

Vertically Federated Graph Neural Network for Privacy-Preserving Node Classification

Chaochao Chen¹, Jun Zhou^{1,2}, Longfei Zheng², Huiwen Wu², Lingjuan Lyu³,
 Jia Wu⁴, Bingzhe Wu⁵, Ziqi Liu², Li Wang², Xiaolin Zheng^{1,6*}

¹College of Computer Science and Technology, Zhejiang University, Hangzhou, China

²Ant Group, Hangzhou, China

³Sony AI, Tokyo, Japan

⁴Macquarie University, Sydney, NSW 2109, Australia

⁵Peking University, Beijing, China

⁶JZTData Technology, Hangzhou, China

{zjuccc, xlzheng}@zju.edu.cn, Lingjuan.Lv@sony.com, jia.wu@mq.edu.au, wubingzhe@pku.edu.cn,
 {jun.zhoujun, zlf206411, huiwen.whw, ziqiliu, raymond.wang}@antgroup.com

Abstract

Graph Neural Network (GNN) has achieved remarkable progresses in various real-world tasks on graph data. High-performance GNN models always depend on both rich features and complete edge information in graph. However, such information could possibly be isolated by different data holders in practice, which is the so-called data isolation problem. To solve this problem, in this paper, we propose Vertically Federated Graph Neural Network (VFGNN), a federated GNN learning paradigm for privacy-preserving node classification task under data vertically partitioned setting, which can be generalized to existing GNN models. Specifically, we split the computation graph into two parts. We leave the private data (i.e., features, edges, and labels) related computations on data holders, and delegate the rest of computations to a semi-honest server. We also propose to apply differential privacy to prevent potential information leakage from the server. We conduct experiments on three benchmarks and the results demonstrate the effectiveness of VFGNN.

1 Introduction

Graph Neural Network (GNN) has gained increasing attentions from both academy and industry due to its ability to model high-dimensional feature information and high-order adjacent information on both homogeneous and heterogeneous graphs [Wu *et al.*, 2019]. An important ingredient for high-performance GNN models is high-quality graph data including rich node features and complete adjacent information. However, in practice, such information could possibly be isolated by different data holders, which is the so-called *data isolation* problem. Such a data isolation problem

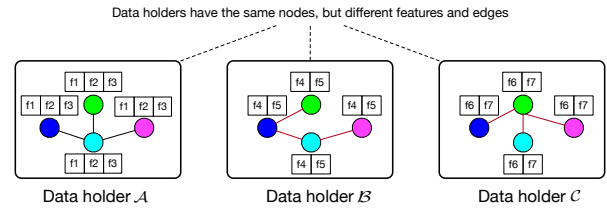


Figure 1: Problem description of vertically federated GNN.

presents a serious challenge for the development of Artificial Intelligence, which becomes a hot research topic recently.

Problem. Figure 1 shows a privacy-preserving node classification problem under vertically partitioned data setting. Here, we assume there are three data holders (\mathcal{A} , \mathcal{B} , and \mathcal{C}) and they have four same nodes. The node features are vertically split, i.e., \mathcal{A} has f_1 , f_2 , and f_3 , \mathcal{B} has f_4 and f_5 , and \mathcal{C} has f_6 and f_7 . Meanwhile, \mathcal{A} , \mathcal{B} , and \mathcal{C} may have their own edges. For example, \mathcal{A} has social relation between nodes while \mathcal{B} and \mathcal{C} have payment relation between nodes. We also assume \mathcal{A} is the party who has the node labels. The problem is to build federated GNN models using the graph data of \mathcal{A} , \mathcal{B} , and \mathcal{C} .

Related work. To date, many kinds of privacy-preserving machine learning models have been proposed, e.g., logistic regression [Chen *et al.*, 2021], decision tree [Fang *et al.*, 2021], and neural network [Wagh *et al.*, 2019]. There are also several work on studying the privacy issues in GNN, e.g., graph publishing [Sajadmanesh and Gatica-Perez, 2020], GNN inference [He *et al.*, 2020], and federated GNN when data are horizontally partitioned [Zheng *et al.*, 2021; Wu *et al.*, 2021]. Unlike previous privacy-preserving machine learning models that assume only samples (nodes) are held by different parties and they have no relationship, our task is more challenging because GNN relies on the relationships between samples, which are also private to data holders.

Naive solution. A direct way to build privacy-preserving GNN is employing advanced cryptographic algorithms, such

*Corresponding author

as homomorphic encryption (HE) and secure multi party computation (MPC) [Mohassel and Zhang, 2017]. Such a pure cryptographic way can provide high security guarantees, however, it suffers high computation and communication costs, which limits their efficiency [Osia *et al.*, 2019].

Our solution. We propose VFGNN, a vertically federated GNN learning paradigm for privacy-preserving node classification task. Motivated by the existing work in split learning [Vepakomma *et al.*, 2018; Osia *et al.*, 2019; Gu *et al.*, 2018], we split the computation graph of GNN into two parts for privacy and efficiency concern, i.e., the private data related computations carried out by data holders and non-private data related computations conducted by a semi-honest server. Specifically, data holders first apply MPC techniques to compute the initial layer of the GNN using private node feature information collaboratively, which acts as the feature extractor module, and then perform neighborhood aggregation using private edge information individually, similar as the existing GNNs [Veličković *et al.*, 2017], and finally get the local node embeddings. Next, we propose different combination strategies for a semi-honest server to combine local node embeddings from data holders and generate global node embeddings, based on which the server can conduct the successive non-private data related computations, e.g., the non-linear operations in deep network structures that are time-consuming for MPC techniques. Finally, the server returns the final hidden layer to the party who has labels to compute prediction and loss. Data holders and the server perform forward and back propagations to complete model training and prediction, during which the private data (i.e., features, edges, and labels) are always kept by data holders themselves. Moreover, we adopt differential privacy, on the exchanged information between server and data holders, to further protect potential information leakage from the server.

Contributions. We summarize our main contributions as:

- We propose a novel VFGNN learning paradigm, which not only can be generalized to most existing GNNs, but also enjoys good accuracy and efficiency.
- We propose different combination strategies for the server to combine local node embeddings from data holders.
- We evaluate our proposals on three real-world datasets, and the results demonstrate the effectiveness of VFGNN.

2 Preliminaries

2.1 Security Model

In this paper, we assume the adversary is honest-but-curious (semi-honest). That is, data holders and the server strictly follow the protocol, but they also use all intermediate computation results to infer as much information as possible. We also assume that the server does not collude with any data holders. This security setting is similar as most existing work [Mohassel and Zhang, 2017; Hardy *et al.*, 2017].

2.2 Additive Secret Sharing

Additive Secret Sharing has two main procedures [Shamir, 1979]. To additively share $\mathbf{Shr}(\cdot)$ an ℓ -bit value a for party $i \in \mathcal{P} = \{1, \dots, I\}$, party i generates $\{a_j \in \mathbb{Z}_{2^\ell}, j \in \mathcal{P}$

$j \neq i\}$ uniformly at random, sends a_j to party j , and keeps $a_i = a - \sum_{j \neq i} a_j \bmod 2^\ell$. We use $\langle a \rangle_i = a_i$ to denote the share of party i . To reconstruct $\mathbf{Rec}(\cdot, \cdot)$ a shared value $\langle a \rangle$, each party i sends $\langle a \rangle_i$ to one who computes $\sum_i a_i \bmod 2^\ell$. For simplification, we denote additive sharing by $\langle \cdot \rangle$. Addition in secret sharing can be done by participants locally. Multiplication in secret sharing usually relies on Beaver’s triplet technique [Beaver, 1991].

2.3 Differential Privacy

Definition 1. (Differential Privacy [Dwork *et al.*, 2014]). A randomized algorithm \mathcal{M} that takes as input a dataset consisting of individuals is (ϵ, δ) -differentially private (DP) if for any pair of neighbouring data x, y that differ in a single entry, and any event E ,

$$P[\mathcal{M}(x) \in E] \leq \exp(\epsilon)P[\mathcal{M}(y) \in E] + \delta, \quad (1)$$

and if $\delta = 0$, we say that \mathcal{M} is ϵ -differentially private.

In [Dwork *et al.*, 2014], the authors pointed out that the ℓ_2 -sensitivity of a function f measures the magnitude by which a single individual’s data can change the function in the worst case.

Definition 2. (ℓ_2 -sensitivity [Dwork *et al.*, 2014]). Suppose x and y are neighbouring inputs that differ in one entry. The ℓ_2 -sensitivity of a function $f : \mathcal{D} \rightarrow \mathbb{R}^d$ is:

$$\Delta_2 f = \max_{x, y \in \mathcal{D}, \|x - y\|_2 = 1} \|f(x) - f(y)\|_2. \quad (2)$$

Definition 3. (The Gaussian Mechanism [Dwork *et al.*, 2014]) Given a function $f : \mathcal{D} \rightarrow \mathbb{R}^d$ over a dataset \mathcal{D} , the Gaussian mechanism is defined as:

$$\mathcal{M}_G(x, f(\cdot), \epsilon) = f(x) + (Y_1, \dots, Y_k), \quad (3)$$

where Y_i are i.i.d. random variables drawn from $\mathcal{N}(0, \sigma^2 \Delta_2 f^2)$ and $\sigma = \frac{\sqrt{2 \ln(1.25/\delta)}}{\epsilon}$.

Theorem 1. [Dwork *et al.*, 2014] The Gaussian mechanism defined in Definition 3 preserves (ϵ, δ) -DP for each publication step.

3 The Proposed Model

3.1 Overview of VFGNN

As described in Section 1, for the sake of privacy and efficiency, we design a vertically federated GNN (VFGNN) learning paradigm by *splitting the computational graph of GNN* into two parts. That is, we keep the *private data related computations* to data holders for privacy concern, and delegate the *non-private data related computations* to a semi-honest server for efficiency concern. In the context of GNN, the private data refers to node features, labels, and edges (node relations). To be specific, we divide the computational graph into the following three sub-Computational Graphs (CG), as is shown in Figure 2.

CG1: private feature and edge related computations. The first step of GNN is generating initial node embeddings using node’s private features, e.g., user features in social networks. In vertically data split setting, each data holder has partial

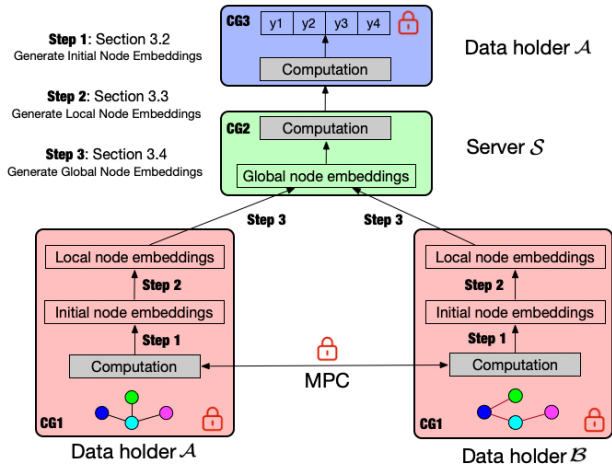


Figure 2: Overview of our proposed VFGNN.

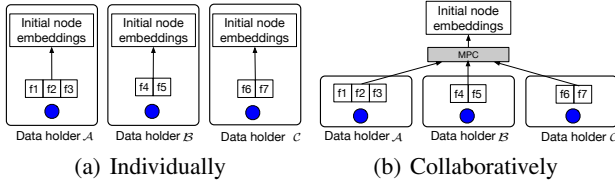


Figure 3: Methods of generating initial node embeddings.

node features, as shown in Figure 1. We will present how data holders collaboratively learn initial node embeddings later. In the next step, data holders generate local node embeddings by aggregating multi-hop neighbors’ information using different aggregator functions.

CG2: non-private data related computations. We delegate non-private data related computations to a semi-honest server for efficiency concern. First, the server combines the local node embeddings from data holders with different COMBINE strategies, and obtains the global node embeddings. Next, the server can perform the successive computations using plaintext data. Note that this part has many non-linear computations such as max-pooling and activation functions, which are not cryptographically friendly and need to be approximated [Mohassel and Zhang, 2017; Hardy *et al.*, 2017]. Delegating these plaintext computations to server will not only improve our model accuracy, but also significantly improve our model efficiency. After this, the server gets a final hidden layer (\mathbf{z}_L) and sends it back to the data holder who has label to calculate the prediction, where L is the total number of layers of the deep neural network.

CG3: private label related computations. The data holder who has label can compute prediction using the final hidden layer it receives from the server. For node classification task, the Softmax activation function is used for the output layer [Kipf and Welling, 2016], which is defined as $\text{softmax}(z_c) = \frac{1}{Z} \exp(z_c)$ with $c \in C$ be the node class and $Z = \sum_c \exp(z_c)$.

In the following subsections, we will describe three important components of VFGNN.

3.2 Generate Initial Node Embeddings

Initial node embeddings are generated by using node features. Under vertically partitioned data setting, each data holder has partial node features. There are two methods for data holders to generate initial node embeddings, i.e., *individually* and *collaboratively*, as shown in Figure 3.

The ‘*individually*’ method means that data holders generate initial node embeddings using their own node features, individually. For data holder $i \in \mathcal{P}$, this can be done by $\mathbf{h}_0^i = (\mathbf{x}^i)^T \cdot \mathbf{W}^i$, where \mathbf{x}^i and \mathbf{W}^i are node features and weight matrix of data holder i . As the example in Figure 3 (a), \mathcal{A} , \mathcal{B} , and \mathcal{C} generate their initial node embeddings using their own features separately. Although this method is simple and data holders do not need to communicate with each other, it cannot capture the relationship between features of data holders and thus causes information loss.

To solve the shortcoming of ‘*individually*’ method, we propose a ‘*collaboratively*’ method. It indicates that data holders generate initial node embeddings using their node features, collaboratively, and meanwhile keep their private features secure. Technically, this can be done by using cryptographic methods such as secret sharing and homomorphic encryption [Acar *et al.*, 2018]. In this paper, we choose additive secret sharing due to its high efficiency.

3.3 Generate Local Node Embeddings

We generate local node embeddings by using multi-hop neighborhood aggregation on graphs, based on initial node embeddings. Note that, *neighborhood aggregation should be done by data holders separately*, rather than cooperatively, to protect the private edge information. This is because one may infer the neighborhood information of v given the neighborhood aggregation results of k -hop ($\mathbf{h}_v^k(i)$) and $k+1$ -hop ($\mathbf{h}_v^{k+1}(i)$), if neighborhood aggregation is done by data holders jointly. For $\forall v \in \mathcal{V}$ at each data holder, neighborhood aggregation is the same as the traditional GNN. Take GraphSAGE [Hamilton *et al.*, 2017] for example, it introduces aggregator functions to update hidden embeddings by sampling and aggregating features from a node’s local neighborhood:

$$\begin{aligned} \mathbf{h}_{\mathcal{N}(v)}^k &\leftarrow \text{AGG}_k(\{\mathbf{h}_u^{k-1}, \forall u \in \mathcal{N}(v)\}), \\ \mathbf{h}_v^k &\leftarrow (\mathbf{W}^k \cdot \text{CONCAT}(\mathbf{h}_v^{k-1}, \mathbf{h}_{\mathcal{N}(v)}^k)), \end{aligned} \quad (4)$$

where we follow the same notations as GraphSAGE, and the aggregator functions AGG are of three types, i.e., Mean, LSTM, and Pooling. After it, data holders send local node embeddings to a semi-honest server for combination and further non-private data related computations.

3.4 Generate Global Node Embeddings

The server combines the local node embeddings from data holders and gets global node embeddings. The combination strategy (COMBINE) should be trainable and maintaining high representational capacity, and we design three of them.

Concat. The concat operator can fully preserve local node embeddings learnt from different data holders. That is, Line 14 in Algorithm 2 becomes

$$\mathbf{h}_v^K \leftarrow \text{CONCAT}(\mathbf{h}_v^K(1), \mathbf{h}_v^K(2), \dots, \mathbf{h}_v^K(I)). \quad (5)$$

Algorithm 1 Information publishing mechanisms of data holders to server using differential privacy

Input: Local information of data holders \mathbf{x} , dimension of local information d , noise multiplier σ , clipping value C .

Output: Differentially private node embeddings.

- 1: Scale local information $\bar{\mathbf{x}} = \min(1, C/\|\mathbf{x}\|)\mathbf{x}$;
- 2: Draw i.i.d. samples from $\mathcal{N}(0, \sigma^2 C^2)$, which forms a d -dimension noise vector \mathbf{n} ;
- 3: # **Gaussian Mechanism**
- 4: Add noise $\tilde{\mathbf{x}} = \mathbf{x}^K + \mathbf{n}$;
- 5: # **James-Stein Estimator**
- 6: Compute James-Stein Estimator

$$\tilde{\mathbf{x}}_{\text{JS}} = \left(1 - \frac{(d-2)\sigma^2 C^2}{\|\tilde{\mathbf{x}}\|^2}\right) \tilde{\mathbf{x}}$$
- 7: **return** $\tilde{\mathbf{x}}$ or $\tilde{\mathbf{x}}_{\text{JS}}$.

Mean. The mean operator takes the elementwise mean of the vectors in $(\{\mathbf{h}_v^K(i), \forall i \in \mathcal{P}\})$, assuming data holders contribute equally to the global node embeddings, i.e.,

$$\mathbf{h}_v^K \leftarrow \text{MEAN}(\mathbf{h}_v^K(1) \cup \mathbf{h}_v^K(2) \cup \dots \cup \mathbf{h}_v^K(I)). \quad (6)$$

Regression. The above two strategies treat data holders equally. In reality, the local node embeddings from different data holder may contribute diversely to the global node embeddings. We propose a Regression strategy to handle this kind of situation. Let ω_i be the weight vector of local node embeddings from data holder $i \in \mathcal{P}$, then

$$\mathbf{h}_v^K \leftarrow \omega_1 \odot \mathbf{h}_v^K(1) + \omega_2 \odot \mathbf{h}_v^K(2) \dots + \omega_I \odot \mathbf{h}_v^K(I), \quad (7)$$

where \odot is element-wise multiplication.

These different combination operators can utilize local node embeddings in diverse ways, and we will empirically study their effects on model performances in experiments.

3.5 Enhancing Privacy by Adopting DP

Data holders directly send the *local information*, e.g., local node embeddings during forward propagation and gradient update during back propagation, to the server may cause potential information leakage [Lyu *et al.*, 2020], and we propose to apply differential privacy to further enhance privacy. In this section, we introduce two DP based data publishing mechanisms, to further enhance the privacy of our proposed VFGNN. Such that with a single entry modification in the local information of data holders, there is a large probability that the server cannot distinguish the difference before or after the modification. We present the two mechanisms, i.e., Gaussian Mechanism and James-Stein Estimator, in Algorithm 1. We have described Gaussian mechanism in Section 2.3, we present James-Stein Estimator as follows.

Theorem 2. (*James-Stein Estimator and its adaptivity [Balle and Wang, 2018]*). Suppose d is the dimension of local information \mathbf{x} . When $d \geq 3$, substituting w in $\tilde{\mathbf{x}}_{\text{Bayes}}$ with its maximum likelihood estimate under $\mathbf{x} \sim \mathcal{N}(0, w^2 I)$, $\tilde{\mathbf{x}}|\mathbf{x} \sim \mathcal{N}(\mathbf{x}, \sigma^2 C^2 I)$, and $\tilde{\mathbf{x}}_{\text{Bayes}} = \text{argmin}_{\tilde{\mathbf{x}}} \|\tilde{\mathbf{x}} - \mathbf{x}\|^2$ produces James-Stein Estimator $\tilde{\mathbf{x}}_{\text{JS}} = \left(1 - \frac{(d-2)\sigma^2 C^2}{\|\tilde{\mathbf{x}}\|^2}\right) \tilde{\mathbf{x}}$. Moreover, it has an Mean Squared Error (MSE) of

$$E[\|\tilde{\mathbf{x}}_{\text{JS}} - \mathbf{x}\|^2] = d\sigma^2 \left(1 - \frac{(d-2)^2}{d^2} \frac{\sigma^2 C^2}{w^2 + \sigma^2 C^2}\right). \quad (8)$$

The MSE of Gaussian Mechanism $\tilde{\mathbf{x}}$ to exact \mathbf{x} is $E\|\tilde{\mathbf{x}} - \mathbf{x}\|^2 = d\sigma^2 C^2$, while the MSE of James-Stein Estimator is reduced with a factor of $(1 - \frac{(d-2)^2}{d^2} \frac{\sigma^2 C^2}{w^2 + \sigma^2 C^2})$. Both methods preserve (ϵ, δ) -DP while James-Stein estimator shows reductions in MSE, thus improves utility. By the definition of Gaussian mechanism (Definition 3), we have the privacy loss for both information publishing mechanisms in Algorithm 1. By combining it with Moment Accountant (MA) [Abadi *et al.*, 2016], we present the overall privacy for T iterations.

Theorem 3. Suppose each iteration of Algorithm 1 is (ϵ, δ) -DP. There exist constants c_1 and c_2 so that given the sampling probability q and the number of iterations T , and $\epsilon < c_1 q \sqrt{T}$, Algorithm 1 over T iteration is (ϵ', δ) -DP, with $\epsilon' = c_2 q \sqrt{T} \epsilon$.

Proof. By Definition 3 and Theorem 1, to ensure one iteration (ϵ, δ) -DP, we set $\sigma = \frac{\sqrt{2 \ln(1.25/\delta)}}{\epsilon}$. By Theorem 1 in [Abadi *et al.*, 2016], with $\sigma = \frac{\sqrt{2 \ln(1.25/\delta)}}{\epsilon}$ and the appropriate choice of ϵ, q, T , such that $\epsilon < c_1 q \sqrt{T}$, the privacy loss over T iterations is $\epsilon' = \frac{c_1 q \sqrt{T \log(1/\delta)}}{\sigma} = c_2 q \sqrt{T} \epsilon$. \square

3.6 Putting Together

By combining **CG1-CG3**, we complete the forward propagation of VFGNN. To describe the procedures in details, without loss of generality, we take GraphSAGE [Hamilton *et al.*, 2017] for example and present its forward propagation process in Algorithm 2. VFGNN can be learnt by gradient descent through minimizing the cross-entropy error over all labeled training examples. As can be seen, in VFGNN, both private data and model are hold by data holders themselves, thus data privacy can be better guaranteed.

4 Experiments

We conduct experiments to answer the following questions. **Q1:** whether VFGNN outperforms the GNN models that are trained on the isolated data. **Q2:** how does VFGNN behave comparing with the traditional insecure model trained on the plaintext mixed data. **Q3:** how does VFGNN perform comparing with the naive solution in Section 1. **Q4:** are our proposed combination strategies effective to VFGNN. **Q5:** what is the effect of the number of data holders on VFGNN. **Q6:** what is the effect of differential privacy on VFGNN.

4.1 Experimental Setup

Datasets. We use four benchmark datasets, i.e., Cora, Pubmed, Citeseer [Sen *et al.*, 2008], and arXiv [Hu *et al.*, 2020]. We use exactly the same dataset partition of training, validate, and test following the prior work [Kipf and Welling, 2016; Hu *et al.*, 2020]. Besides, in data isolated GNN setting, both node features and edges are hold by different parties. For all the experiments, we use five-fold cross validation and adopt average accuracy as the evaluation metric.

Comparison methods. We compare VFGNN with GraphSAGE models [Hamilton *et al.*, 2017] that are trained using isolated data and mixed plaintext data to answer **Q1** and **Q2**.

Algorithm 2 Privacy-preserving GraphSAGE for node label prediction (forward propagation)

Input: Data holder $\forall i \in \mathcal{P}$; Graph $\mathcal{G}(\mathcal{V}, \mathcal{E}^i)$ and node features $\{\mathbf{x}_v^i, \forall v \in \mathcal{V}\}$; depth K ; aggregator functions $\text{AGG}_k, \forall k \in \{1, \dots, K\}$; weight matrices $\mathbf{W}_i^k, \forall k \in \{1, \dots, K\}$; max layer L ; weight matrices $\mathbf{W}_l, \forall l \in \{0, \dots, L\}$; non-linearity σ ; neighborhood functions $\mathcal{N}^i : v \rightarrow 2^{\mathcal{V}}$; node labels on data holder $p \in \mathcal{P}$ and $c \in \mathcal{C}$

Output: Node label predictions $\{\hat{y}_{vc}, \forall v \in \mathcal{V}, \forall c \in \mathcal{C}\}$

- 1: # **CG1:** private feature and edge related computations
- 2: **Data holders:** jointly calculate initial node embeddings $\mathbf{h}_v^0(i) \leftarrow \mathbf{x}_v^i, \forall i \in \mathcal{P}, \forall v \in \mathcal{V}$
- 3: **for** $i \in \mathcal{P}$ **in parallel do**
- 4: **for** $k = 1$ to K **do**
- 5: **for** $v \in \mathcal{V}$ **do**
- 6: **Data holder:** calculates $\mathbf{h}_{\mathcal{N}^i(v)}^k(i) \leftarrow \text{AGG}_k(\{\mathbf{h}_u^{k-1}(i), \forall u \in \mathcal{N}^i(v)\})$
- 7: **end for**
- 8: **Data holder:** calculates $\mathbf{h}_v^k(i) \leftarrow \sigma(\mathbf{W}_i^k \cdot \text{CONCAT}(\mathbf{h}_v^{k-1}(i), \mathbf{h}_{\mathcal{N}^i(v)}^k(i)))$
- 9: **end for**
- 10: **Data holder:** calculates local node embeddings $\mathbf{h}_v^K(i) \leftarrow \mathbf{h}_v^K(i) / \|\mathbf{h}_v^K(i)\|_2, \forall v \in \mathcal{V}$ and sends (publishes) it to server using differential privacy
- 11: **end for**
- 12: # **CG2:** non-private data related computations
- 13: **for** $v \in \mathcal{V}$ **do**
- 14: **Server:** combines the local node embeddings from data holders $\mathbf{h}_v^K = \text{COMBINE}(\{\mathbf{h}_v^K(i), \forall i \in \mathcal{P}\})$
- 15: **Server:** forward propagation based on the global node embeddings $\mathbf{z}_L = \sigma(\mathbf{W}_{L-1} \cdot \sigma(\dots \sigma(\mathbf{W}_0 \cdot \mathbf{h}_v^K)))$
- 16: **Server:** sends \mathbf{z}_L to data holder p
- 17: **end for**
- 18: # **CG3:** private label related computations
- 19: **Data holder** p : makes prediction by $\hat{y}_{vc} \leftarrow \text{softmax}(\mathbf{W}_L \cdot \mathbf{z}_L), \forall v \in \mathcal{V}, \forall c \in \mathcal{C}$

We also compare VFGNN with the naive solution described in Section 1 to answer **Q3**. To answer **Q4**, we vary the proportion of the data (features and edges) hold by \mathcal{A} and \mathcal{B} , and change VFGNN with different combination strategies. We vary the number of data holders in VFGNN to answer **Q5**, and vary the parameters of differential privacy to answer **Q6**. For all these models, we choose Mean operator as the aggregator function.

Parameter settings. For all the models, we use TanH as the active function of neighbor propagation, and Sigmoid as the active function of hidden layers. For the deep neural network on server, we set the dropout rate to 0.5 and network structure as $(d, d, |C|)$, where $d \in \{32, 64, 128\}$ is the dimension of node embeddings and $|C|$ the number of classes. We vary $\epsilon \in \{1, 2, 4, 8, 16, 32, 64, \infty\}$, set $\delta = 1e^{-4}$ and the clip value $C = 1$ to study the effects of differential privacy on our model. Since we have many comparison and ablation models, and they achieve the best performance with different parameters, we cannot report all the best parameters. In-

Dataset	#Node	#Edge	#Features	#Classes
Cora	2,708	5,429	1,433	7
Pubmed	19,717	44,338	500	3
Citeseer	3,327	4,732	3,703	6
arXiv	169,343	2,315,598	128	40

Table 1: Dataset statistics.

Dataset	Cora	Pubmed	Citeseer	arXiv
GraphSAGE _A	0.611	0.672	0.541	0.471
GraphSAGE _B	0.606	0.703	0.457	0.482
VFGNN_C	0.790	0.774	0.685	0.513
VFGNN_M	0.809	0.781	0.695	0.522
VFGNN_R	0.802	0.782	0.693	0.518
GraphSAGE _{A+B}	0.815	0.791	0.700	0.529

Table 2: Comparison results on three datasets (Q1 and Q2).

stead, we report the range of the best parameters. We vary the propagation depth $K \in \{2, 3, 4, 5\}$, L2 regularization in $\{10^{-2} - 10^{-4}\}$, and learning rate in $\{10^{-2} - 10^{-3}\}$. We tune parameters based on the validate dataset and evaluate model performance on the test dataset.

4.2 Comparison Results and Analysis

To answer **Q1-Q3**, we assume there are two data holders (\mathcal{A} and \mathcal{B}) who have equal number of node features and edges, i.e., the proportion of data held by \mathcal{A} and \mathcal{B} is 5:5, and compare our models with GraphSAGEs that are trained on isolated data individually and on mixed plaintext data. We also set $\epsilon = \infty$ during comparison and will study its effects later. We summarize the results in Table 2, where VFGNN_C, VFGNN_M, and VFGNN_R denote VFGNN with Concat, Mean, and Regression combination strategies.

Result1: answer to Q1. We first compare VFGNNs with the GraphSAGEs that are trained on isolated feature and edge data, i.e., GraphSAGE_A and GraphSAGE_B. From Table 2, we find that, VFGNNs with different combination strategies significantly outperforms GraphSAGE_A and GraphSAGE_B on all the three datasets. Take Citeseer for example, our VFGNN_R improves GraphSAGE_A and GraphSAGE_B by as high as 28.10% and 51.64%, in terms of accuracy.

Analysis of Result1. The reason of result1 is straightforward. GraphSAGE_A and GraphSAGE_B can only use partial feature and edge information held by \mathcal{A} and \mathcal{B} . In contrast, VFGNNs provide a solution for \mathcal{A} and \mathcal{B} to jointly train GNNs without compromising their own data. By doing this, VFGNNs can use the information from the data of both \mathcal{A} and \mathcal{B} simultaneously, and therefore achieve better performance.

Result2: answer to Q2. We then compare VFGNNs with GraphSAGE that is trained on the mixed plaintext data, i.e., GraphSAGE_{A+B}. It can be seen from Table 2 that VFGNNs have comparable performance with GraphSAGE_{A+B}, e.g., 0.8090 vs. 0.8150 on Cora dataset and 0.6950 vs. 0.7001 on Citeseer dataset.

Analysis of Result2. It is easy to explain why our proposal has comparable performance with the model that are trained

Model	VFGNN_C	VFGNN_M	VFGNN_R
Prop.=9:1	0.809	0.805	0.809
Prop.=8:2	0.802	0.796	0.807
Prop.=7:3	0.793	0.793	0.803

Table 3: Comparison of combination operators on Cora by varying the proportion of data hold by \mathcal{A} and \mathcal{B} (Q4).

No. of DH	VFGNN_C	VFGNN_M	VFGNN_R
2	0.790	0.809	0.802
3	0.749	0.774	0.760
4	0.712	0.733	0.722

Table 4: Comparison results on Cora by varying the number of data holders (Q5).

on the mixed plaintext data. First, we propose a secret sharing based protocol for \mathcal{A} and \mathcal{B} to generate the initial node embeddings from their node features, which are the same as those generated by using mixed plaintext features. Second, although \mathcal{A} and \mathcal{B} generate local node embeddings by using their own edge data to do neighbor aggregation separately (for security concern), we propose different combination strategies to combine their local node embeddings. Eventually, the edge information from both \mathcal{A} and \mathcal{B} is used for training the classification model. Therefore, VFGNN achieves comparable performance with GraphSAGE _{$\mathcal{A}+\mathcal{B}$} .

Result3: answer to Q3. In VFGNN, we delegate the non-private data related computations to server. One would be curious about what if these computations are also performed by data holders using existing secure neural network protocols, i.e., SecureML [Mohassel and Zhang, 2017]. To answer this question, we compare VFGNN with the secure GNN model that is implemented using SecureML, which we call as SecureGNN, where we use 3-degree Taylor expansion to approximate TanH and Sigmoid. The accuracy and running time per epoch (in seconds) of VFGNN vs. SecureGNN on Pubmed are 0.8090 vs. 0.7970 and 18.65 vs. 166.81, respectively, where we use local area network.

Analysis of Result3. From the above comparison results, we find that our proposed VFGNN learning paradigm not only achieves better accuracy, but also has much better efficiency. This is because the non-private data related computations involve many non-linear functions that are not cryptographically friendly, which have to be approximately calculated using time-consuming MPC techniques in SecureML.

4.3 Ablation Study

We now study the effects of different combination operators and different number of data holders on VFGNN.

Result4: answer to Q4. Different combination operators can utilize local node embeddings in diverse ways and make our proposed VFGNN adaptable to different scenarios, we study this by varying the proportion (Prop.) of data (node features and edges) hold by \mathcal{A} and \mathcal{B} in {9: 1, 8: 2, 7: 3}. The results on Cora dataset are shown in Table 3.

Analysis of Result4. From Table 3, we find that with the

ϵ	4	8	16	32	64
Gaussian	0.502	0.702	0.772	0.789	0.794
James-Stein	0.510	0.706	0.781	0.799	0.804

Table 5: Effect of DP on VFGNN using Cora dataset (Q6).

proportion of data hold by \mathcal{A} and \mathcal{B} being even, i.e., from {9: 1} to {7: 3}, the performances of most strategies tend to decrease. This is because the neighbor aggregation is done by data holders individually, and with a bigger proportion of data hold by a single holder, it is easier for this party to generate better local node embeddings. Moreover, we also find that Mean operator works well when data are evenly split, and Regression operator is good at handling the situations where data holders have different quality of data, since it treats the local node embeddings from each data holder differently, and assigns weights to them intelligently.

Result5: answer to Q5. We vary the number of data holders in {2, 3, 4} and study the performance of VFGNN. We report the results in Table 4, where we use the Cora dataset and assume data holders have even feature and edge data.

Analysis of Result5. From Table 4, we find that, as the number of data holders increases, the accuracy of all the models decreases. This is because the neighborhood aggregation in VFGNN is done by each holder individually for privacy concern, and each data holder will have less edge data when there are more data holders, since they split the original edge information evenly. Therefore, when more participants are involved, more information will be lost during the neighborhood aggregation procedure.

Result6: answer to Q6. We present the privacy loss of each iteration in Table 5 and the over all privacy in Theorem 3. We vary ϵ and set $\delta = 1e^{-4}$ to study the effects of DP on VFGNN. We report the results in Table 5, where we use Cora dataset, use MEAN as the combination operator, and assume data holders have even feature and edge data.

Analysis of Result6. From Table 5, we can see that the accuracy of VFGNN increases with ϵ . In other words, there is a trade-off between accuracy and privacy. The smaller ϵ , the more noise will be added into the local node embeddings, which causes stronger privacy guarantee but lower accuracy. We also find James-Stein estimator consistently works better than Gaussian mechanism, since it can reduce MSE, as we have analyzed in Section 3.5.

5 Conclusion

We propose VFGNN, a vertically federated GNN learning paradigm for privacy-preserving node classification task. We finish this by splitting the computation graph of GNN. We leave the private data related computations on data holders and delegate the rest computations to a server. Experiments on real world datasets demonstrate that our model significantly outperforms the GNNs by using the isolated data and has comparable performance with the traditional GNN by using the mixed plaintext data insecurely.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (No. 62172362) and “Leading Goose” R&D Program of Zhejiang (No. 2022C01126).

References

- [Abadi *et al.*, 2016] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *CCS*, pages 308–318. ACM, 2016.
- [Acar *et al.*, 2018] Abbas Acar, Hidayet Aksu, A Selcuk Uluagac, and Mauro Conti. A survey on homomorphic encryption schemes: Theory and implementation. *ACM Computing Surveys (CSUR)*, 51(4):79, 2018.
- [Balle and Wang, 2018] Borja Balle and Yu-Xiang Wang. Improving the gaussian mechanism for differential privacy: Analytical calibration and optimal denoising. *arXiv preprint arXiv:1805.06530*, 2018.
- [Beaver, 1991] Donald Beaver. Efficient multiparty protocols using circuit randomization. In *Annual International Cryptology Conference*, pages 420–432. Springer, 1991.
- [Chen *et al.*, 2021] Chaochao Chen, Jun Zhou, Li Wang, Xibin Wu, Wenjing Fang, Jin Tan, Lei Wang, Alex X. Liu, Hao Wang, and Cheng Hong. When homomorphic encryption marries secret sharing: Secure large-scale sparse logistic regression and applications in risk control. In *SIGKDD*, pages 2652–2662. ACM, 2021.
- [Dwork *et al.*, 2014] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [Fang *et al.*, 2021] Wenjing Fang, Derun Zhao, Jin Tan, Chaochao Chen, Chaofan Yu, Li Wang, Lei Wang, Jun Zhou, and Benyu Zhang. Large-scale secure XGB for vertical federated learning. In *CIKM*, pages 443–452. ACM, 2021.
- [Gu *et al.*, 2018] Zhongshu Gu, Heqing Huang, Jialong Zhang, Dong Su, Ankita Lamba, Dimitrios Pendarakis, and Ian Molloy. Securing input data of deep learning inference systems via partitioned enclave execution. *arXiv preprint arXiv:1807.00969*, 2018.
- [Hamilton *et al.*, 2017] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *NeurIPS*, pages 1024–1034, 2017.
- [Hardy *et al.*, 2017] Stephen Hardy, Wilko Henecka, Hamish Ivey-Law, Richard Nock, Giorgio Patrini, Guillaume Smith, and Brian Thorne. Private federated learning on vertically partitioned data via entity resolution and additively homomorphic encryption. *arXiv preprint arXiv:1711.10677*, 2017.
- [He *et al.*, 2020] Xinlei He, Jinyuan Jia, Michael Backes, Neil Zhenqiang Gong, and Yang Zhang. Stealing links from graph neural networks. *arXiv preprint arXiv:2005.02131*, 2020.
- [Hu *et al.*, 2020] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. *arXiv preprint arXiv:2005.00687*, 2020.
- [Kipf and Welling, 2016] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [Lyu *et al.*, 2020] Lingjuan Lyu, Han Yu, Jun Zhao, and Qiang Yang. Threats to federated learning. In *Federated Learning*, pages 3–16. Springer, 2020.
- [Mohassel and Zhang, 2017] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *S&P*, pages 19–38, 2017.
- [Osia *et al.*, 2019] Seyed Ali Osia, Ali Shahin Shamsabadi, Ali Taheri, Kleomenis Katevas, Sina Sajadmanesh, Hamid R Rabiee, Nicholas D Lane, and Hamed Haddadi. A hybrid deep learning architecture for privacy-preserving mobile analytics. *arXiv preprint arXiv:1703.02952*, 2019.
- [Sajadmanesh and Gatica-Perez, 2020] Sina Sajadmanesh and Daniel Gatica-Perez. Locally private graph neural networks. *CoRR abs/2006.05535*, 12, 2020.
- [Sen *et al.*, 2008] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- [Shamir, 1979] Adi Shamir. How to share a secret. *Communications of the ACM*, 22(11):612–613, 1979.
- [Veličković *et al.*, 2017] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [Vepakomma *et al.*, 2018] Praneeth Vepakomma, Otkrist Gupta, Tristan Swedish, and Ramesh Raskar. Split learning for health: Distributed deep learning without sharing raw patient data. *arXiv preprint arXiv:1812.00564*, 2018.
- [Wagh *et al.*, 2019] Sameer Wagh, Divya Gupta, and Nishanth Chandran. Securenn: 3-party secure computation for neural network training. *PETs*, 1:24, 2019.
- [Wu *et al.*, 2019] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S Yu. A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*, 2019.
- [Wu *et al.*, 2021] Chuhan Wu, Fangzhao Wu, Yang Cao, Yongfeng Huang, and Xing Xie. Fedgnn: Federated graph neural network for privacy-preserving recommendation. *arXiv preprint arXiv:2102.04925*, 2021.
- [Zheng *et al.*, 2021] Longfei Zheng, Jun Zhou, Chaochao Chen, Bingzhe Wu, Li Wang, and Benyu Zhang. Asfgnn: Automated separated-federated graph neural network. *PPNA*, pages 1–13, 2021.