldsymbol{\alpha}}^k$. Here $\tilde{\boldsymbol{\alpha}}^k$ denotes the posterior mean of $\boldsymbol{\alpha}^k$ for the k -th global latent factor, and D_{KL} denotes the Kullback-Leibler (KL) divergence (Moreno, Ho, and Vasconcelos 2003). The detailed derivation of ELBO and $\tilde{\boldsymbol{\alpha}}^k$ are provided in the Appendix C and D, respectively. Eq. (5) shows that the derived ELBO is the sum of two terms, where the first term is related to the global latent factors on the server, and the second term can be parallelly optimized on M clients.

Adaptation to FL However, since the server does not have access to the private data of the clients, the optimization process on the server (*i.e.*, $\sum_{k=1}^K \log p(\tilde{\boldsymbol{\alpha}}^k)$) is challenging. To adapt Eq. (5) to the federated learning scenario, we make some modifications to $\sum_{k=1}^K \log p(\tilde{\boldsymbol{\alpha}}^k)$. On one hand, we substitute $\sum_{k=1}^K \log p(\tilde{\boldsymbol{\alpha}}^k)$ by $\sum_{m=1}^M \sum_{k=1}^K \log p(\hat{\boldsymbol{\alpha}}_m^k)$, where $\hat{\boldsymbol{\alpha}}_m^k$ denotes the learnable parameter deployed on the m -th client. On the other hand, $\tilde{\boldsymbol{\alpha}}^k$ on each client is constrained by the KL divergence between $p(\hat{\boldsymbol{\alpha}}_m^k)$ and $p(\tilde{\boldsymbol{\alpha}}^k)$. Finally, we have that

$$\begin{aligned} \hat{\mathcal{L}}_{\text{ELBO}} &= \sum_{m=1}^M \sum_{k=1}^K \left\{ \log p(\hat{\boldsymbol{\alpha}}_m^k) - D_{\text{KL}}(p(\hat{\boldsymbol{\alpha}}_m^k) || p(\tilde{\boldsymbol{\alpha}}^k)) \right\} + \\ &\quad \sum_{m'=1}^M \left\{ \mathbb{E}_{q(\tilde{\mathbf{H}}_{m'} | \mathcal{G}_{m'})} [\log p(\mathcal{G}_{m'} | \tilde{\mathbf{H}}_{m'})] - \right. \\ &\quad \left. \sum_{k'=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_{m'}^{k'} | \mathcal{G}_{m'}) || p(\tilde{\mathbf{H}}_{m'}^{k'} | \tilde{\boldsymbol{\alpha}}^{k'})) \right\}. \end{aligned} \quad (6)$$

Eq. (6) implies that the optimizations of ELBO can be distributedly performed on M clients in a parallel way.

Variational graph autoencoder To implement the approximate posterior distribution $q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m)$ and the mathematical expectation $\mathbb{E}_{q(\tilde{\mathbf{H}}_m | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m)]$ in Eq. (6), we introduce a simple yet effective framework (*i.e.*, Variational Graph AutoEncoder (VGAE) (Kipf and Welling 2016)) on each client, which includes an inference network (*a.k.a.* encoder) and a generative network (*a.k.a.* decoder). First, we take an inference model parameterized by two DisenGCNs (*i.e.*, $\text{DisenGCN}_{\boldsymbol{\mu}_m}(\mathcal{G}_m)$ and $\text{DisenGCN}_{\boldsymbol{\sigma}_m}(\mathcal{G}_m)$). Note that $\text{DisenGCN}_{\boldsymbol{\mu}_m}(\mathcal{G}_m)$ and $\text{DisenGCN}_{\boldsymbol{\sigma}_m}(\mathcal{G}_m)$ are used to infer the means and standard deviations of \mathcal{G}_m for K latent factors (*i.e.*, $\mathbf{H}_{m,\boldsymbol{\mu}}$ and $\mathbf{H}_{m,\boldsymbol{\sigma}}$), respectively. For ease of expression, we rewrite them as $\mathbf{H}_{m,\boldsymbol{\mu}}^1, \mathbf{H}_{m,\boldsymbol{\mu}}^2, \dots, \mathbf{H}_{m,\boldsymbol{\mu}}^K$ and $\mathbf{H}_{m,\boldsymbol{\sigma}}^1, \mathbf{H}_{m,\boldsymbol{\sigma}}^2, \dots, \mathbf{H}_{m,\boldsymbol{\sigma}}^K$, respectively. By using the reparameterization trick (Kingma and Welling 2013), we can have

$$\tilde{\mathbf{H}}_m^k = \mathbf{H}_{m,\boldsymbol{\mu}}^k + \mathbf{H}_{m,\boldsymbol{\sigma}}^k \odot \boldsymbol{\epsilon}, \quad (7)$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, and \odot denotes the element-wise product. Second, our generative network is constructed by an inner product between latent variables. Since $\mathbb{E}_{q(\tilde{\mathbf{H}}_m | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m)]$ in Eq. (6) can be regarded as a reconstruction loss, we implement it by

$$\begin{aligned} \mathbb{E}_{q(\tilde{\mathbf{H}}_m | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m)] &= p(\mathbf{A}_m | \tilde{\mathbf{H}}_m) \\ &= \prod_{i=1}^{n_m} \prod_{j=1}^{n_m} p(\mathbf{A}_m^{ij} | \mathbf{r}_i, \mathbf{r}_j), \end{aligned} \quad (8)$$

where $p(\mathbf{A}_m^{ij} = 1 | \mathbf{r}_i, \mathbf{r}_j) = \sigma(\mathbf{r}_i^\top \mathbf{r}_j)$, and \mathbf{A}_m^{ij} denotes the element of \mathbf{A}_m . In Eq. (8), \mathbf{r}_i and \mathbf{r}_j are the i -th and j -th rows of the matrix $\tilde{\mathbf{H}}_m$, respectively. The detailed implementation of HVGAE can be found in the Appendix E.

Similarity calculation Since the similarities between clients are essentially determined by the similarities of the local data distributions, we seek to calculate the similarities between clients based on the divergences of the inferred subgraph data distributions. Moreover, since HVGAE has K disentangled latent factors, we can compute the similarities corresponding to each latent factor separately. Specifically, the similarity with respect to the k -th latent factor between clients m and j can be computed as

$$S(m, j)^k = 1 - D_{\text{JS}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || q(\tilde{\mathbf{H}}_j^k | \mathcal{G}_j)), \quad (9)$$

where $S(m, j)^k$ denotes the similarity, and D_{JS} denotes the Jensen-Shannon (JS) divergence (Sutter, Daunhawer, and Vogt 2020). Note that Eq. (9) uses only approximate posterior distributions sent to the server, which does not compromise data privacy.

Federated Aggregation

Before describing the federation process, let us analyze the learnable parameters that should be federated in our FedIIIH. The first part of the federated parameters are

$\mathbf{W}_m^1, \mathbf{W}_m^2, \dots, \mathbf{W}_m^K$ and $\mathbf{b}_m^1, \mathbf{b}_m^2, \dots, \mathbf{b}_m^K$ in Eq. (1). The second part of the federated parameters are $\mathbf{W}_m^{\text{cls}} \in \mathbb{R}^{c \times d_{\text{out}}}$ and $\mathbf{b}_m^{\text{cls}} \in \mathbb{R}^c$, which come from the node classifier. Here c is the number of node classes. The detailed analysis can be found in the Appendix G. For the federated aggregation of $\mathbf{W}_m^1, \mathbf{W}_m^2, \dots, \mathbf{W}_m^K, \mathbf{b}_m^1, \mathbf{b}_m^2, \dots, \mathbf{b}_m^K$, and $\mathbf{W}_m^{\text{cls}}$, we propose a separate federation method. Specifically, with the computed similarity $S(i, j)^k$, we use the weighted averaging of parameters across different clients. Moreover, since HVGAE has K disentangled latent factors, and parameters (e.g., \mathbf{W}^k and \mathbf{b}^k in Eq. (1)) correspond exactly to the k -th latent factor, we can therefore perform the separate federation according to each latent factor. For ease of expression, we rewrite $\mathbf{W}_m^{\text{cls}}$ as $\mathbf{W}_m^{\text{cls},1}, \mathbf{W}_m^{\text{cls},2}, \dots, \mathbf{W}_m^{\text{cls},K}$, where $\mathbf{W}_m^{\text{cls},k} \in \mathbb{R}^{c \times \frac{d_{\text{out}}}{K}}$ denotes the parameters with respect to the k -th latent factor. Our proposed separate federation can be defined as

$$\begin{aligned} \overline{\mathbf{W}}_m^k &\leftarrow \sum_{j=1}^M \beta_{mj}^k \cdot \mathbf{W}_j^k, & \overline{\mathbf{b}}_m^k &\leftarrow \sum_{j=1}^M \beta_{mj}^k \cdot \mathbf{b}_j^k, \\ \overline{\mathbf{W}}_m^{\text{cls},k} &\leftarrow \sum_{j=1}^M \beta_{mj}^k \cdot \mathbf{W}_j^{\text{cls},k}, \end{aligned} \quad (10)$$

where

$$\beta_{mj}^k = \frac{\exp(\tau \cdot S(m, j)^k)}{\sum_{l=1}^M \exp(\tau \cdot S(m, l)^k)}, \quad (11)$$

and β_{mj}^k denotes the normalized similarity with respect to the k -th latent factor between clients m and j . Here ‘ \leftarrow ’ denotes the assignment operation, and τ denotes a hyperparameter for scaling the similarity score. Our proposed separate federation not only allows different clients to obtain personalized parameters, which is beneficial for dealing with inter-heterogeneity, but also allows parameters in the same client to be federated separately according to each latent factor, which is beneficial for dealing with intra-heterogeneity. For the federation of the remaining parameter $\mathbf{b}_m^{\text{cls}} \in \mathbb{R}^c$, we use the federation method as proposed in FedAvg (McMahan et al. 2017), which is simple but effective. The pseudocode of our FedIIIH is listed in the Appendix H.

Experiments

To validate the effectiveness of our FedIIIH, we perform extensive experiments on eleven widely used benchmark datasets. These datasets include both homophilic and heterophilic graphs, the statistical information and descriptions of which are included in the Appendix I.2. We use both the non-overlapping and overlapping subgraph partitioning settings. To ensure a fair comparison, we compute the mean accuracy (or mean AUC), and the corresponding standard deviation over three independent runs according to (Baek et al. 2023). See Appendix I for more details.

Main Results

Homophilic datasets Tab. 1 and Tab. 2 show the node classification results on the homophilic datasets in two partitioning settings, respectively. We observe that our FedIIIH achieves the best performance among all the methods, and

the standard deviations are also relatively small as well, suggesting that FedIIIH is more effective and stable than the compared methods. Moreover, the average accuracy of FedIIIH for all six datasets in the non-overlapping scenario is 83.10%, which is 1.51% higher than the second-best method (i.e., FED-PUB). In the overlapping scenario, the average accuracy of FedIIIH is 81.01%, which is 1.48% higher than the second-best method (i.e., FED-PUB). Generally, the non-overlapping scenario is more challenging than the overlapping scenario (Baek et al. 2023; Zhang et al. 2024) due to the increasing heterogeneity. However, our FedIIIH can still outperform the second-best method in both the non-overlapping and overlapping scenarios, and this validates the effectiveness of modeling inter-intra heterogeneity. For example, as shown in Tab. 1, the classification accuracy of FedIIIH increases from 93.42% to 93.55% when the number of clients increases from 5 to 20 on the *Amazon-Photo* dataset. This is non-trivial because the data heterogeneity increases from 0.664 to 0.759 (Zhang et al. 2024) during this change, indicating that FedIIIH is robust to the heterogeneity.

Heterophilic datasets Tab. 3 and Tab. 4 show the node classification results on the heterophilic datasets in two partitioning settings, respectively. We can find that our proposed FedIIIH not only achieves the best average performance among all baseline methods, but also outperforms the second-best method (i.e., FedSage+) by 5.79% and 4.53% in the non-overlapping and overlapping scenarios, respectively. This is because our FedIIIH not only considers the inter-heterogeneity as these methods do, but also successfully deals with the intra-heterogeneity that they ignore. Moreover, although the intra-heterogeneity on the heterophilic datasets is stronger than that on the homophilic datasets (Platonov et al. 2023), FedIIIH achieves a larger performance improvement than on the homophilic datasets, which further validates the effectiveness of our FedIIIH.

Ablation Study

Our FedIIIH employs the hierarchical model and variational inference to compute the inter-subgraph similarities. In addition, we use the disentanglement to learn disentangled representations. To shed light on the contributions of these components, we report the experimental results of FedIIIH when each of these components is removed on the *Cora* and *Amazon-ratings* datasets in Tab. 5. For simplicity, ‘w/o HM’, ‘w/o VI’, and ‘w/o Dis’ denote the reduced models by removing the hierarchical model, variational inference, and disentanglement, respectively. We can clearly find that the performance decreases when any component is removed, showing that each component contributes a lot to the final performance. For example, the accuracies on the *Cora* dataset are obviously reduced when the disentanglement component is removed. The ablation studies on other datasets are shown in the Appendix J.1.

Effectiveness of Inter-Subgraph Similarity and K

First, we validate the effectiveness of the inter-subgraph similarities calculated by our method. In Fig. 3, we show the similarity heatmaps on the *Amazon-ratings* dataset in the overlapping setting with 20 clients. The similarity heatmaps

Table 1: Node classification results of different methods on the **homophilic** graph datasets in the **non-overlapping** subgraph partitioning setting. The best results are shown in **bold**.

Methods	Cora			CiteSeer			PubMed			-
	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	-
Local	81.30±0.21	79.94±0.24	80.30±0.25	69.02±0.05	67.82±0.13	65.98±0.17	84.04±0.18	82.81±0.39	82.65±0.03	-
FedAvg (McMahan et al. 2017)	74.45±5.64	69.19±0.67	69.50±3.58	71.06±0.60	63.61±3.59	64.68±1.83	79.40±0.11	82.71±0.29	80.97±0.26	-
FedProx (Li et al. 2020)	72.03±4.56	60.18±7.04	48.22±6.18	71.73±1.11	63.33±3.25	64.85±1.35	79.45±0.25	82.55±0.24	80.50±0.25	-
FedPer (Arivazhagan et al. 2019)	81.68±0.40	79.35±0.04	78.01±0.32	70.41±0.32	70.53±0.28	66.64±0.27	85.80±0.21	84.20±0.28	84.72±0.31	-
GCFL (Xie et al. 2021)	81.47±0.65	78.66±0.27	79.21±0.70	70.34±0.57	69.01±0.12	66.33±0.05	85.14±0.33	84.18±0.19	83.94±0.36	-
FedGNN (Wu et al. 2021)	81.51±0.68	70.12±0.99	70.10±3.52	69.06±0.92	55.52±3.17	52.23±6.00	79.52±0.23	83.25±0.45	81.61±0.59	-
FedSage+(Zhang et al. 2021)	72.97±5.94	69.05±1.59	57.97±12.6	70.74±0.69	65.63±3.10	65.46±0.74	79.57±0.24	82.62±0.31	80.82±0.25	-
FED-PUB (Baek et al. 2023)	83.70±0.19	81.54±0.12	81.75±0.56	72.68±0.44	72.35±0.53	67.62±0.12	86.79±0.09	86.28±0.18	85.53±0.30	-
FedGTA (Li et al. 2023)	80.06±0.63	80.59±0.38	79.01±0.31	70.12±0.10	71.57±0.34	69.94±0.14	87.75±0.01	86.80±0.01	87.12±0.05	-
AdaFGL (Li et al. 2024)	82.01±0.51	80.09±0.00	79.74±0.05	71.44±0.27	72.34±0.00	70.95±0.45	86.91±0.28	86.97±0.10	86.59±0.21	-
FedIIIH (Ours)	84.11±0.17	81.85±0.09	83.01±0.15	72.86±0.25	76.50±0.06	73.36±0.41	87.80±0.18	87.65±0.18	87.19±0.25	-
Amazon-Computer										
Methods	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	Avg.
Local	89.22±0.13	88.91±0.17	89.52±0.20	91.67±0.09	91.80±0.02	90.47±0.15	66.76±0.07	64.92±0.09	65.06±0.05	79.57
FedAvg (McMahan et al. 2017)	84.88±1.96	79.54±0.23	74.79±0.24	89.89±0.83	83.15±3.71	81.35±1.04	65.54±0.07	64.44±0.10	63.24±0.13	74.58
FedProx (Li et al. 2020)	85.25±1.27	83.81±1.09	73.05±1.30	90.38±0.48	80.92±4.64	82.32±0.29	65.21±0.20	64.37±0.18	63.04±0.04	72.84
FedPer (Arivazhagan et al. 2019)	89.67±0.34	89.73±0.04	87.86±0.43	91.44±0.37	91.76±0.23	90.59±0.06	66.87±0.05	64.99±0.18	64.66±0.11	79.94
GCFL (Xie et al. 2021)	89.07±0.91	90.03±0.16	89.08±0.25	91.99±0.29	92.06±0.25	90.79±0.17	66.80±0.12	65.09±0.08	65.08±0.04	79.90
FedGNN (Wu et al. 2021)	88.08±0.15	88.18±0.41	83.16±0.13	90.25±0.70	87.12±2.01	81.00±4.48	65.47±0.22	64.21±0.32	63.80±0.05	75.23
FedSage+(Zhang et al. 2021)	85.04±0.61	80.50±1.13	70.42±0.85	90.77±0.44	76.81±8.24	80.58±1.15	65.69±0.09	64.52±0.14	63.31±0.20	73.47
FED-PUB (Baek et al. 2023)	90.74±0.05	90.55±0.13	90.12±0.09	93.29±0.19	92.73±0.18	91.92±0.12	67.77±0.09	66.58±0.08	66.64±0.12	81.59
FedGTA (Li et al. 2023)	86.69±0.18	86.66±0.23	85.01±0.87	93.33±0.12	93.50±0.21	92.61±0.15	60.32±0.04	60.22±0.09	58.74±0.14	79.45
AdaFGL (Li et al. 2024)	80.20±0.05	83.62±0.26	84.53±0.23	86.69±0.19	89.85±0.83	88.11±0.05	62.73±0.19	51.77±0.36	50.94±0.08	76.97
FedIIIH (Ours)	90.74±0.13	90.86±0.23	90.44±0.05	93.42±0.02	94.22±0.08	93.55±0.09	70.30±0.06	69.34±0.02	68.65±0.04	83.10

Table 2: Node classification results of different methods on the **homophilic** graph datasets in the **overlapping** subgraph partitioning setting. The best results are shown in **bold**.

Methods	Cora			CiteSeer			PubMed			-
	10 Clients	30 Clients	50 Clients	10 Clients	30 Clients	50 Clients	10 Clients	30 Clients	50 Clients	-
Local	73.98±0.25	71.65±0.12	76.63±0.10	65.12±0.08	64.54±0.42	66.68±0.44	82.32±0.07	80.72±0.16	80.54±0.11	-
FedAvg (McMahan et al. 2017)	76.48±0.36	53.99±0.98	53.99±4.53	69.48±0.15	66.15±0.64	66.51±1.00	82.67±0.11	82.05±0.12	80.24±0.35	-
FedProx (Li et al. 2020)	77.85±0.50	51.38±1.74	56.27±9.04	69.39±0.35	66.11±0.75	66.53±0.43	82.63±0.17	82.13±0.13	80.50±0.46	-
FedPer (Arivazhagan et al. 2019)	78.73±0.31	74.18±0.24	74.42±0.37	69.81±0.28	65.19±0.81	67.64±0.44	85.31±0.06	84.35±0.38	83.94±0.10	-
GCFL (Xie et al. 2021)	78.84±0.26	73.41±0.27	76.63±0.16	69.48±0.39	64.92±0.18	65.98±0.30	83.59±0.25	80.77±0.12	81.36±0.11	-
FedGNN (Wu et al. 2021)	70.63±0.83	61.38±2.33	56.91±0.82	68.72±0.39	59.98±1.52	58.98±0.98	84.25±0.07	82.02±0.22	81.85±0.10	-
FedSage+(Zhang et al. 2021)	77.52±0.46	51.99±0.42	55.48±1.15	68.75±0.48	65.97±0.02	65.93±0.30	82.77±0.08	82.14±0.11	80.31±0.68	-
FED-PUB (Baek et al. 2023)	79.60±0.12	75.40±0.54	77.84±0.23	70.58±0.20	68.33±0.45	69.21±0.30	85.70±0.08	85.16±0.10	84.84±0.12	-
FedGTA (Li et al. 2023)	76.42±0.62	75.63±0.33	77.69±0.14	70.43±0.08	71.71±0.33	69.19±0.32	85.34±0.42	84.99±0.05	84.47±0.06	-
AdaFGL (Li et al. 2024)	78.50±0.19	75.80±0.23	74.41±0.00	72.63±0.15	68.18±0.31	62.90±0.75	85.58±0.23	85.85±0.41	84.45±0.07	-
FedIIIH (Ours)	80.57±0.23	76.82±0.24	78.58±0.25	73.16±0.18	72.27±0.21	69.56±0.11	85.87±0.03	86.65±0.11	85.65±0.12	-
Amazon-Computer										
Methods	10 Clients	30 Clients	50 Clients	10 Clients	30 Clients	50 Clients	10 Clients	30 Clients	50 Clients	All
Local	88.50±0.20	86.66±0.00	87.04±0.02	92.17±0.12	90.16±0.12	90.42±0.15	62.52±0.07	61.32±0.04	60.04±0.04	76.72
FedAvg (McMahan et al. 2017)	88.99±0.19	83.37±0.47	76.34±0.12	92.91±0.07	89.30±0.22	74.19±0.57	63.56±0.02	59.72±0.06	60.94±0.24	73.38
FedProx (Li et al. 2020)	88.84±0.20	83.84±0.89	76.60±0.47	92.67±0.19	89.17±0.40	72.36±2.06	63.52±0.11	59.86±0.16	61.12±0.04	73.38
FedPer (Arivazhagan et al. 2019)	89.30±0.04	87.99±0.23	88.22±0.27	92.88±0.24	91.23±0.16	90.92±0.38	63.97±0.08	62.29±0.04	61.24±0.11	78.42
GCFL (Xie et al. 2021)	89.01±0.22	87.24±0.09	87.02±0.22	92.45±0.10	90.58±0.11	90.54±0.08	63.24±0.02	61.66±0.10	60.32±0.01	77.61
FedGNN (Wu et al. 2021)	88.15±0.09	87.00±0.10	83.96±0.88	91.47±0.11	87.91±1.34	78.90±6.46	63.08±0.19	60.09±0.04	60.51±0.11	73.66
FedSage+(Zhang et al. 2021)	89.24±0.15	81.33±1.20	76.72±0.39	92.76±0.05	88.69±0.99	72.41±1.36	63.24±0.02	59.90±0.12	60.95±0.09	73.12
FED-PUB (Baek et al. 2023)	89.98±0.08	89.15±0.06	88.76±0.14	93.22±0.07	92.01±0.07	91.71±0.11	64.18±0.04	63.34±0.12	62.55±0.12	79.53
FedGTA (Li et al. 2023)	90.10±0.18	88.79±0.27	88.15±0.21	93.13±0.14	92.49±0.06	91.77±0.06	55.98±0.09	56.76±0.07	57.89±0.09	74.40
AdaFGL (Li et al. 2024)	80.49±0.00	80.42±0.00	82.12±0.00	89.24±0.00	87.68±0.00	86.81±0.00	56.81±0.06	55.17±0.00	54.82±0.00	75.74
FedIIIH (Ours)	90.15±0.04	89.56±0.19	89.99±0.00	93.38±0.00	94.17±0.04	93.25±0.16	66.69±0.09	66.10±0.03	65.67±0.06	81.01

Table 3: Node classification results of different methods on the **heterophilic** graph datasets in the **non-overlapping** subgraph partitioning setting. Accuracy (%) is reported for *Roman-empire* and *Amazon-ratings*, and AUC (%) is reported for *Minesweeper*, *Tolokers*, and *Questions*. The best and second-best results are shown in **bold** and underlined, respectively.

Methods	Roman-empire			Amazon-ratings			Minesweeper			-
	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	-
Local	33.65±0.13	28.42±0.26	23.89±0.32	45.03±0.31	45.89±0.19	46.02±0.02	71.35±0.17	69.96±0.16	69.31±0.09	-
FedAvg (McMahan et al. 2017)	38.93±0.32	35.43±0.32	32.00±0.39	41.26±0.53	41.66±0.14	42.20±0.21	72.60±0.08	71.84±0.02	71.36±0.16	-
FedProx (Li et al. 2020)	27.95±0.59	26.43±1.41	23.12±0.49	36.92±0.00	36.86±0.14	36.96±0.00	71.91±0.27	70.66±0.20	71.50±0.37	-
FedPer (Arivazhagan et al. 2019)	20.75±1.75	15.51±1.13	15.45±2.76	36.62±0.30	32.34±1.01	36.96±0.00	58.73±10.45	65.35±7.02	53.80±11.40	-
GCFL (Xie et al. 2021)	30.40±0.16	29.44±0.49	26.73±0.19	36.92±0.00	36.86±0.14	36.96±0.00	72.04±0.13	71.14±0.09	47.77±0.14	-
FedGNN (Wu et al. 2021)	30.26±0.11	29.09±0.01	26.60±0.02	36.92±0.00	36.72±0.00	36.96±0.00	72.03±0.13	71.12±0.09	71.71±0.21	-
FedSage+(Zhang et al. 2021)	57.26±0.00	49.07±0.00	38.36±0.00	36.82±0.00	36.71±0.00	37.03±0.00	77.74±0.00	72.80±0.00	69.70±0.00	-
FED-PUB (Baek et al. 2023)	40.80±0.26	36.77±0.30	32.67±0.39	44.41±0.41	44.85±0.17	45.39±0.50	72.18±0.02	71.56±0.05	70.72±0.40	-
FedGTA (Li et al. 2023)	61.56±0.27	60.94±0.19	59.65±0.28	41.22±0.66	39.40±0.44	39.24±0.12	45.60±1.41	64.97±0.35	49.63±8.64	-
AdaFGL (Li et al. 2024)	67.64±0.18	64.55±0.00	62.42±0.26	41.70±0.06	42.30±0.00	42.59±0.14	47.45±2.10	65.59±0.56	51.48±7.14	-
FedIIIH (Ours)	68.32±0.05	66.44±0.28	64.61±0.13	44.26±0.24	44.24±0.10	45.19±0.04	74.29±0.02	73.23±0.04	72.81±0.02	-
Amazon-Computer										
Methods	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	5 Clients	10 Clients	20 Clients	All
Local	67.81±0.17	70.04±0.23	62.34±0.67	66.73±0.57	57.96±0.10	60.00±0.21	56.91	54.45	52.31	54.56
FedAvg (McMahan et al. 2017)	60.74±0.31	54.73±0.50	56.36±0.39	65.68±0.23	58.91±0.22	60.33±0.15	55.84	52.51	52.45	53.60
FedProx (Li et al. 2020)	42.90±0.24	41.15±0.22	40.42±0.62	47.36±0.38	45.46±0.34	46.83±0.11	45.41	44.11	43.77	44.43
FedPer (Arivazhagan et al. 2019)	46.61±9.88	54.97±13.23	44.82±11.61	58.38±9.39	59.40±9.71	62.32±1.56	44.22	45.51	42.67	44.13
GCFL (Xie et al. 2021)	27.61±2.55	19.81±0.57	17.53±0.04	47.94±0.41	45.71±0.25	47.47±0.21	42.98	40.59	35.29	39.62
FedGNN (Wu et al. 2021)	43.10±0.27	41.57±0.07	40.70±0.74	47.55±0.02	45.73±0.26	47.46±0.25	45.97	44.85	44.69	45.17
FedSage+(Zhang et al. 2021)	75.06±0.00	71.31±0.00	69.73±0.00	64.95±0.00	65.06±0.00	59.3				

Table 4: Node classification results of different methods on the **heterophilic** graph datasets in the **overlapping** subgraph partitioning setting. Accuracy (%) is reported for *Roman-empire* and *Amazon-ratings*, and AUC (%) is reported for *Minesweeper*, *Tolokers*, and *Questions*. The best and second-best results are shown in **bold** and underlined, respectively.

Methods	Roman-empire			Amazon-ratings			Minesweeper			-
	10 Clients	30 Clients	50 Clients	10 Clients	30 Clients	50 Clients	10 Clients	30 Clients	50 Clients	-
Local	39.47±0.03	34.43±0.14	31.28±0.18	41.43±0.04	41.81±0.14	42.57±0.12	67.98±0.07	64.39±0.10	62.73±0.23	-
FedAvg (McMahan et al. 2017)	40.89±0.25	38.66±0.08	36.71±0.20	39.86±0.06	41.40±0.02	41.02±0.16	69.06±0.07	67.95±0.04	66.89±0.08	-
FedProx (Li et al. 2020)	36.63±0.14	35.31±0.17	33.61±0.59	37.00±0.00	36.60±0.00	36.89±0.00	68.27±0.05	66.75±0.19	66.03±0.16	-
FedPer (Arivazhagan et al. 2019)	23.66±3.27	23.27±3.09	22.23±3.58	32.33±4.23	31.58±0.54	34.48±2.25	61.85±1.02	60.13±1.38	60.06±3.61	-
GCFL (Xie et al. 2021)	37.65±0.27	36.32±0.19	34.80±0.09	37.00±0.00	36.60±0.00	36.89±0.00	68.47±0.06	67.13±0.10	57.41±12.56	-
FedGNN (Wu et al. 2021)	37.46±0.12	36.30±0.16	34.84±0.13	37.00±0.00	36.60±0.00	36.89±0.00	68.47±0.06	67.12±0.11	66.41±0.23	-
FedSage+(Zhang et al. 2021)	57.48±0.00	42.55±0.00	33.99±0.00	36.86±0.00	36.71±0.00	37.03±0.00	76.64±0.00	70.56±0.00	70.34±0.00	-
FED-PUB (Baek et al. 2023)	43.80±0.25	40.46±0.16	37.73±0.09	42.25±0.25	42.25±0.06	42.88±0.34	68.80±0.09	67.43±0.25	65.98±0.15	-
FedGTA (Li et al. 2023)	59.86±0.04	58.32±0.09	57.57±0.21	40.81±0.24	39.44±0.06	39.37±0.04	54.35±0.73	48.20±1.28	52.94±1.77	-
AdaFGL (Li et al. 2024)	64.44±0.03	61.77±0.02	59.55±0.01	39.39±0.05	41.19±0.15	40.71±0.25	55.15±0.84	50.15±1.63	54.18±2.15	-
FedIIIH (Ours)	65.48±0.12	63.32±0.06	62.42±0.10	42.63±0.02	42.40±0.05	42.65±0.21	69.35±0.25	68.09±0.26	67.37±0.14	-
Methods	Tolokers			Questions			Avg.			All
Local	73.83±0.03	69.01±0.31	66.63±0.20	63.17±0.02	57.17±0.08	56.13±0.02	57.18	53.36	51.87	54.14
FedAvg (McMahan et al. 2017)	72.99±0.40	58.51±0.27	55.47±0.42	62.80±0.63	58.88±0.18	60.78±0.27	57.12	53.08	52.17	54.12
FedProx (Li et al. 2020)	54.49±1.69	45.59±0.41	41.49±0.45	52.53±0.34	51.54±0.41	50.72±0.40	49.78	47.16	45.75	47.56
FedPer (Arivazhagan et al. 2019)	39.60±0.11	59.44±0.79	41.92±0.06	61.31±0.29	53.41±1.53	50.29±0.10	43.75	45.57	41.80	43.70
GCFL (Xie et al. 2021)	55.91±1.13	23.26±0.70	18.40±0.25	53.04±0.47	51.84±0.38	51.10±0.38	50.41	43.03	39.72	44.39
FedGNN (Wu et al. 2021)	56.21±1.20	46.85±0.31	42.18±0.45	53.04±0.47	51.86±0.36	51.11±0.38	50.44	47.75	46.29	48.16
FedSage+(Zhang et al. 2021)	74.54±0.00	70.88±0.00	69.61±0.00	64.22±0.00	65.34±0.00	62.76±0.00	61.95	57.21	54.75	57.97
FED-PUB (Baek et al. 2023)	74.17±0.29	70.35±0.54	66.80±0.85	65.39±2.44	58.38±1.19	58.76±0.16	58.88	55.77	54.43	56.36
FedGTA (Li et al. 2023)	40.02±1.70	47.34±0.75	45.81±1.96	35.56±5.46	50.43±1.05	53.33±0.40	46.12	48.75	49.80	48.22
AdaFGL (Li et al. 2024)	45.15±2.15	49.18±0.84	47.54±2.48	41.05±6.49	52.18±2.16	56.46±0.92	49.04	50.89	51.69	50.54
FedIIIH (Ours)	71.67±0.02	71.69±0.12	69.99±0.03	68.79±0.09	66.98±0.04	64.73±0.35	63.58	62.50	61.43	62.50

Table 5: Ablation studies in both non-overlapping and overlapping partitioning settings on two datasets with 10 clients.

Methods	Cora		Amazon-ratings	
	non-overlapping	overlapping	non-overlapping	overlapping
w/o HM	78.67±1.17 (↓3.18)	78.58±0.03 (↓1.99)	41.69±0.09 (↓2.55)	38.04±0.09 (↓4.59)
w/o VI	73.51±0.52 (↓8.34)	78.40±0.25 (↓2.17)	41.68±0.05 (↓2.56)	40.70±0.42 (↓1.93)
w/o Dis	78.78±0.63 (↓3.07)	77.18±0.23 (↓3.39)	41.20±0.14 (↓3.04)	39.98±0.06 (↓2.65)
FedIIIH	81.85±0.09	80.57±0.23	44.24±0.10	42.63±0.02




Figure 3: Similarity heatmaps on the *Amazon-ratings* dataset in the overlapping setting with 20 clients.

on other datasets are shown in the Appendix K.2. Baek *et al.* (Baek et al. 2023) regard two clients have similar data distributions if they have similar label distributions. We argue that the label similarity cannot fully reflect the similarity of the subgraph data distribution, because both the node features and the structure are crucial properties of subgraphs. Therefore, we compute the subgraph distribution similarity by using the JS divergence, taking into account both semantic and structural information, and we treat such distribution similarity (Fig. 3a) as the ground truth. Both Fig. 3d and Fig. 3e are close to the ground truth, verifying the effectiveness of our similarity calculation scheme based on the inferred subgraph data distributions. In contrast, the similarity heatmaps of FED-PUB (Fig. 3b) and FedGTA (Fig. 3c) are both different from the ground truth, which implies that their calculated similarities are not good enough. Moreover, as shown in the Appendix K.1, our calculated similarities are much more stable than the similarities calculated by FED-PUB. Second, we investigate the effectiveness of K . As shown in Fig. 4, the performances on the *Tolokers* dataset increase consistently as the value of K increases. More results




Figure 4: Since the intra-heterogeneity of the *Tolokers* dataset is high, the performances increase consistently as the value of K increases.

are provided in the Appendix J.2. We can find that for most datasets the performances are usually bad when $K = 1$, and it means that disentangling the subgraph into several latent factors is beneficial to improve the performances. Last but not least, the efficiency and robustness analysis in the Appendix L and Appendix M clearly shows that our FedIIIH is more efficient and robust than the baseline methods.

Conclusion

In this paper, we proposed a novel method FedIIIH, which naturally integrates the inter- and intra- heterogeneity in GFL. On one hand, our new method characterizes the inter-heterogeneity from a multi-level global perspective, and thus it can properly compute the inter-subgraph similarities based on the whole distribution. On the other hand, it disentangles the subgraph into several latent factors, so that we can further consider the critical intra-heterogeneity. To the best of our knowledge, this is the first time in GFL that combines both inter- and intra- heterogeneity into a unified framework. Due to the adequate consideration of inter-intra heterogeneity, our method achieves satisfactory results on eleven datasets and outperforms the second-best method by a large margin of 5.79% on the heterophilic graph data.

Acknowledgments

This research is supported by NSF of China (Nos: 62336003, 12371510), and NSF for Distinguished Young Scholar of Jiangsu Province (No: BK20220080).

References

- Arivazhagan, M. G.; Aggarwal, V.; Singh, A. K.; and Choudhary, S. 2019. Federated learning with personalization layers. *arXiv:1912.00818*.
- Baek, J.; Jeong, W.; Jin, J.; Yoon, J.; and Hwang, S. J. 2023. Personalized subgraph federated learning. In *International Conference on Machine Learning*, 1396–1415.
- Chen, F.; Luo, M.; Dong, Z.; Li, Z.; and He, X. 2018. Federated meta-learning with fast convergence and efficient communication. *arXiv:1802.07876*.
- Fallah, A.; Mokhtari, A.; and Ozdaglar, A. 2020. Personalized federated learning with theoretical guarantees: A model-agnostic meta-learning approach. In *Advances in Neural Information Processing Systems*, 3557–3568.
- Gelman, A.; and Hill, J. 2006. *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.
- Guo, J.; Huang, K.; Yi, X.; and Zhang, R. 2024. Learning disentangled graph convolutional networks locally and globally. *IEEE Transactions on Neural Networks and Learning Systems*, 35(3): 3640–3651.
- He, C.; Ceyani, E.; Balasubramanian, K.; Annavam, M.; and Avestimehr, S. 2022. SpreadGNN: Decentralized multi-task federated learning for graph neural networks on molecular data. In *AAAI Conference on Artificial Intelligence*, 1–9.
- Huang, W.; Wan, G.; Ye, M.; and Du, B. 2023. Federated graph semantic and structural learning. In *International Joint Conference on Artificial Intelligence*, 1–9.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv:1312.6114*.
- Kipf, T. N.; and Welling, M. 2016. Variational graph auto-encoders. *arXiv:1611.07308*.
- Li, H.; Wang, X.; Zhang, Z.; Yuan, Z.; Li, H.; and Zhu, W. 2021. Disentangled contrastive learning on graphs. In *Advances in Neural Information Processing Systems*, 21872–21884.
- Li, T.; Sahu, A. K.; Zaheer, M.; Sanjabi, M.; Talwalkar, A.; and Smith, V. 2020. Federated optimization in heterogeneous networks. In *Proceedings of Machine Learning and Systems*, 429–450.
- Li, X.; Wu, Z.; Zhang, W.; Sun, H.; Li, R.-H.; and Wang, G. 2024. AdaFGL: A new paradigm for federated node classification with topology heterogeneity. In *International Conference on Data Engineering*, 2517–2530.
- Li, X.; Wu, Z.; Zhang, W.; Zhu, Y.; Li, R.-H.; and Wang, G. 2023. FedGTA: Topology-aware averaging for federated graph learning. In *Proceedings of the VLDB Endowment*, 41–50.
- Ma, J.; Cui, P.; Kuang, K.; Wang, X.; and Zhu, W. 2019. Disentangled graph convolutional networks. In *International Conference on Machine Learning*, 4212–4221.
- McMahan, B.; Moore, E.; Ramage, D.; Hampson, S.; and y Arcas, B. A. 2017. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics*, 1273–1282.
- Moreno, P.; Ho, P.; and Vasconcelos, N. 2003. A Kullback-Leibler divergence based kernel for SVM classification in multimedia applications. In *Advances in Neural Information Processing Systems*, 1–8.
- Pillutla, K.; Malik, K.; Mohamed, A.-R.; Rabbat, M.; Sanjabi, M.; and Xiao, L. 2022. Federated learning with partial model personalization. In *International Conference on Machine Learning*, 17716–17758.
- Platonov, O.; Kuznedelev, D.; Diskin, M.; Babenko, A.; and Prokhorenkova, L. 2023. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In *International Conference on Learning Representations*, 1–15.
- Sutter, T.; Daunhawer, I.; and Vogt, J. 2020. Multimodal generative learning utilizing Jensen-Shannon-Divergence. In *Advances in Neural Information Processing Systems*, 6100–6110.
- T. Dinh, C.; Tran, N.; and Nguyen, J. 2020. Personalized federated learning with moreau envelopes. In *Advances in Neural Information Processing Systems*, 21394–21405.
- Tan, A. Z.; Yu, H.; Cui, L.; and Yang, Q. 2023a. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12): 9587–9603.
- Tan, Y.; Liu, Y.; Long, G.; Jiang, J.; Lu, Q.; and Zhang, C. 2023b. Federated learning on non-iid graphs via structural knowledge sharing. In *AAAI Conference on Artificial Intelligence*, 1–9.
- Tran, D.; Ranganath, R.; and Blei, D. 2017. Hierarchical implicit models and likelihood-free variational inference. In *Advances in Neural Information Processing Systems*, 1–11.
- Wang, Z.; Fan, X.; Qi, J.; Jin, H.; Yang, P.; Shen, S.; and Wang, C. 2023. FedGS: Federated graph-based sampling with arbitrary client availability. In *AAAI Conference on Artificial Intelligence*, 1–8.
- Wu, C.; Wu, F.; Cao, Y.; Huang, Y.; and Xie, X. 2021. FedGNN: Federated graph neural network for privacy-preserving recommendation. *arXiv:2102.04925*.
- Xie, H.; Ma, J.; Xiong, L.; and Yang, C. 2021. Federated graph classification over non-iid graphs. In *Advances in Neural Information Processing Systems*, 18839–18852.
- Ye, M.; Fang, X.; Du, B.; Yuen, P. C.; and Tao, D. 2023. Heterogeneous federated learning: State-of-the-art and research challenges. *ACM Computing Surveys*, 56(3): 1–44.
- Zhang, C.; Bütepage, J.; Kjellström, H.; and Mandt, S. 2019. Advances in variational inference. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(8): 2008–2026.
- Zhang, J.; Hua, Y.; Wang, H.; Song, T.; Xue, Z.; Ma, R.; and Guan, H. 2023. FedALA: Adaptive local aggregation for personalized federated learning. In *AAAI Conference on Artificial Intelligence*, 1–8.

Zhang, K.; Yang, C.; Li, X.; Sun, L.; and Yiu, S. M. 2021. Subgraph federated learning with missing neighbor generation. In *Advances in Neural Information Processing Systems*, 6671–6682.

Zhang, Z.; Hu, Q.; Yu, Y.; Gao, W.; and Liu, Q. 2024. FedGT: Federated node classification with scalable graph transformer. *arXiv:2401.15203*.

Modeling Inter-Intra Heterogeneity for Graph Federated Learning (Appendix)

Wentao Yu¹ Shuo Chen^{2*} Yongxin Tong³ Tianlong Gu⁴ Chen Gong^{5*}

¹School of Computer Science and Engineering,
Nanjing University of Science and Technology, China

²Center for Advanced Intelligence Project, RIKEN, Japan

³School of Computer Science and Engineering, Beihang University, China

⁴Engineering Research Center of Trustworthy AI (Ministry of Education),
Jinan University, China

⁵Department of Automation, Institute of Image Processing and Pattern Recognition,
Shanghai Jiao Tong University, China

wentao.yu@njust.edu.cn, shuo.chen.ya@riken.jp, yxtong@buaa.edu.cn
gutianlong@jnu.edu.cn, goodgongchen@gmail.com

A Details of Neighborhood Routing Mechanism

Recall that we use the node u as an example to describe the disentangling process of DisenGCN. By using Eq. (1) in the main manuscript, the obtained $\mathbf{z}^{i,k}$ describes the aspect of node i that are related with the k -th latent factor. However, $\mathbf{z}^{i,k}$ only denotes the information of the node i itself. Therefore, for the node u , we have to mine information from the neighborhoods, which are connected with node u due to the k -th latent factor. Specifically, we aim to identify the latent factor that causes the connection between node u and its neighbor node v , and accordingly extract features of v that are specific to that factor. This is exactly the purpose of the neighborhood routing mechanism.

First, we define $p^{v,k}$, which represents the probability that the latent factor k is the reason why the node u reaches its neighbor v . Then, we can have $p^{v,k} \geq 0$, and $\sum_{k'=1}^K p^{v,k'} = 1$. Moreover, $p^{v,k}$ is also the probability that we should use the neighbor v to construct $\mathbf{c}^{u,k}$. The neighborhood routing mechanism will infer $p^{v,k}$ and construct $\mathbf{c}^{u,k}$ in an iterative manner. To initialize this process, we make $(p^{v,k})^{(1)} \propto \exp[(\mathbf{z}^{v,k})^\top \mathbf{z}^{u,k} / \tau_p]$. Then, it subsequently iterates to identify the largest cluster within each subspace, ensuring that each neighbor is primarily associated with a single subspace cluster:

$$(\mathbf{c}^{u,k})^{(t)} = \frac{\mathbf{z}^{u,k} + \sum_{v:(u,v) \in \mathcal{G}_m} (p^{v,k})^{(t-1)} \mathbf{z}^{v,k}}{\|\mathbf{z}^{u,k} + \sum_{v:(u,v) \in \mathcal{G}_m} (p^{v,k})^{(t-1)} \mathbf{z}^{v,k}\|_2}, \quad (1)$$

$$(p^{v,k})^{(t)} = \frac{\exp[(\mathbf{z}^{v,k})^\top (\mathbf{c}^{u,k})^{(t)} / \tau_p]}{\sum_{k'=1}^K \exp[(\mathbf{z}^{v,k'})^\top (\mathbf{c}^{u,k'})^{(t)} / \tau_p]}, \quad (2)$$

where $(\mathbf{c}^{u,k})^{(t)}$ denotes the center of each subspace cluster at the t -th iteration, $t = 2, 3, \dots, T$, and τ_p is a hyperparameter that controls the hardness of the assignment. Here τ_p is set to 1 according to [1]. After T times of iterations, the final output is $\mathbf{c}^{u,k} = (\mathbf{c}^{u,k})^{(T)}$.

B Instantiations of Distributions

We summarize the family of distributions instantiated by our proposed FedIIIH in Tab. 1. Specifically, the priors $p(\boldsymbol{\alpha}^k)$ and $p(\tilde{\mathbf{H}}_m^k)$ are both centered isotropic multivariate Gaussian distributions. Besides,

*Corresponding authors: Chen Gong, Shuo Chen.

Table 1: Family of distributions instantiated by our proposed FedIH.

$p(\tilde{\boldsymbol{\alpha}}^k)$	$\mathcal{N}(0, \sigma_{\tilde{\boldsymbol{\alpha}}^k}^2 \mathbf{I})$
$p(\tilde{\mathbf{H}}_m^k)$	$\mathcal{N}(0, \sigma_{\tilde{\mathbf{H}}_m^k}^2 \mathbf{I})$
$p(\tilde{\mathbf{H}}_m^k \boldsymbol{\alpha}^k)$	$\mathcal{N}(\boldsymbol{\alpha}^k, \sigma_{\tilde{\mathbf{H}}_m^k}^2 \mathbf{I})$
$q(\tilde{\boldsymbol{\alpha}}^k)$	$\mathcal{N}(\tilde{\boldsymbol{\alpha}}^k, \sigma_{\tilde{\boldsymbol{\alpha}}^k}^2 \mathbf{I})$
$q(\tilde{\mathbf{H}}_m^k \mathcal{G}_m)$	$\mathcal{N}(\hat{\boldsymbol{\mu}}_{\tilde{\mathbf{H}}_m^k}, \hat{\sigma}_{\tilde{\mathbf{H}}_m^k}^2)$

the prior over the local latent factor $\tilde{\mathbf{H}}_m^k$ conditioned on $\boldsymbol{\alpha}^k$ (i.e., $p(\tilde{\mathbf{H}}_m^k | \boldsymbol{\alpha}^k)$) is an isotropic multi-variate Gaussian distribution centered at $\boldsymbol{\alpha}^k$. Similarly, the marginal distribution of $\boldsymbol{\alpha}^k$ (i.e., $q(\boldsymbol{\alpha}^k)$) is a multivariate diagonal Gaussian distribution. As shown in Tab. 1, $\tilde{\boldsymbol{\alpha}}^k$ and $\sigma_{\tilde{\boldsymbol{\alpha}}^k}^2$ denote the posterior mean and variance of $\boldsymbol{\alpha}^k$, respectively. Moreover, $\hat{\boldsymbol{\mu}}_{\tilde{\mathbf{H}}_m^k}$ and $\hat{\sigma}_{\tilde{\mathbf{H}}_m^k}^2$ denote the variational mean and variance evaluated at \mathcal{G}_m , respectively.

C Derivation of the ELBO

First, the ELBO for the marginal likelihood of $\mathcal{G}_{1:M}$ (i.e., $\log p(\mathcal{G}_{1:M})$) can be obtained by using the Jensen inequality:

$$\begin{aligned}
& \log p(\mathcal{G}_{1:M}) \\
&= \log \sum_{\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M}} q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M}) \frac{p(\mathcal{G}_{1:M}, \tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K})}{q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})}, \\
&\geq \sum_{\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M}} q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M}) \log \frac{p(\mathcal{G}_{1:M}, \tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K})}{q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})}, \\
&\triangleq \text{ELBO}(q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M}), \mathcal{G}_{1:M}), \\
&= \mathbb{E}_{q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})} [\log p(\mathcal{G}_{1:M}, \tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K}) \\
&\quad - \log q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})].
\end{aligned} \tag{3}$$

Second, inspired by [2], the ELBO can be derived as follows:

$$\begin{aligned}
& \mathbb{E}_{q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})} [\log p(\mathcal{G}_{1:M}, \tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K}) \\
&\quad - \log q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})] \\
&= \sum_{m=1}^M \mathbb{E}_{q(\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K)] \\
&\quad - \sum_{m=1}^M \sum_{k=1}^K \mathbb{E}_{q(\boldsymbol{\alpha}^k)} [D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^k | \boldsymbol{\alpha}^k))] \\
&\quad - \sum_{k=1}^K D_{\text{KL}}(q(\boldsymbol{\alpha}^k) || p(\boldsymbol{\alpha}^k)).
\end{aligned} \tag{4}$$

Third, we compute the expected KL divergence of two Gaussian distributions (*i.e.*, $q(\tilde{\mathbf{H}}_m^k|\mathcal{G}_m)$ and $p(\tilde{\mathbf{H}}_m^k|\alpha^k)$) over a Gaussian distribution (*i.e.*, $q(\alpha^k)$) analytically. According to Tab. 1, we can have

$$\begin{aligned}
& \mathbb{E}_{q(\alpha^k)} [D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k|\mathcal{G}_m)||p(\tilde{\mathbf{H}}_m^k|\alpha^k))] \\
&= \mathbb{E}_{q(\alpha^k)} [D_{\text{KL}}(\mathcal{N}(\hat{\boldsymbol{\mu}}_{\tilde{\mathbf{H}}_m^k}, \hat{\boldsymbol{\sigma}}_{\tilde{\mathbf{H}}_m^k}^2)||\mathcal{N}(\alpha^k, \sigma_{\tilde{\mathbf{H}}_m^k}^2 \mathbf{I}))] \\
&= \mathbb{E}_{q(\alpha^k)} \left[-\frac{1}{2} \sum_{j=1}^J \left(1 + \log \frac{\hat{\sigma}_{\tilde{\mathbf{H}}_m^k, j}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} - \frac{(\hat{\mu}_{\tilde{\mathbf{H}}_m^k, j} - \alpha^{k, j})^2 + \hat{\sigma}_{\tilde{\mathbf{H}}_m^k, j}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} \right) \right] \\
&= -\frac{1}{2} \sum_{j=1}^J \left(1 + \log \frac{\hat{\sigma}_{\tilde{\mathbf{H}}_m^k, j}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} - \frac{\hat{\sigma}_{\tilde{\mathbf{H}}_m^k, j}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} - \mathbb{E}_{q(\alpha^k)} \left[\frac{(\hat{\mu}_{\tilde{\mathbf{H}}_m^k, j} - \alpha^{k, j})^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} \right] \right) \quad (5) \\
&= D_{\text{KL}}(\mathcal{N}(\hat{\boldsymbol{\mu}}_{\tilde{\mathbf{H}}_m^k}, \hat{\boldsymbol{\sigma}}_{\tilde{\mathbf{H}}_m^k}^2)||\mathcal{N}(\tilde{\boldsymbol{\alpha}}^k, \sigma_{\tilde{\mathbf{H}}_m^k}^2)) + \frac{J}{2} \frac{\sigma_{\tilde{\boldsymbol{\alpha}}^k}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} \\
&= D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k|\mathcal{G}_m)||p(\tilde{\mathbf{H}}_m^k|\tilde{\boldsymbol{\alpha}}^k)) + \frac{J}{2} \frac{\sigma_{\tilde{\boldsymbol{\alpha}}^k}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2},
\end{aligned}$$

where J denotes the dimension of $\tilde{\mathbf{H}}_m^k$. Moreover, $\alpha^{k, j}$, $\hat{\mu}_{\tilde{\mathbf{H}}_m^k, j}$, and $\hat{\sigma}_{\tilde{\mathbf{H}}_m^k, j}^2$ denote the j -th element of α^k , $\hat{\boldsymbol{\mu}}_{\tilde{\mathbf{H}}_m^k}$, and $\hat{\boldsymbol{\sigma}}_{\tilde{\mathbf{H}}_m^k}^2$, respectively.

Fourth, the KL divergence between $q(\alpha^k)$ and $p(\alpha^k)$ can be computed analytically as

$$\begin{aligned}
& D_{\text{KL}}(q(\alpha^k)||p(\alpha^k)) \\
&= D_{\text{KL}}(\mathcal{N}(\tilde{\boldsymbol{\alpha}}^k, \sigma_{\tilde{\boldsymbol{\alpha}}^k}^2 \mathbf{I})||\mathcal{N}(0, \sigma_{\alpha^k}^2 \mathbf{I})) \\
&= -\frac{1}{2} \sum_{j=1}^J \left(1 + \log \frac{\sigma_{\tilde{\boldsymbol{\alpha}}^k}^2}{\sigma_{\alpha^k}^2} - \frac{(\tilde{\alpha}^{k, j} - 0)^2 + \sigma_{\tilde{\boldsymbol{\alpha}}^k}^2}{\sigma_{\alpha^k}^2} \right) \quad (6) \\
&= -\frac{1}{2} \sum_{j=1}^J (1 + \log \sigma_{\tilde{\boldsymbol{\alpha}}^k}^2) - \frac{1}{2} \log 2\pi - \log p(\tilde{\boldsymbol{\alpha}}^k),
\end{aligned}$$

where $\tilde{\alpha}^{k, j}$ denote the j -th element of $\tilde{\boldsymbol{\alpha}}^k$.

Fifth, we substitute the result of Eq. (5) and Eq. (6) into Eq. (4), respectively. Then, we can have

$$\begin{aligned}
& \mathbb{E}_{q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K}|\mathcal{G}_{1:M})} \left[\log p(\mathcal{G}_{1:M}, \tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K}) \right. \\
& \quad \left. - \log q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K}|\mathcal{G}_{1:M}) \right] \\
&= \sum_{m=1}^M \mathbb{E}_{q(\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K|\mathcal{G}_m)} \left[\log p(\mathcal{G}_m|\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K) \right] \quad (7) \\
& \quad - \sum_{m=1}^M \sum_{k=1}^K \mathbb{E}_{q(\alpha^k)} [D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k|\mathcal{G}_m)||p(\tilde{\mathbf{H}}_m^k|\alpha^k))]
\end{aligned}$$

$$\begin{aligned}
& - \sum_{k=1}^K D_{\text{KL}}(q(\boldsymbol{\alpha}^k) || p(\boldsymbol{\alpha}^k)) \\
& = \sum_{m=1}^M \mathbb{E}_{q(\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K)] \\
& - \sum_{m=1}^M \sum_{k=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^k | \tilde{\boldsymbol{\alpha}}^k)) - \sum_{m=1}^M \sum_{k=1}^K \frac{J}{2} \frac{\sigma_{\tilde{\boldsymbol{\alpha}}^k}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} \\
& + \sum_{k=1}^K \left[\frac{1}{2} \sum_{j=1}^J (1 + \log \sigma_{\tilde{\boldsymbol{\alpha}}^k}^2) + \frac{1}{2} \log 2\pi + \log p(\tilde{\boldsymbol{\alpha}}^k) \right] \\
& = \sum_{m=1}^M \left\{ \mathbb{E}_{q(\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K)] \right. \\
& - \sum_{k=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^k | \tilde{\boldsymbol{\alpha}}^k)) \\
& \left. - \underbrace{\sum_{k=1}^K \frac{J}{2} \frac{\sigma_{\tilde{\boldsymbol{\alpha}}^k}^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2}}_{\text{constant}} \right\} \\
& + \sum_{k=1}^K \left[\log p(\tilde{\boldsymbol{\alpha}}^k) + \underbrace{\frac{1}{2} \sum_{j=1}^J (1 + \log \sigma_{\tilde{\boldsymbol{\alpha}}^k}^2) + \frac{1}{2} \log 2\pi}_{\text{constant}} \right].
\end{aligned} \tag{8}$$

Finally, since two constants in the above Eq. (7) can be omitted, we can have

$$\begin{aligned}
& \mathbb{E}_{q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M})} [\log p(\mathcal{G}_{1:M}, \tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K})] \\
& - \log q(\tilde{\mathbf{H}}_{1:M}^1, \tilde{\mathbf{H}}_{1:M}^2, \dots, \tilde{\mathbf{H}}_{1:M}^K, \boldsymbol{\alpha}^{1:K} | \mathcal{G}_{1:M}) \\
& \approx \sum_{m=1}^M \left\{ \mathbb{E}_{q(\tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m^1, \tilde{\mathbf{H}}_m^2, \dots, \tilde{\mathbf{H}}_m^K)] \right. \\
& - \sum_{k=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^k | \tilde{\boldsymbol{\alpha}}^k)) \left. \right\} \\
& + \sum_{k=1}^K \log p(\tilde{\boldsymbol{\alpha}}^k) \\
& = \sum_{m=1}^M \left\{ \mathbb{E}_{q(\tilde{\mathbf{H}}_m | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m)] \right. \\
& - \sum_{k=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^k | \tilde{\boldsymbol{\alpha}}^k)) \left. \right\} \\
& + \sum_{k=1}^K \log p(\tilde{\boldsymbol{\alpha}}^k).
\end{aligned} \tag{9}$$

D Derivation of $\tilde{\alpha}^k$

Since estimating the exact Maximum A Posterior (MAP) of α^k is intractable, inspired by [2], we approximate α^k with the help of ELBO as follows:

$$\begin{aligned}
\tilde{\alpha}^k &= \operatorname{argmax}_{\alpha^k} \log p(\alpha^k | \mathcal{G}_{1:M}) \\
&= \operatorname{argmax}_{\alpha^k} \log \frac{p(\mathcal{G}_{1:M}, \alpha^k)}{p(\mathcal{G}_{1:M})} \\
&= \operatorname{argmax}_{\alpha^k} \log p(\mathcal{G}_{1:M}, \alpha^k) \\
&= \operatorname{argmax}_{\alpha^k} \left(\sum_{m=1}^M \log p(\mathcal{G}_m | \alpha^k) \right) + \log p(\alpha^k) \\
&\approx \operatorname{argmax}_{\alpha^k} \sum_{m=1}^M \left\{ \mathbb{E}_{q(\tilde{\mathbf{H}}_m | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m)] - \sum_{k=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^k | \tilde{\alpha}^k)) \right\} + \log p(\alpha^k) \\
&= \operatorname{argmax}_{\alpha^k} \sum_{m=1}^M \sum_{k=1}^K -D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k | \mathcal{G}_m) || p(\tilde{\mathbf{H}}_m^k | \alpha^k)) + \log p(\alpha^k) \\
&= \operatorname{argmax}_{\alpha^k} \sum_{m=1}^M \sum_{k=1}^K \sum_{j=1}^J \frac{(\hat{\mu}_{\tilde{\mathbf{H}}_m^k, j} - \alpha^{k,j})^2}{\sigma_{\tilde{\mathbf{H}}_m^k}^2} - \sum_{j=1}^J \frac{(\alpha^{k,j} - 0)^2}{\sigma_{\alpha^k}^2} \\
&= \operatorname{argmax}_{\alpha^k} f(\alpha^k),
\end{aligned} \tag{10}$$

where $f(\alpha^k)$ is a concave quadratic function with only one maximum point, and $\alpha^{k,j}$ denotes the j -th element of α^k . The closed-form solution of $\tilde{\alpha}^k$ can be derived by differentiating Eq. (10) with respect to α^k . Specifically, we let

$$\left. \frac{\partial f(\alpha^k)}{\partial \alpha^k} \right|_{\alpha^k = \tilde{\alpha}^k} = 0, \tag{11}$$

and then we can have

$$\tilde{\alpha}^k = \frac{\sum_{m=1}^M \hat{\mu}_{\tilde{\mathbf{H}}_m^k}}{M + \frac{\sigma_{\tilde{\mathbf{H}}_m^k}^2}{\sigma_{\alpha^k}^2}}. \tag{12}$$

E Details of HVGAE

In this section, we introduce the detailed neural network architectures for our proposed HVGAE. Since HVGAE is deployed in each client, we take the m -th client as an example. Recall that \mathcal{G}_m is the subgraph on the m -th client, which contains the node feature matrix \mathbf{X}_m and adjacency matrix \mathbf{A}_m . We use two DisenGCNs (*i.e.*, $\text{DisenGCN}_{\mu_m}(\mathcal{G}_m)$ and $\text{DisenGCN}_{\sigma_m}(\mathcal{G}_m)$) as the encoder and an inner product as the decoder of HVGAE, respectively. Note that $\text{DisenGCN}_{\mu_m}(\mathcal{G}_m)$ and $\text{DisenGCN}_{\sigma_m}(\mathcal{G}_m)$ are used to infer the means and standard deviations of \mathcal{G}_m for K latent factors, respectively. Moreover, $\text{DisenGCN}_{\mu}(\mathcal{G}_m)$ and $\text{DisenGCN}_{\sigma}(\mathcal{G}_m)$ share the same node feature projection layer.

Fig. 1 shows the architecture of our proposed HVGAE. First, we project the node feature to K subspaces according to Eq. (1) in the main manuscript. Meanwhile, the node representations after the node feature projection layer are used to train a local node classifier. Second, we use the neighborhood routing mechanism to obtain $\mathbf{H}_{m,\mu}^{1:K}$ and $\mathbf{H}_{m,\sigma}^{1:K}$, respectively. Third, we use the reparameterization trick [3] to sample $\tilde{\mathbf{H}}_m^{1:K}$ from $\mathbf{H}_{m,\mu}^{1:K}$ and $\mathbf{H}_{m,\sigma}^{1:K}$ (see Eq. (7) in the main manuscript). Fourth, HVGAE decodes from $\tilde{\mathbf{H}}_m^{1:K}$ and then computes $\mathbb{E}_{q(\tilde{\mathbf{H}}_m | \mathcal{G}_m)} [\log p(\mathcal{G}_m | \tilde{\mathbf{H}}_m)]$ (see Eq. (8) in the main manuscript). The implementations of two prior distributions $p(\tilde{\alpha}^k)$ and $p(\tilde{\mathbf{H}}_m^k | \tilde{\alpha}^k)$ are shown in the Appendix I.6.

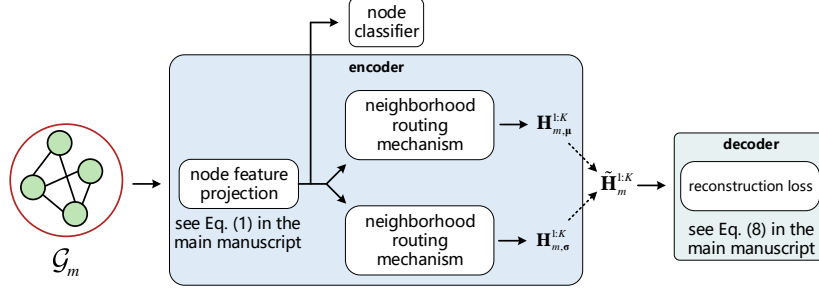


Figure 1: The architecture of our proposed HVGAE. The dashed lines show the process of sampling using the reparameterization trick [3].

F Additional Related Work

Here we review the typical works related to the FL via Bayesian methods. Recently, some methods [4, 5, 6] have tried to model the FL problem by using the Bayesian theories. Specifically, they take the network weights as a whole entity and treat them as a single random variable shared by all clients. For example, FedPA [4] infers the global posterior by averaging the local posteriors. However, FedPA targets a general FL setting, so it does not apply to the personalized FL, and its performance may be degraded. Different from FedPA, pFedBayes [5] assigns a personalized Bayesian neural network to each client. It infers the global posterior from individual posteriors under a regularizer based on the KL divergence. However, it cannot model the posterior of each client from a global perspective. To address this deficiency, Kim *et. al* propose a hierarchical Bayesian approach called FedHB [6], which mitigates the heterogeneity problem in a personalized way. Specifically, it introduces two types of latent random variables, one used as the network weights for each client’s backbone, and the other used as a globally shared random variable to be associated with each client. Unlike FedHB, we model the local subgraph data rather than the network weight, so that we can infer the subgraph data distribution on each client.

G Analysis of Federated Parameters

According to the architecture of our proposed HVGAE, there are several components: a node feature projection layer, two neighborhood routing mechanism layers, a node classifier, and a reconstruction layer. First, recall that there are no learnable parameters in the neighborhood routing mechanism. Second, the reconstruction layer is constructed by an inner product, so there are no learnable parameters. Third, although we use two DisenGCNs (*i.e.*, $\text{DisenGCN}_{\mu_m}(\mathcal{G}_m)$ and $\text{DisenGCN}_{\sigma_m}(\mathcal{G}_m)$) as the encoder of our proposed HVGAE, they share the same node feature projection layer (see Appendix E). Consequently, according to the node feature projection layer defined by Eq. (1) in the main manuscript, we can find that $\mathbf{W}_m^1, \mathbf{W}_m^2, \dots, \mathbf{W}_m^K$ and $\mathbf{b}_m^1, \mathbf{b}_m^2, \dots, \mathbf{b}_m^K$ are learnable parameters that should be federated. Last but not least, since we introduce a node classifier, which is actually a Multi-Layer Perceptron (MLP), we can find that $\mathbf{W}_m^{\text{cls}} \in \mathbb{R}^{c \times d_{\text{out}}}$ and $\mathbf{b}_m^{\text{cls}} \in \mathbb{R}^c$ are learnable parameters that should be federated.

In Eq. (6) of the main manuscript, although we have defined the learnable parameter $\hat{\alpha}_m^k$ on each client, $\hat{\alpha}_m^k$ is only used to approximate the data distribution instead of a parameter in a neural network. Consequently, $\hat{\alpha}_m^k$ does not participate in the federation.

H Pseudocode of FedIIIH

In this section, we show the pseudocode of our proposed FedIIIH for the clients and server in Algorithm 1 and Algorithm 2, respectively. Moreover, the framework of our proposed FedIIIH from the perspective of the clients and the server is shown in Fig. 2 and Fig. 3, respectively. Inspired by [7], we regard the posterior mean of α^k for the k -th global latent factor (*i.e.*, $\tilde{\alpha}^k$) in the last round as the prior for the k -th local latent factor on each client in the current round.

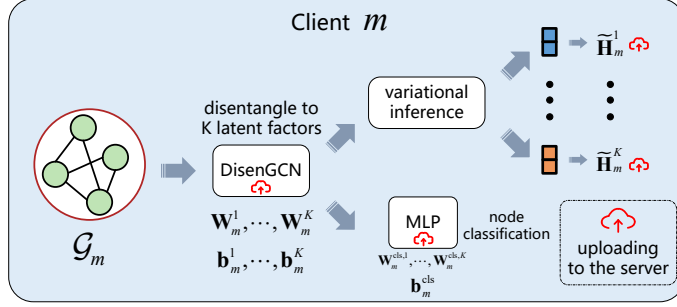


Figure 2: The framework of our proposed FedIIIH from the perspective of the clients. Let us take the client m as an example.

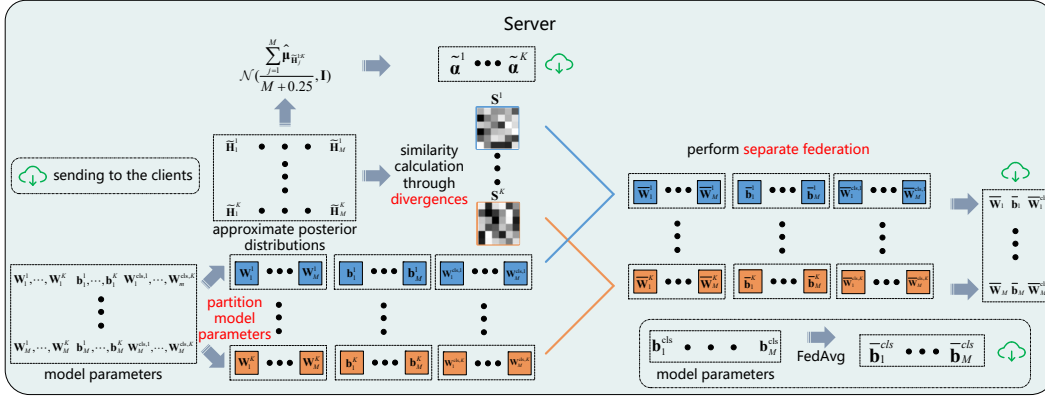


Figure 3: The framework of our proposed FedIIIH from the perspective of the server, where S^k denotes the similarity matrix with respect to the k -th latent factor. Here we use blue and orange to represent the first and the K -th latent factors, respectively.

I Implementation Details

In this section, we provide the implementation details, including the experimental platform, the dataset descriptions, the details of subgraph partitioning, the information on baseline methods, training details, and the implementations of two prior distributions (*i.e.*, $p(\tilde{\alpha}^k)$ and $p(\tilde{\mathbf{H}}_m^k | \tilde{\alpha}^k)$).

I.1 Experimental Platform

All the experiments in this work are conducted on a Linux server with a 2.90 GHz Intel Xeon Gold 6326 CPU, 64 GB of RAM, and two NVIDIA GeForce RTX 4090 GPUs with 48GB of memory. Our proposed method is implemented via Python 3.8.8, PyTorch 1.12.0, and PyTorch Geometric (PyG) 2.3.0 with the Adam optimizer.

I.2 Datasets

To validate the effectiveness of our proposed FedIIIH, we perform extensive experiments on eleven widely used benchmark datasets, including six homophilic and five heterophilic graph datasets. In the homophilic graph datasets, there are *Cora*, *CiteSeer*, *PubMed*, and *ogbn-arxiv* for the citation graphs; *Amazon-Computer* and *Amazon-Photo* for Amazon product graphs. In the heterophilic graph datasets [9], there are *Roman-empire*, *Amazon-ratings*, *Minesweeper*, *Tolokers*, and *Questions*. The statistical information of the above benchmark datasets is described in Tab. 2. Note that we use the Area Under the ROC curve (AUC) as the evaluation metric (higher values are better) for the *Minesweeper*, *Tolokers*, and *Questions* datasets, and use the accuracy as the evaluation metric (higher values are better) for other datasets.

Algorithm 1: FedIIIH Client Algorithm

- Input:** Number of local epochs E ; number of latent factors K ; subgraph \mathcal{G}_m on client m ; node feature matrix \mathbf{X}_m on client m ; label matrix \mathbf{Y}_m on client m ; parameters of two DisenGCNs $\mathbf{W}_m^{1:K}$ and $\mathbf{b}_m^{1:K}$ on client m ; parameters of the node classifier $\mathbf{W}_m^{\text{cls}}$ and $\mathbf{b}_m^{\text{cls}}$ on client m ; federated parameters $\overline{\mathbf{W}}_m^{\text{cls}}$, $\overline{\mathbf{b}}_m^{\text{cls}}$, $\overline{\mathbf{W}}_m^{1:K}$, and $\overline{\mathbf{b}}_m^{1:K}$ from the server; global latent factors $\tilde{\boldsymbol{\alpha}}^{1:K}$ from the server.
- Output:** Predicted label for each unlabeled graph node in subgraph \mathcal{G}_m .
- 1 Download federated parameters $\overline{\mathbf{W}}_m^{\text{cls}}$, $\overline{\mathbf{b}}_m^{\text{cls}}$, $\overline{\mathbf{W}}_m^{1:K}$, and $\overline{\mathbf{b}}_m^{1:K}$ from the server;
 - 2 Download the global latent factors $\tilde{\boldsymbol{\alpha}}^{1:K}$ from the server;
 - 3 $\mathbf{W}_m^{\text{cls}} \leftarrow \overline{\mathbf{W}}_m^{\text{cls}}$, $\mathbf{b}_m^{\text{cls}} \leftarrow \overline{\mathbf{b}}_m^{\text{cls}}$, $\mathbf{W}_m^{1:K} \leftarrow \overline{\mathbf{W}}_m^{1:K}$, $\mathbf{b}_m^{1:K} \leftarrow \overline{\mathbf{b}}_m^{1:K}$;
 - 4 **for** each local epoch e from 1 to E **do**
 - 5 Project the node feature into K subspaces via Eq. (1) in the main manuscript;
 - 6 Optimize $\mathbf{W}_m^{\text{cls}}$ and $\mathbf{b}_m^{\text{cls}}$ via the cross-entropy loss;
 - 7 Disentangle projected node representations into K latent factors via the neighborhood routing mechanism, and then obtain $\mathbf{H}_{m,\mu}^{1:K}$ and $\mathbf{H}_{m,\sigma}^{1:K}$, respectively;
 - 8 Sample $\tilde{\mathbf{H}}_m^{1:K}$ from $\mathbf{H}_{m,\mu}^{1:K}$ and $\mathbf{H}_{m,\sigma}^{1:K}$ via Eq. (7) in the main manuscript so as to obtain $q(\tilde{\mathbf{H}}_m^{1:K}|\mathcal{G}_m)$;
 - 9 Compute $\mathbb{E}_{q(\tilde{\mathbf{H}}_m|\mathcal{G}_m)}[\log p(\mathcal{G}_m|\tilde{\mathbf{H}}_m)]$ via Eq. (8) in the main manuscript;
 - 10 Approximate $\tilde{\boldsymbol{\alpha}}^{1:K}$ by using $\hat{\boldsymbol{\alpha}}_m^{1:K}$ so as to compute $\sum_{k=1}^K \left\{ \log p(\hat{\boldsymbol{\alpha}}_m^k) - D_{\text{KL}}(p(\hat{\boldsymbol{\alpha}}_m^k)||p(\tilde{\boldsymbol{\alpha}}^k)) \right\}$;
 - 11 Compute $\sum_{k=1}^K D_{\text{KL}}(q(\tilde{\mathbf{H}}_m^k|\mathcal{G}_m)||p(\tilde{\mathbf{H}}_m^k|\tilde{\boldsymbol{\alpha}}^k))$, where $p(\tilde{\mathbf{H}}_m^{1:K}|\tilde{\boldsymbol{\alpha}}^{1:K}) \sim \mathcal{N}(\tilde{\boldsymbol{\alpha}}^{1:K}, \sigma_{\tilde{\mathbf{H}}_m^{1:K}}^2 \mathbf{I})$;
 - 12 Optimize $\mathbf{W}_m^{1:K}$ and $\mathbf{b}_m^{1:K}$ via Eq. (6) in the main manuscript;
 - 13 **end**
 - 14 Upload $\mathbf{W}_m^{\text{cls}}$, $\mathbf{b}_m^{\text{cls}}$, $\mathbf{W}_m^{1:K}$, $\mathbf{b}_m^{1:K}$, and $\tilde{\mathbf{H}}_m^{1:K}$ to the server;
 - 15 Predict labels based on the trained node classifier.
-

In order to facilitate the division of datasets, a random sample of 20% of nodes is selected for training purposes, 40% for the purpose of validation, and 40% for testing, with the exception of the *ogbn-arxiv* dataset. This is due to the fact that the *ogbn-arxiv* dataset comprises a relatively large number of nodes in comparison to the other datasets, as reported in Tab. 2. Consequently, for the *ogbn-arxiv* dataset, a random sample of 5% of the nodes is used for training, while the remaining half of the nodes are used for validation and testing, respectively.

I.3 Subgraph Partitioning

Inspired by real-world requirements and following [10], we consider two subgraph partitioning settings: non-overlapping and overlapping. In the non-overlapping setting, $\cup_{m=1}^M \mathcal{V}_m = \mathcal{V}$ and $\mathcal{V}_m \cap \mathcal{V}_n = \emptyset$ for $\forall m \neq n \in \{1, 2, \dots, M\}$, where \mathcal{V} represents the node set of the global graph. Partitioning without this property is called overlapping. Here we present the details of how to partition the original graph into multiple subgraphs. It should be noted that the number of subgraphs is equal to the number of clients. Both non-overlapping and overlapping subgraph partitioning settings are used in the experiments for all datasets.

Non-overlapping partitioning First, if there are M clients, the number of non-overlapping subgraphs to be generated is specified as M . Second, the METIS graph partitioning algorithm, as described in [11], is used to divide the original graph into M subgraphs. In other words, the non-overlapping partitioning subgraph for each client is directly obtained by the output of the METIS algorithm.

Overlapping partitioning First, if there are M clients, the number of overlapping subgraphs to be generated is specified as M . Second, the METIS [11] graph partitioning algorithm is used to divide

Algorithm 2: FedIIIH Server Algorithm

Input: Number of rounds R ; number of clients M ; number of latent factors K ; parameters of two DisenGCNs $\mathbf{W}_m^{1:K}$ and $\mathbf{b}_m^{1:K}$ from client m ; parameters of the node classifier $\mathbf{W}_m^{\text{cls}}$ and $\mathbf{b}_m^{\text{cls}}$ from client m ; approximate posterior distributions $\tilde{\mathbf{H}}_m^{1:K}$ from client m .

Output: Federated model parameters for client m .

```
1 Initialize parameters  $(\overline{\mathbf{W}}^{\text{cls}})^{(1)}$ ,  $(\overline{\mathbf{b}}^{\text{cls}})^{(1)}$ ,  $(\overline{\mathbf{W}}^{1:K})^{(1)}$ , and  $(\overline{\mathbf{b}}^{1:K})^{(1)}$ ;
2 Initialize  $(\tilde{\boldsymbol{\alpha}}^{1:K})^{(1)}$ ;
3 for each round  $r$  from 1 to  $R$  do
4   for client  $m \in \{1, 2, \dots, M\}$  in parallel do
5     if  $r = 1$  then
6       Send  $(\overline{\mathbf{W}}^{\text{cls}})^{(r)}$ ,  $(\overline{\mathbf{b}}^{\text{cls}})^{(r)}$ ,  $(\overline{\mathbf{W}}^{1:K})^{(r)}$ , and  $(\overline{\mathbf{b}}^{1:K})^{(r)}$  to client  $m$ ;
7       Send  $(\tilde{\boldsymbol{\alpha}}^{1:K})^{(r)}$  to client  $m$ ;
8     end
9     else
10      Receive  $\mathbf{W}_m^{\text{cls}}$ ,  $\mathbf{b}_m^{\text{cls}}$ ,  $\mathbf{W}_m^{1:K}$ ,  $\mathbf{b}_m^{1:K}$ , and  $\tilde{\mathbf{H}}_m^{1:K}$  from client  $m$ ;
11      Compute  $S(m, j)^{1:K}$  via Eq. (9) in the main manuscript, where  $j \in \{1, 2, \dots, M\}$ ;
12      Compute  $\beta_{m,j}^{1:K}$  via Eq. (11) in the main manuscript, where  $j \in \{1, 2, \dots, M\}$ ;
13      Perform the separate federation to obtain  $(\overline{\mathbf{W}}_m^{\text{cls}})^{(r)}$ ,  $(\overline{\mathbf{W}}_m^{1:K})^{(r)}$ , and  $(\overline{\mathbf{b}}_m^{1:K})^{(r)}$  via
        Eq. (10) in the main manuscript;
14      Perform the FedAvg [8] to obtain  $(\overline{\mathbf{b}}_m^{\text{cls}})^{(r)}$ ;
15      Sample  $(\tilde{\boldsymbol{\alpha}}^{1:K})^{(r)}$  from  $\mathcal{N}(\frac{\sum_{j=1}^M \hat{\boldsymbol{\mu}}_j^{1:K}}{M+0.25}, \mathbf{I})$ ;
16      Send  $(\overline{\mathbf{W}}_m^{\text{cls}})^{(r)}$ ,  $(\overline{\mathbf{b}}_m^{\text{cls}})^{(r)}$ ,  $(\overline{\mathbf{W}}_m^{1:K})^{(r)}$ , and  $(\overline{\mathbf{b}}_m^{1:K})^{(r)}$  to client  $m$ ;
17      Send  $(\tilde{\boldsymbol{\alpha}}^{1:K})^{(r)}$  to client  $m$ ;
18    end
19    Perform Algorithm 1 on client  $m$ ;
20  end
21 end
```

Table 2: Statistical information of eleven used graph datasets.

Types	Datasets	# Nodes	# Edges	# Classes	# Node Features
homophilic graph	<i>Cora</i>	2,708	5,429	7	1,433
	<i>CiteSeer</i>	3,327	4,732	6	3,703
	<i>PubMed</i>	19,717	44,324	3	500
	<i>Amazon-Computer</i>	13,752	491,722	10	767
	<i>Amazon-Photo</i>	7,650	238,162	8	745
	<i>ogbn-arxiv</i>	169,343	2,315,598	40	128
heterophilic graph	<i>Roman-empire</i>	22,662	32,927	18	300
	<i>Amazon-ratings</i>	24,492	93,050	5	300
	<i>Minesweeper</i>	10,000	39,402	2	7
	<i>Tolokers</i>	11,758	519,000	2	10
	<i>Questions</i>	48,921	153,540	2	301

the original graph into $\lfloor \frac{M}{5} \rfloor$ subgraphs. Third, in each subgraph generated by METIS, half of the nodes and their associated edges are randomly sampled. This procedure is performed five times to generate five different yet overlapped subgraphs. By doing so, the number of overlapping subgraphs is equal to the number of clients.

I.4 Baseline Methods

We compare our proposed FedIIIH with the following baseline methods, which can be categorized into two groups. The first group comprises general FL baseline methods, including FedAvg [8],

FedProx [12], and FedPer [13]. The second group consists of six GFL baseline methods, namely GCFL [14], FedGNN [15], FedSage+ [16], FED-PUB [10], FedGTA [17], and AdaFGL [18]. Moreover, we perform experiments with local training, that is, training each client without federated aggregation. The detailed descriptions of these baseline methods are provided below.

FedAvg This method [8] represents one of the fundamental baseline methods in the field of FL. First, each client independently trains a model, which is subsequently transmitted to a server. Then, the server aggregates the locally updated models by averaging and transmits the aggregated model back to the clients.

FedProx This method [12] is one of the personalized FL baseline methods. It customizes a personalized model for each client by adding a proximal term as a subproblem that minimizes weight differences between local and global models.

FedPer This method [13] is one of the personalized FL baseline methods. It only federates the weights of the backbone while training the personalized classification layer in each client.

GCFL This method [14] is one of the basic GFL methods. Specifically, GCFL is designed for vertical GFL (*e.g.*, GFL for molecular graphs). In particular, it uses the bi-partitioning scheme, which divides a set of clients into two disjoint groups of clients based on the similarity of their gradients. This is similar to the mechanism proposed for image classification in clustered-FL [19]. Then, after partitioning, the model weights are shared only among clustered clients with similar gradients.

FedGNN This method [15] is one of the GFL baseline methods. It extends local subgraphs by exchanging node embeddings from other clients. Specifically, if two nodes in two different clients have exactly the same neighbors, FedGNN transfers the nodes with the same neighbors from other clients and expands them.

FedSage+ This method [16] is one of the GFL baseline methods. It generates the missing edges between subgraphs and the corresponding neighbor nodes by using the missing neighbor generator. To train this neighbor generator, each client first receives node representations from other clients, and then computes the gradient of the distances between the local node features and the node representations of the other clients. After that, the gradient is sent back to the other clients, and this gradient is then used to train the neighbor generator.

FED-PUB This method [10] is one of the personalized GFL baseline methods. It estimates the similarities between the subgraphs based on the outputs of the local models that are given the same test graph. Then, based on the similarities, it performs a weighted averaging of the local models for each client. In addition, it learns a personalized sparse mask at each client in order to select and update only the subgraph-relevant subset of the aggregated parameters.

FedGTA This method [17] is one of the personalized GFL baseline methods. In this method, each client first computes topology-aware local smoothing confidence and mixed moments of neighbor features. They are then used to compute the inter-subgraph similarities, which are uploaded to the server along with the model parameters. Finally, the server is able to perform a weighted federation for each client.

AdaFGL This method [18] is one of the personalized GFL baseline methods. Actually, it is a decoupled two-step personalized approach. First, it uses standard multi-client federated collaborative training to acquire the federated knowledge extractor by aggregating uploaded models in the final round at the server. Second, each client performs personalized training based on the local subgraph and the federated knowledge extractor.

Local This method is the non-FL baseline, where the model is trained only locally for each client, with no weight sharing.

I.5 Training Details

Training rounds and epochs For the *Cora*, *CiteSeer*, *PubMed*, *Roman-empire*, *Amazon-ratings*, *Minesweeper*, *Tolokers*, and *Questions* datasets, we set the number of local training epochs and total rounds to 1 and 100, respectively. For larger datasets, such as *Amazon-Computer*, *Amazon-Photo*, and *ogbn-arxiv*, we set the number of total rounds to 200. Note that the number of local epochs is set to 2 for the *Amazon-Photo* and *ogbn-arxiv* datasets, and to 3 for the *Amazon-Computer* dataset. Finally, we report the test performance of all models at the best validation epoch, and the performance is measured by averaging across all clients in terms of node classification accuracies (or AUCs).

Hyperparameters We report the detailed hyperparameters used to train our proposed FedIIIH in Tab. 3 and Tab. 4. These hyperparameters are determined by grid search. Note that we set the similarity scaling hyperparameter (*i.e.*, τ in Eq. (11) of the main manuscript) to 10 in both non-overlapping and overlapping subgraph partitioning scenarios according to the recommendation of FED-PUB [10]. We provide the detailed reasons in the Appendix K.7. The search range of hyperparameters is shown in Tab. 5. The code of our proposed FedIIIH will be released on GitHub upon acceptance of the paper. The detailed hyperparameter sensitivity analysis can be found in the Appendix M.2.

Hyperparameter tuning guidelines There are four vital hyperparameters (*i.e.*, number of latent factors, number of neighborhood routing layers, number of neighborhood routing iterations, and τ) in our proposed FedIIIH. First, the number of latent factors (*i.e.*, K) is related to whether the graph tends to be homophilic or heterophilic. If the graph dataset tends to be homophilic, a small K is recommended. If the graph dataset tends to be heterophilic, a large K is recommended. Second, the variations in performance under different numbers of neighborhood routing layers, different numbers of neighborhood routing iterations, and different values of τ are all small. Therefore, these three hyperparameters can be tuned by grid search.

Network architectures For the experiments of all baseline methods, except FedSage+, FedGTA, and our proposed FedIIIH, we use two layers of the Graph Convolutional Network (GCN) [20] and a linear classifier layer as their network architectures. For the hyperparameter settings of the baseline methods, we use the default settings given in their original papers. Because of the inductive and scalability advantages of GraphSAGE [21], FedSage+ uses GraphSAGE as the encoder and then trains a missing neighbor generator to handle missing links across local subgraphs. For our proposed FedIIIH, we use the node feature projection layer of DisenGCN [1] to obtain the node representations (see Fig. 1) and a linear classifier layer (*i.e.*, MLP) to perform node classifications. In contrast, FedGTA uses a Graph Attention Multi-Layer Perceptron (GAMLP) [22] as its backbone and a linear classifier layer to classify nodes. Note that GAMLP [22] is one of the scalable Graph Neural Network (GNN) models, which can capture the underlying correlations between different scales of graph knowledge.

I.6 Implementations of Two Prior Distributions

Here we describe the detailed implementations of $p(\tilde{\alpha}^k)$ and $p(\tilde{\mathbf{H}}_m^k|\tilde{\alpha}^k)$, respectively. Since $\tilde{\alpha}^k$ denotes the posterior mean of α^k for the k -th global latent factor, we assume that $p(\tilde{\alpha}^k) \sim \mathcal{N}(\tilde{\alpha}^k, \sigma_{\alpha^k}^2 \mathbf{I})$, where $\tilde{\alpha}^k$ is given in Eq. (12). Moreover, $\sigma_{\alpha^k}^2$ and $\sigma_{\tilde{\mathbf{H}}_m^k}^2$ is set to 1 and 0.25, respectively. In other words, $p(\tilde{\alpha}^k) \sim \mathcal{N}(\frac{\sum_{m=1}^M \hat{\mu}_{\tilde{\mathbf{H}}_m^k}}{M+0.25}, \mathbf{I})$. According to Tab. 1, $p(\tilde{\mathbf{H}}_m^k|\alpha^k) \sim \mathcal{N}(\alpha^k, \sigma_{\tilde{\mathbf{H}}_m^k}^2 \mathbf{I})$. Consequently, we can have

$$\begin{aligned} p(\tilde{\mathbf{H}}_m^k|\tilde{\alpha}^k) &\sim \mathcal{N}(\tilde{\alpha}^k, 0.25\mathbf{I}) \\ &\sim \mathcal{N}\left(\frac{\sum_{m=1}^M \hat{\mu}_{\tilde{\mathbf{H}}_m^k}}{M+0.25}, 0.25\mathbf{I}\right). \end{aligned} \quad (13)$$

J Additional Experiments

In this section, we provide the additional experiments related to the ablation studies and hyperparameter analysis, respectively.

Table 3: Hyperparameters used in our proposed FedIIH on the **homophilic** graph datasets in both the non-overlapping and overlapping subgraph partitioning settings.

Hyperparameters	Cora		CiteSeer		PubMed	
	non-overlapping	overlapping	non-overlapping	overlapping	non-overlapping	overlapping
# latent factors	2	2	2	2	4	2
learning rate	0.02	0.01	0.01	0.01	0.01	0.015
# hidden dimensions	128	256	256	256	256	256
dropout rate	0.3	0.35	0.35	0.35	0.25	0.4
weight decay	0.005	1e-6	1e-6	1e-6	0.0045	1e-6
# neighborhood routing layers	4	5	5	5	1	1
# neighborhood routing iterations	6	6	6	6	6	6
	Amazon-Computer		Amazon-Photo		ogbn-arxiv	
Hyperparameters	non-overlapping	overlapping	non-overlapping	overlapping	non-overlapping	overlapping
# latent factors	6	4	6	10	2	2
learning rate	0.015	0.015	0.015	0.01	0.01	0.01
# hidden dimensions	128	128	256	128	128	128
dropout rate	0.4	0.35	0.4	0.35	0.35	0.35
weight decay	1e-6	1e-6	1e-6	1e-6	1e-6	1e-6
# neighborhood routing layers	1	5	1	1	5	5
# neighborhood routing iterations	6	6	6	5	6	6

Table 4: Hyperparameters used in our proposed FedIIH on the **heterophilic** graph datasets in both the non-overlapping and overlapping subgraph partitioning settings.

Hyperparameters	Roman-empire		Amazon-ratings		Minesweeper	
	non-overlapping	overlapping	non-overlapping	overlapping	non-overlapping	overlapping
# latent factors	4	4	4	4	6	4
learning rate	0.015	0.015	0.01	0.01	0.01	0.01
# hidden dimensions	128	128	128	256	256	128
dropout rate	0.35	0.35	0.35	0.35	0.35	0.35
weight decay	1e-6	1e-6	1e-6	1e-6	1e-6	1e-6
# neighborhood routing layers	1	1	3	5	5	6
# neighborhood routing iterations	6	6	7	6	6	7
	Tolokers		Questions			
Hyperparameters	non-overlapping	overlapping	non-overlapping	overlapping		
# latent factors	10	10	8	2		
learning rate	0.01	0.01	0.01	0.01		
# hidden dimensions	128	128	256	256		
dropout rate	0.35	0.35	0.35	0.35		
weight decay	0.0045	0.0045	1e-6	1e-6		
# neighborhood routing layers	1	1	5	5		
# neighborhood routing iterations	2	2	6	6		

J.1 Ablation Studies on Other Datasets

To further analyze the contribution of each component, we conduct ablation studies on the remaining datasets in both non-overlapping and overlapping partitioning settings with 10 clients. As shown in Tab. 6, we can find that the performance of FedIIH is significantly better than the three variants. It validates that each component indeed contributes a lot to the final performance.

J.2 Hyperparameter Sensitivity Analysis of K

The hyperparameter sensitivity analysis of K on the remaining datasets are shown in Fig. 5, Fig. 6, ..., to Fig. 13. According to the experimental results, we have the following observations and insights: 1) In general, the performance variation under different K is small except for the *Roman-empire* dataset (see Fig. 10). The *Roman-empire* dataset has low homophily [9], which is based on the Roman Empire article from the English Wikipedia. Each node in the graph corresponds to a (non-unique) word in the text, and the node label is determined by the syntactic role of the word. Due to the syntactic dependencies within neighboring words, there exists strong heterogeneity within the *Roman-empire* graph. Consequently, if we ignore this intra-heterogeneity (*i.e.*, without disentangling), the performance will decrease significantly on the *Roman-empire* dataset. This is consistent with

Table 5: The search range of the hyperparameters used in our proposed FedIIIH.

Hyperparameters	# latent factors	learning rate	# hidden dimensions	dropout rate
Range	{1, 2, 4, 6, 8, 10}	{0.01, 0.015, 0.02}	{128, 256}	[0.25, 0.4]
Hyperparameters	weight decay	# neighborhood routing layers	# neighborhood routing iterations	
Range	[1e-6, 5e-3]	{1, 2, 3, 4, 5, 6}	{2, 3, 4, 5, 6, 7}	

Table 6: Ablation studies in both non-overlapping and overlapping partitioning settings on other datasets with 10 clients.

	CiteSeer		PubMed		Amazon-Computer	
Methods	non-overlapping	overlapping	non-overlapping	overlapping	non-overlapping	overlapping
FedIIIH (w/o HM)	74.91±0.27 (↓1.59)	72.29±0.16 (↓0.87)	85.19±0.05 (↓2.46)	85.16±0.17 (↓0.71)	85.75±0.90 (↓5.11)	87.71±0.20 (↓2.44)
FedIIIH (w/o VI)	72.27±2.16 (↓4.23)	72.60±0.16 (↓0.56)	81.35±0.20 (↓6.30)	84.56±0.04 (↓1.31)	69.16±1.15 (↓21.70)	74.90±1.24 (↓15.25)
FedIIIH (w/o Dis)	75.71±0.38 (↓0.79)	71.54±0.12 (↓1.62)	85.71±0.12 (↓1.94)	84.30±0.03 (↓1.57)	88.96±0.08 (↓1.90)	88.51±0.04 (↓1.64)
FedIIIH	76.50±0.06	73.16±0.18	87.65±0.18	85.87±0.03	90.86±0.23	90.15±0.04
	Amazon-Photo		ogbn-arxiv		Roman-empire	
Methods	non-overlapping	overlapping	non-overlapping	overlapping	non-overlapping	overlapping
FedIIIH (w/o HM)	91.84±0.05 (↓2.38)	91.88±0.27 (↓1.50)	65.18±0.41 (↓4.16)	63.54±0.15 (↓3.15)	56.28±0.20 (↓10.16)	60.90±0.21 (↓4.58)
FedIIIH (w/o VI)	79.25±0.96 (↓14.97)	82.16±0.56 (↓11.22)	49.26±0.74 (↓20.08)	46.20±1.73 (↓20.49)	57.38±0.18 (↓9.06)	61.69±0.09 (↓3.79)
FedIIIH (w/o Dis)	92.43±0.01 (↓1.79)	92.01±0.04 (↓1.37)	61.51±0.15 (↓7.83)	60.64±0.16 (↓6.05)	40.51±0.27 (↓25.93)	42.84±0.09 (↓22.64)
FedIIIH	94.22±0.08	93.38±0.00	69.34±0.02	66.69±0.09	66.44±0.28	65.48±0.12
	Minesweeper		Tolokers		Questions	
Methods	non-overlapping	overlapping	non-overlapping	overlapping	non-overlapping	overlapping
FedIIIH (w/o HM)	70.88±0.01 (↓2.35)	66.56±0.07 (↓2.79)	64.56±0.17 (↓6.76)	69.21±0.22 (↓2.46)	65.90±0.09 (↓2.09)	67.77±0.10 (↓1.02)
FedIIIH (w/o VI)	71.34±0.09 (↓1.89)	68.47±0.06 (↓0.88)	64.34±0.12 (↓6.98)	68.03±0.34 (↓3.64)	66.63±0.08 (↓1.36)	68.24±0.06 (↓0.55)
FedIIIH (w/o Dis)	70.67±0.02 (↓2.56)	68.75±0.19 (↓0.60)	62.53±0.29 (↓8.79)	68.10±0.12 (↓3.57)	66.42±0.01 (↓1.57)	67.52±0.17 (↓1.27)
FedIIIH	73.23±0.04	69.35±0.25	71.32±0.09	71.67±0.02	67.99±0.09	68.79±0.09

the experimental results in Tab. 3 and Tab. 4 in the main manuscript, where our FedIIIH outperforms other methods by a large margin. 2) As the value of K increases, the performance may increase or decrease depending on the datasets. For example, on the *Roman-empire* (see Fig. 10) and *Tolokers* (see Fig. 4 in the main manuscript) datasets, their performances increase consistently as the value of K increases. On the contrary, on the *Cora* dataset (see Fig. 4), the accuracy reaches its highest value when $K = 2$, and then it decreases as the value of K is further increased.

K Discussions

K.1 Similarity Heatmaps of FED-PUB and FedIIIH over Three Independent Runs

As shown in Fig. 14, Fig. 15, Fig. 16, and Fig. 17, we present the similarity heatmaps of FED-PUB and our FedIIIH over three independent runs on the *Cora* and *Amazon-ratings* datasets in the overlapping setting with 20 clients, respectively. We can find that our calculated similarities are fairly much more stable than the similarities calculated by FED-PUB. This is because FED-PUB estimates the similarities between subgraphs based on the outputs of local models given the same random test graph. Since the random test graph varies over three independent runs, the outputs of the local models also change. In contrast, our FedIIIH successfully infers the whole distribution of subgraph data in a multi-level global perspective, such that we can stably characterize the inter-subgraph similarities. Note that the similarity ground truth of the *Cora* and *Amazon-ratings* datasets in the overlapping setting with 20 clients are presented in Fig. 18a and Fig. 3a of the main manuscript, respectively.

K.2 Similarity Heatmaps on Other Datasets

The similarity heatmaps on other datasets are shown in Fig. 18, Fig. 19, ..., to Fig. 38. We present the similarity heatmaps for each dataset (except the *Cora* dataset) in two settings, namely, non-overlapping with 20 clients and overlapping with 30 clients. However, for the *Cora* dataset, we present the similarity heatmaps in the overlapping setting with 20 clients. This is because in [10], Baek *et al.* present the heatmaps on the *Cora* dataset in the overlapping setting with 20 clients, and we specifically want to be consistent with that here. According to the experimental results on these datasets, we can find that the similarity heatmaps of our FedIIIH are always much closer to the ground truth than FED-PUB and FedGTA, verifying the effectiveness of our similarity calculation scheme based on the inferred subgraph data distribution.




Figure 4: Sensitivity of the number of latent factors K on the *Cora* dataset.




Figure 5: Sensitivity of the number of latent factors K on the *CiteSeer* dataset.

K.3 Case Study of Different Latent Factors

Since our proposed FedIIIH disentangles the subgraph into several latent factors, one may ask what is the difference between the similarity heatmaps under different latent factors. To illustrate this, we take the *Roman-empire* dataset as an example to perform a case study. As demonstrated in Fig. 10, the accuracies of our FedIIIH obviously increase as K changes from 1, to 2, to 4. Consequently, in Fig. 39, we present the similarity heatmaps of FedIIIH on the *Roman-empire* dataset when K is set to 1, 2, and 4, respectively. From Fig. 39, it can be observed that there are indeed differences between different similarity heatmaps. Although these differences may seem trivial, they have a very large impact on the separate federation and thus on the final performance.

K.4 Disentangled Latent Factors

Here we provide some insights into the interpretability of the disentangled latent factors. Disentangled latent factors, as widely explored in [1, 23, 24], are used to explore the reasons why a node is connected to others. In other words, the interpretability of the disentangled latent factors can be considered as the relationship from a given node to one of its neighbors. For example, a user in a social graph is connected to others for various different reasons, such as families, hobbies, studies, and work. Each disentangled latent factor is capable of capturing mutually exclusive information. For example, the correlation plot of DisenGCN on the eight-factor synthetic graph dataset (Figure 3 in [1]) clearly shows eight diagonal blocks, verifying that the latent factors have indeed been successfully disentangled.

One might ask: Is it possible that the parameter positions corresponding to the same disentangled latent factor differ between clients? We would like to clarify that the parameter positions corresponding to the same latent factor usually remain the same in the disentangled GNNs of different clients. Taking the DisenGCN as an example, there are two crucial processes: node feature projection (Eq. (1) in the main manuscript) and neighborhood routing mechanism (Eq. (1) and Eq. (2) in the Appendix A). According to Eq. (1) in the main manuscript, we can find that the position of parameters corresponding to the k -th latent factor (*i.e.*, \mathbf{W}^k) is determined and then fixed by the node feature \mathbf{x}^i . Since the distributions of \mathbf{x}^i are similar in different clients, the parameter positions corresponding to K latent factors on each client are determined and then fixed in the same way. Moreover, the parameter




Figure 6: Sensitivity of the number of latent factors K on the *PubMed* dataset.




Figure 7: Sensitivity of the number of latent factors K on the *Amazon-Computer* dataset.

positions corresponding to K latent factors are not changed by the neighborhood routing mechanism, because there are no learnable parameters in the neighborhood routing mechanism. Therefore, the parameter positions corresponding to the same latent factor usually remain the same. This can be verified by the experiments, where the standard deviations of our FedIIH are quite small on different datasets (see Tables 1, 2, 3, and 4 in the main manuscript).

K.5 Disentangled Graph Neural Networks

The current implementation of our proposed FedIIH is based on the existing method DisenGCN [1], which is used to instantiate the inference network in our HVGAE. Instantiation in variational inference is very common in many existing approaches [25, 26, 27]. For example, [25] and [26] use the GCN and Graph Attention neTworks (GAT) to instantiate their inference networks, respectively.

Although many disentangled graph neural networks [1, 23, 24] can be chosen flexibly, we directly choose a simple but popular model (*i.e.*, DisenGCN). This is because DisenGCN is a pioneering work in the field. Moreover, our proposed FedIIH is flexible since other disentangled graph neural networks [23, 24] can easily be used.

K.6 Differences Between FED-PUB and Our FedIIH

First, in FED-PUB [10], each client simply feeds the randomly generated graph to the local model and sends the output to the server. In stark contrast, our FedIIH differs noticeably in that it uses HVGAE to disentangle the subgraph into multiple latent factors and accurately infer the distribution of the subgraph data. Then, each client sends this inferred data distribution to the server.

Second, in FED-PUB, the server only measures the similarities of the clients by computing the cosine similarities of the outputs of local models. Conversely, in our FedIIH, the server specifically measures the similarities by computing the JS divergences of the inferred subgraph data distributions, therefore providing a more accurate and stable measure of client similarities.

In summary, our FedIIH differs noticeably from FED-PUB, as FED-PUB mainly focuses on computing the cosine similarities based on the outputs of local models. Unlike FED-PUB, we measure the similarities of clients by computing the JS divergences of the inferred subgraph data distribution.




Figure 8: Sensitivity of the number of latent factors K on the *Amazon-Photo* dataset.




Figure 9: Sensitivity of the number of latent factors K on the *ogbn-arxiv* dataset.

K.7 Reasons for Setting τ to 10

On one hand, τ is recommended by FED-PUB [10] to be set to 10. On the other hand, according to the experimental results of grid search, the performance can achieve the best result when τ is around 10. Tab. 7 provides the accuracies when using different values of τ .

K.8 Why do local model outputs not accurately reflect the distribution of subgraph?

Most existing methods [10, 17, 28] compute the inter-subgraph similarities based on the simplex outputs of local models. However, we argue that the outputs of local models cannot accurately reveal the whole distribution of subgraph data. Here we provide the reason.

According to the universal approximation theorem of neural networks [29], even if two neural networks (*e.g.*, two local models on different clients) have different inputs, they can still produce the completely same output. Therefore, the outputs of local models cannot accurately reveal the whole distribution of subgraph data. This can be verified by experiments. In Fig. 3 of the main manuscript, the similarity heatmap of FED-PUB (Fig. 3b) is quite different from the ground truth (Fig. 3a), which means that the calculated similarities based on the local model outputs cannot accurately reflect the overall distribution of subgraph data.

L Efficiency Analysis

In this section, we present the spatial and temporal complexity of our proposed FedIIIH and that of two baseline methods (*i.e.*, FedAvg and FED-PUB) on the client and server sides, respectively. Furthermore, we also compare the training time of our proposed FedIIIH and that of two baseline methods (*i.e.*, FedAvg and FED-PUB).

L.1 Spatial Complexity

First, the spatial complexity of our FedIIIH on each client side and that of two baseline methods (*i.e.*, FedAvg and FED-PUB) are both $\mathcal{O}(n_m \times d + L \times d^2)$, where n_m , d , and L denote the node




Figure 10: Sensitivity of the number of latent factors K on the *Roman-empire* dataset.




Figure 11: Sensitivity of the number of latent factors K on the *Amazon-ratings* dataset.

number, feature dimensionality, and layer number, respectively. Therefore, the spatial complexity of our FedIIIH on each client is the same as that of two baseline methods.

Second, the spatial complexity of our method on the server side is $\mathcal{O}(M \times d \times (L \times d^2 + K \times M))$, where M and K denote the number of clients and the number of disentangled latent factors, respectively. The spatial complexity of the baseline method (*i.e.*, FED-PUB) on each client is $\mathcal{O}(M \times d \times (L \times d^2 + M))$. By comparing the spatial complexity of our FedIIIH with that of the baseline method (*i.e.*, FED-PUB), we find that the only difference is that the third factor in FedIIIH's spatial complexity is $L \times d^2 + K \times M$ while the third factor in FED-PUB's spatial complexity is $L \times d^2 + M$. Since $K \leq 10$ (as mentioned in the Appendix I.5), we can find that the spatial complexity of our FedIIIH on the server side is similar to that of the baseline method.

L.2 Temporal Complexity

First, the temporal complexity of HVGAE on each client is $\mathcal{O}(n_m \times d \times (L + n_m + d))$. Since the model deployed on each client is actually a HVGAE, the total temporal complexity of our FedIIIH on each client is $\mathcal{O}(n_m \times d \times (L + n_m + d))$. The temporal complexity of two baseline methods (*i.e.*, FedAvg and FED-PUB) on each client is $\mathcal{O}(n_m \times d \times (L + L \times d + d))$. By comparing the temporal complexity of our FedIIIH with that of two baseline methods (*i.e.*, FedAvg and FED-PUB), the only difference is that the third factor in FedIIIH's temporal complexity is $L + n_m + d$ while the third factor in their temporal complexity is $L + L \times d + d$. Since $n_m \leq L \times d$ in most situations (can be verified in the Appendix I.5), we can find that the temporal complexity of our FedIIIH on the client side is similar to that of two baseline methods (*i.e.*, FedAvg and FED-PUB). Therefore, we can find that the introduction of HVGAE does not increase too much computational overhead of our FedIIIH.

Second, the temporal complexity of the divergence computation on the server is $\mathcal{O}(K \times M^2 \times d)$, and the total temporal complexity of our FedIIIH on the server side is $\mathcal{O}(M \times d \times (K \times M + L \times d))$. The temporal complexity of the baseline method (*i.e.*, FED-PUB) on the server side is $\mathcal{O}(M \times d \times (M + L \times d))$. By comparing the temporal complexity of our FedIIIH with that of the baseline method (*i.e.*, FED-PUB), we find that the only difference is that the third factor in FedIIIH's temporal complexity is $K \times M + L \times d$ while the third factor in FED-PUB's temporal complexity is $M + L \times d$. Since




Figure 12: Sensitivity of the number of latent factors K on the *Minesweeper* dataset.




Figure 13: Sensitivity of the number of latent factors K on the *Questions* dataset.

$K \leq 10$ (as mentioned in the Appendix I.5), we can find that the temporal complexity of our FedIIIH on the server side is similar to that of the baseline method (*i.e.*, FED-PUB).

L.3 Training Time

We report the training time of one communication round of our FedIIIH and two baseline methods (*i.e.*, FedAvg and FED-PUB). The experimental platform is shown in the Appendix I.1. We conduct experiments on a homophilic graph dataset (*i.e.*, *Cora*) and a heterophilic graph dataset (*i.e.*, *Roman-empire*). As shown in Tab. 8, we can find that our FedIIIH is more efficient in practice than the baseline methods (*i.e.*, FedAvg and FED-PUB).

M Robustness Analysis

In this section, we analyze the robustness of our FedIIIH in terms of client sparsity and hyperparameter sensitivity, respectively.

M.1 Client Sparsity Analysis

Robustness to client sparsity is a critical aspect in evaluating the effectiveness of a similarity-based personalized federated optimization strategy. Here we conduct experiments on a homophilic graph dataset (*i.e.*, *Cora*) and a heterophilic graph dataset (*i.e.*, *Roman-empire*) as the percentage of participating clients increases from 10% to 100%. As shown in Fig. 40, the experimental results clearly show that the performance of our FedIIIH is more robust and stable than baseline methods, and is not too affected when the number of participating clients decreases.

M.2 Hyperparameter Sensitivity Analysis

There are four vital hyperparameters (*i.e.*, number of latent factors, number of neighborhood routing layers, number of neighborhood routing iterations, and τ) in our proposed FedIIIH. Here we perform experiments to analyze the hyperparameter sensitivity.




Figure 14: The similarity heatmaps of FED-PUB over three independent runs on the *Cora* dataset in the overlapping setting with 20 clients.

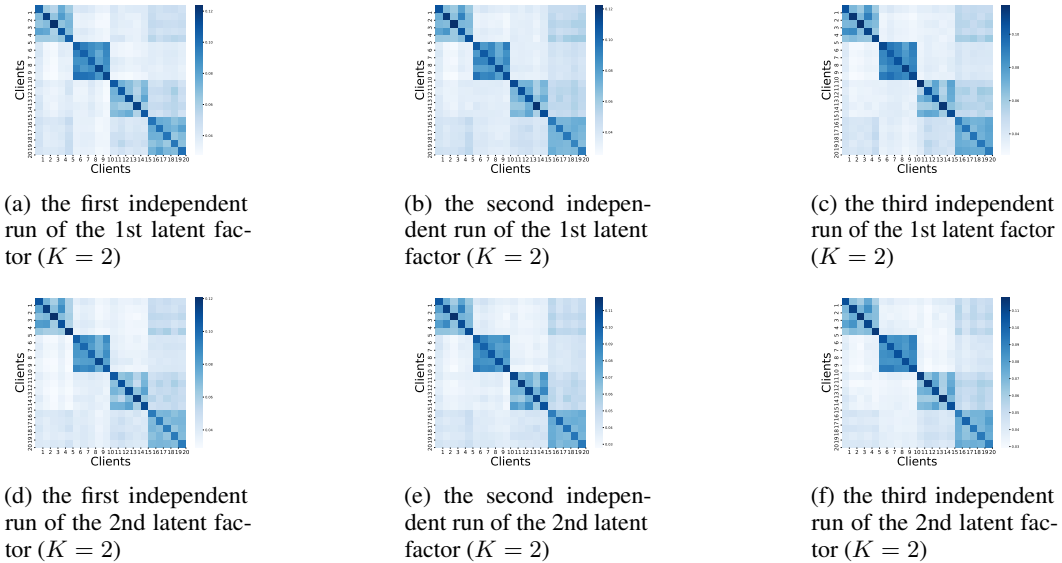


Figure 15: The similarity heatmaps of FedIIH over three independent runs on the *Cora* dataset in the overlapping setting with 20 clients.

First, as shown in the Appendix J.2, the variation in performance under different numbers of latent factors (*i.e.*, K) is small. Second, as shown in Tab. 9, Tab. 10, and Tab. 11, the variations in performance under different numbers of neighborhood routing layers, different numbers of neighborhood routing iterations, and different values of τ are all small. These experimental results clearly demonstrate that the performances of our proposed FedIIH are very stable within a given range of hyperparameters, therefore the hyperparameter of our FedIIH can be easily tuned in practical use.

N Broader Impact

Our work could have the following positive impacts: (1) We provide a new method for GFL to deal with the inter-intra heterogeneity. (2) Our proposed method can greatly improve the performance of GFL on homophilic and heterophilic graph datasets in both non-overlapping and overlapping settings.

The proposed method can be used for both good and bad, similar to many other FL methods. Note that most existing methods and our proposal are not immune to such misuse. We believe that such a problem can be solved in the future, although we do not have an optimal solution.

In summary, we believe that our proposed method can benefit society because many important real-world graphs face inter-intra heterogeneity when performing the GFL. Therefore, they can benefit from our proposed FedIIH.




Figure 16: The similarity heatmaps of FED-PUB over three independent runs on the *Amazon-ratings* dataset in the overlapping setting with 20 clients.

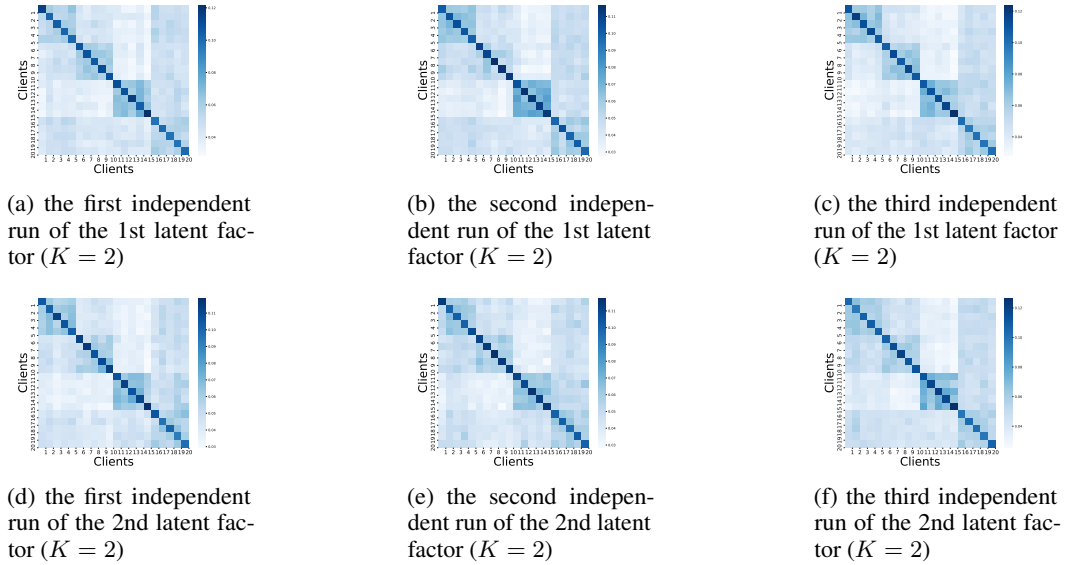


Figure 17: The similarity heatmaps of FedIIH over three independent runs on the *Amazon-ratings* dataset in the overlapping setting with 20 clients.

References

- [1] Jianxin Ma, Peng Cui, Kun Kuang, Xin Wang, and Wenwu Zhu. Disentangled graph convolutional networks. In *International Conference on Machine Learning*, pages 4212–4221, 2019.
- [2] Wei-Ning Hsu, Yu Zhang, and James Glass. Unsupervised learning of disentangled and interpretable representations from sequential data. In *Advances in Neural Information Processing Systems*, pages 1–12, 2017.
- [3] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv:1312.6114*, 2013.
- [4] Maruan Al-Shedivat, Jennifer Gillenwater, Eric Xing, and Afshin Rostamizadeh. Federated learning via posterior averaging: A new perspective and practical algorithms. In *International Conference on Learning Representations*, pages 1–23, 2021.
- [5] Xu Zhang, Yinchuan Li, Wenpeng Li, Kaiyang Guo, and Yunfeng Shao. Personalized federated learning via variational Bayesian inference. In *International Conference on Machine Learning*, pages 26293–26310, 2022.
- [6] Minyoung Kim and Timothy Hospedales. FedHB: Hierarchical Bayesian federated learning. *arXiv:2305.04979*, 2023.




Figure 18: Similarity heatmaps on the *Cora* dataset in the overlapping setting with 20 clients.




Figure 19: Similarity heatmaps on the *CiteSeer* dataset in the non-overlapping setting with 20 clients.

- [7] Liangxi Liu, Xi Jiang, Feng Zheng, Hong Chen, Guo-Jun Qi, Heng Huang, and Ling Shao. A Bayesian federated learning framework with online Laplace approximation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–16, 2023.
- [8] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *International Conference on Artificial Intelligence and Statistics*, pages 1273–1282, 2017.
- [9] Oleg Platonov, Denis Kuznedelev, Michael Diskin, Artem Babenko, and Liudmila Prokhorenkova. A critical look at the evaluation of GNNs under heterophily: Are we really making progress? In *International Conference on Learning Representations*, pages 1–15, 2023.
- [10] Jinheon Baek, Wonyong Jeong, Jiongdoo Jin, Jaehong Yoon, and Sung Ju Hwang. Personalized subgraph federated learning. In *International Conference on Machine Learning*, pages 1396–1415, 2023.
- [11] George Karypis. METIS: Unstructured graph partitioning and sparse matrix ordering system. *Technical report*, 1997.
- [12] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. In *Proceedings of Machine Learning and Systems*, pages 429–450, 2020.
- [13] Manoj Ghuhun Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Federated learning with personalization layers. *arXiv:1912.00818*, 2019.
- [14] Han Xie, Jing Ma, Li Xiong, and Carl Yang. Federated graph classification over non-iid graphs. In *Advances in Neural Information Processing Systems*, pages 18839–18852, 2021.
- [15] Chuhan Wu, Fangzhao Wu, Yang Cao, Yongfeng Huang, and Xing Xie. FedGNN: Federated graph neural network for privacy-preserving recommendation. *arXiv:2102.04925*, 2021.
- [16] Ke Zhang, Carl Yang, Xiaoxiao Li, Lichao Sun, and Siu Ming Yiu. Subgraph federated learning with missing neighbor generation. In *Advances in Neural Information Processing Systems*, pages 6671–6682, 2021.
- [17] Xunkai Li, Zhengyu Wu, Wentao Zhang, Yinlin Zhu, Rong-Hua Li, and Guoren Wang. FedGTA: Topology-aware averaging for federated graph learning. In *Proceedings of the VLDB Endowment*, pages 41–50, 2023.




Figure 20: Similarity heatmaps on the *CiteSeer* dataset in the overlapping setting with 30 clients.




Figure 21: Similarity heatmaps on the *PubMed* dataset in the non-overlapping setting with 20 clients.

- [18] Xunkai Li, Zhengyu Wu, Wentao Zhang, Henan Sun, Rong-Hua Li, and Guoren Wang. AdaFGL: A new paradigm for federated node classification with topology heterogeneity. In *International Conference on Data Engineering*, pages 2517–2530, 2024.
- [19] Felix Sattler, Klaus-Robert Müller, and Wojciech Samek. Clustered federated learning: Model-agnostic distributed multitask optimization under privacy constraints. *IEEE Transactions on Neural Networks and Learning Systems*, 32(8):3710–3722, 2021.
- [20] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, pages 1–14, 2017.
- [21] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, pages 1–11, 2017.
- [22] Wentao Zhang, Ziqi Yin, Zeang Sheng, Yang Li, Wen Ouyang, Xiaosen Li, Yangyu Tao, Zhi Yang, and Bin Cui. Graph attention multi-layer perceptron. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4560–4570, 2022.
- [23] Haoyang Li, Xin Wang, Ziwei Zhang, Zehuan Yuan, Hang Li, and Wenwu Zhu. Disentangled contrastive learning on graphs. In *Advances in Neural Information Processing Systems*, pages 21872–21884, 2021.
- [24] Jingwei Guo, Kaizhu Huang, Xinping Yi, and Rui Zhang. Learning disentangled graph convolutional networks locally and globally. *IEEE Transactions on Neural Networks and Learning Systems*, 35(3):3640–3651, 2024.
- [25] Jiaqi Ma, Weijing Tang, Ji Zhu, and Qiaozhu Mei. A flexible generative framework for graph-based semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 1–10, 2019.
- [26] Haibo Wang, Chuan Zhou, Xin Chen, Jia Wu, Shirui Pan, and Jilong Wang. Graph stochastic neural networks for semi-supervised learning. In *Advances in Neural Information Processing Systems*, pages 1–10, 2020.
- [27] Sheng Wan, Yibing Zhan, Liu Liu, Baosheng Yu, Shirui Pan, and Chen Gong. Contrastive graph poisson networks: Semi-supervised learning with extremely limited labels. In *Advances in Neural Information Processing Systems*, pages 1–12, 2021.
- [28] Zaixi Zhang, Qingyong Hu, Yang Yu, Weibo Gao, and Qi Liu. FedGT: Federated node classification with scalable graph transformer. *arXiv:2401.15203*, 2024.




Figure 22: Similarity heatmaps on the *PubMed* dataset in the overlapping setting with 30 clients.




Figure 23: Similarity heatmaps on the *Amazon-Computer* dataset in the non-overlapping setting with 20 clients.

[29] Rickard Brül Gabrielsson. Universal function approximation on graphs. In *Advances in Neural Information Processing Systems*, pages 1–11, 2020.

Reproducibility Checklist

1. This paper includes a conceptual outline and/or pseudocode description of AI methods introduced. (yes/partial/no/NA)

Answer: [Yes]

Justification: Our paper includes a conceptual outline and pseudocode description of methods in this Appendix.

2. This paper clearly delineates statements that are opinions, hypothesis, and speculation from objective facts and results. (yes/no)

Answer: [Yes]

Justification: First, all the formulas and proofs in the paper are numbered and cross-referenced. Second, due to the space limitation, the detailed proof of ELBO is presented in the Appendix C. Third, we provide the extensive experimental results.

3. This paper provides well marked pedagogical references for less-familiare readers to gain background necessary to replicate the paper. (yes/no)

Answer: [Yes]

Justification: We provide the detailed preliminaries.

4. Does this paper make theoretical contributions? (yes/no)

Answer: [Yes]

Justification: This paper make theoretical contributions related to the ELBO.

5. All assumptions and restrictions are stated clearly and formally. (yes/partial/no)

Answer: [Yes]

Justification: This paper clearly and formally states all assumptions and restrictions.

6. All novel claims are stated formally (e.g., in theorem statements). (yes/partial/no)

Answer: [Yes]

Justification: This paper clearly and formally states all all novel claims.

7. Proofs of all novel claims are included. (yes/partial/no)




Figure 24: Similarity heatmaps on the *Amazon-Computer* dataset in the overlapping setting with 30 clients.




Figure 25: Similarity heatmaps on the *Amazon-Photo* dataset in the non-overlapping setting with 20 clients.

Answer: [Yes]

Justification: We provide the proofs in this Appendix.

8. Proof sketches or intuitions are given for complex and/or novel results. (yes/partial/no)

Answer: [Yes]

Justification: We provide the proofs in this Appendix.

9. Appropriate citations to theoretical tools used are given. (yes/partial/no)

Answer: [Yes]

Justification: We provide all the citations.

10. All theoretical claims are demonstrated empirically to hold. (yes/partial/no/NA)

Answer: [Yes]

Justification: The experimental results verify all theoretical claims.

11. All experimental code used to eliminate or disprove claims is included. (yes/no/NA)

Answer: [Yes]

Justification: Our code is available at <https://github.com/blgpb/FedIIH>. We also provide a README.md file to describe the detailed instructions on how to replicate the results.

12. Does this paper rely on one or more datasets? (yes/no)

Answer: [Yes]

Justification: We use several datasets.

13. A motivation is given for why the experiments are conducted on the selected datasets (yes/partial/no/NA)

Answer: [Yes]

Justification: We explain it in the paper.

14. All novel datasets introduced in this paper are included in a data appendix. (yes/partial/no/NA)

Answer: [Yes]

Justification: We provide them in this Appendix.




Figure 26: Similarity heatmaps on the *Amazon-Photo* dataset in the overlapping setting with 30 clients.




Figure 27: Similarity heatmaps on the *ogbn-arxiv* dataset in the non-overlapping setting with 20 clients.

15. All novel datasets introduced in this paper will be made publicly available upon publication of the paper with a license that allows free usage for research purposes. (yes/partial/no/NA)

Answer: [Yes]

Justification: We provide them in this Appendix.

16. All datasets drawn from the existing literature (potentially including authors' own previously published work) are accompanied by appropriate citations. (yes/no/NA)

Answer: [Yes]

Justification: We cite all the papers that produced the datasets.

17. All datasets drawn from the existing literature (potentially including authors' own previously published work) are publicly available. (yes/partial/no/NA)

Answer: [Yes]

Justification: All datasets are publicly available.

18. Does this paper include computational experiments? (yes/no)

Answer: [Yes]

Justification: We provide them in the paper.

19. Any code required for pre-processing data is included in the appendix. (yes/partial/no).

Answer: [Yes]

Justification: Our code is available at <https://github.com/blgpb/FedIIH>.

20. All source code required for conducting and analyzing the experiments is included in a code appendix. (yes/partial/no)

Answer: [Yes]

Justification: Our code is available at <https://github.com/blgpb/FedIIH>.

21. All source code required for conducting and analyzing the experiments will be made publicly available upon publication of the paper with a license that allows free usage for research purposes. (yes/partial/no)

Answer: [Yes]

Justification: All source code will be made publicly available upon publication of the paper with a license that allows free usage for research purposes.




Figure 28: Similarity heatmaps on the *ogbn-arxiv* dataset in the overlapping setting with 30 clients.




Figure 29: Similarity heatmaps on the *Roman-empire* dataset in the non-overlapping setting with 20 clients.

22. All source code implementing new methods have comments detailing the implementation, with references to the paper where each step comes from (yes/partial/no)

Answer: [Yes]

Justification: We provide them in this Appendix.

23. If an algorithm depends on randomness, then the method used for setting seeds is described in a way sufficient to allow replication of results. (yes/partial/no/NA)

Answer: [Yes]

Justification: Our code is available at <https://github.com/blgpb/FedIIH>. We also provide a README.md file to describe the detailed instructions on how to replicate the results.

24. This paper specifies the computing infrastructure used for running experiments (hardware and software), including GPU/CPU models; amount of memory; operating system; names and versions of relevant software libraries and frameworks. (yes/partial/no)

Answer: [Yes]

Justification: We provide them in this Appendix.

25. This paper formally describes evaluation metrics used and explains the motivation for choosing these metrics. (yes/partial/no)

Answer: [Yes]

Justification: We provide them in this Appendix.

26. This paper states the number of algorithm runs used to compute each reported result. (yes/no)

Answer: [Yes]

Justification: We provide them in the paper.

27. Analysis of experiments goes beyond single-dimensional summaries of performance (e.g., average; median) to include measures of variation, confidence, or other distributional information. (yes/no)

Answer: [Yes]

Justification: We provide them in this Appendix.

28. The significance of any improvement or decrease in performance is judged using appropriate statistical tests (e.g., Wilcoxon signed-rank). (yes/partial/no)




Figure 30: Similarity heatmaps on the *Roman-empire* dataset in the overlapping setting with 30 clients.




Figure 31: Similarity heatmaps on the *Amazon-ratings* dataset in the non-overlapping setting with 20 clients.

Answer: [\[Yes\]](#)

Justification: We provide them in the paper.

29. This paper lists all final (hyper-)parameters used for each model/algorithm in the paper's experiments. (yes/partial/no/NA)

Answer: [\[Yes\]](#)

Justification: We provide them in this Appendix.

30. This paper states the number and range of values tried per (hyper-) parameter during development of the paper, along with the criterion used for selecting the final parameter setting. (yes/partial/no/NA)

Answer: [\[Yes\]](#)

Justification: We provide them in this Appendix.




Figure 32: Similarity heatmaps on the *Amazon-ratings* dataset in the overlapping setting with 30 clients.




Figure 33: Similarity heatmaps on the *Minesweeper* dataset in the non-overlapping setting with 20 clients.




Figure 34: Similarity heatmaps on the *Minesweeper* dataset in the overlapping setting with 30 clients.




Figure 35: Similarity heatmaps on the *Tolokers* dataset in the non-overlapping setting with 20 clients.




Figure 36: Similarity heatmaps on the *Tolokers* dataset in the overlapping setting with 30 clients.




Figure 37: Similarity heatmaps on the *Questions* dataset in the non-overlapping setting with 20 clients.



Figure 38: Similarity heatmaps on the *Questions* dataset in the overlapping setting with 30 clients.

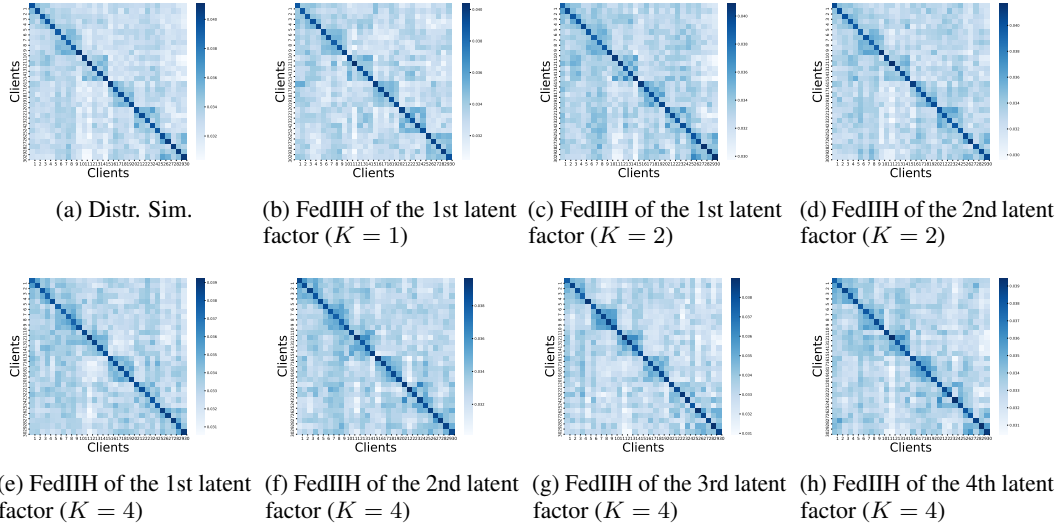


Figure 39: Similarity heatmaps on the *Roman-empire* dataset in the overlapping setting with 30 clients. There are indeed differences between different similarity heatmaps when K is set to 1, 2, and 4.

Table 7: The hyperparameter sensitivity analysis of τ . The best results are shown in **bold**.

τ	<i>Cora</i> Non-overlapping 10 Clients	<i>Cora</i> Overlapping 30 Clients	<i>Roman-empire</i> Non-overlapping 10 Clients	<i>Roman-empire</i> Overlapping 30 Clients
1	81.58±0.15	76.74±0.34	66.12±0.25	63.15±0.15
2	81.61±0.12	76.52±0.16	66.41±0.34	63.23±0.12
3	81.76±0.16	76.69±0.20	66.34±0.32	63.19±0.25
4	81.80±0.17	76.72±0.23	66.40±0.36	63.29±0.30
5	81.65±0.05	76.66±0.31	66.36±0.27	63.26±0.31
6	81.82±0.10	76.75±0.26	66.21±0.19	63.28±0.24
7	81.75±0.11	76.80±0.15	66.37±0.33	63.13±0.15
8	81.81±0.18	76.79±0.25	66.36±0.40	63.22±0.22
9	81.82±0.14	76.74±0.26	66.34±0.21	63.30±0.15
10	81.85±0.09	76.82±0.24	66.44±0.28	63.32±0.06

Table 8: The training time of one communication round of our FedIH and two baseline methods.

Methods	<i>Cora</i> Nonoverlapping			<i>Cora</i> Overlapping		
	5 Clients	10 Clients	20 Clients	10 Clients	30 Clients	50 Clients
FedAvg	20.81s	24.90s	59.58s	30.63s	72.55s	141.82s
FED-PUB	22.04s	27.34s	60.31s	33.46s	80.54s	147.84s
FedIH (Ours)	19.57s	22.76s	56.09s	19.67s	65.49s	139.03s
Methods	<i>Roman-empire</i> Nonoverlapping			<i>Roman-empire</i> Overlapping		
	5 Clients	10 Clients	20 Clients	10 Clients	30 Clients	50 Clients
FedAvg	18.25s	29.85s	55.21s	28.30s	79.12s	127.61s
FED-PUB	18.81s	28.03s	61.75s	28.45s	83.07s	133.05s
FedIH (Ours)	17.45s	24.90s	47.01s	28.19s	61.17s	100.10s

Table 9: The hyperparameter sensitivity analysis of the number of neighborhood routing layers.


# neighborhood routing layers	<i>Cora</i> Non-overlapping	<i>Cora</i> Overlapping	<i>Roman-empire</i> Non-overlapping	<i>Roman-empire</i> Overlapping
	10 Clients	30 Clients	10 Clients	30 Clients
1	81.19±0.27	76.18±0.46	66.44±0.28	63.32±0.06
2	81.41±0.15	76.34±0.31	66.16±0.41	63.15±0.15
3	81.46±0.12	76.48±0.15	66.23±0.36	63.12±0.21
4	81.85±0.09	76.67±0.18	66.15±0.46	63.15±0.42
5	81.62±0.08	76.82±0.24	66.11±0.55	63.10±0.24
6	81.37±0.24	76.46±0.45	66.08±0.39	63.04±0.38
max - min	0.66	0.64	0.36	0.28

Table 10: The hyperparameter sensitivity analysis of the number of neighborhood routing iterations.


# neighborhood routing iterations	<i>Cora</i> Non-overlapping	<i>Cora</i> Overlapping	<i>Roman-empire</i> Non-overlapping	<i>Roman-empire</i> Overlapping
	10 Clients	30 Clients	10 Clients	30 Clients
2	81.24±0.50	76.22±0.30	66.04±0.43	63.01±0.45
3	81.38±0.45	76.31±0.44	66.10±0.46	63.10±0.34
4	81.45±0.41	76.37±0.45	66.16±0.44	63.18±0.29
5	81.57±0.35	76.45±0.41	66.15±0.35	63.24±0.17
6	81.85±0.09	76.82±0.24	66.44±0.28	63.32±0.06
7	81.36±0.42	76.68±0.39	66.24±0.29	63.12±0.22
max - min	0.61	0.60	0.40	0.31

Table 11: The hyperparameter sensitivity analysis of τ .


τ	<i>Cora</i> Non-overlapping	<i>Cora</i> Overlapping	<i>Roman-empire</i> Non-overlapping	<i>Roman-empire</i> Overlapping
	10 Clients	30 Clients	10 Clients	30 Clients
1	81.58±0.15	76.74±0.34	66.12±0.25	63.15±0.15
2	81.61±0.12	76.52±0.16	66.41±0.34	63.23±0.12
3	81.76±0.16	76.69±0.20	66.34±0.32	63.19±0.25
4	81.80±0.17	76.72±0.23	66.40±0.36	63.29±0.30
5	81.65±0.05	76.66±0.31	66.36±0.27	63.26±0.31
6	81.82±0.10	76.75±0.26	66.21±0.19	63.28±0.24
7	81.75±0.11	76.80±0.15	66.37±0.33	63.13±0.15
8	81.81±0.18	76.79±0.25	66.36±0.40	63.22±0.22
9	81.82±0.14	76.74±0.26	66.34±0.21	63.30±0.15
10	81.85±0.09	76.82±0.24	66.44±0.28	63.32±0.06
max - min	0.27	0.30	0.32	0.19




(a) Cora Non-overlapping 20 Clients



(b) Cora Overlapping 50 Clients



(c) Roman-empire Non-overlapping 20 Clients



(d) Roman-empire Overlapping 50 Clients

Figure 40: Performances with different percentages of participating clients.