

Towards Private Learning on Decentralized Graphs with Local Differential Privacy

Wanyu Lin, *Member, IEEE*, Baochun Li, *Fellow, IEEE*, and Cong Wang, *Fellow, IEEE*

Abstract—Many real-world networks are inherently decentralized. For example, in social networks, each user maintains a local view of a social graph, such as a list of friends and her profile. It is typical to collect these local views of social graphs and conduct graph learning tasks. However, learning over graphs can raise privacy concerns as these local views often contain sensitive information. In this paper, we seek to ensure private graph learning on a decentralized network graph. Towards this objective, we propose *Solitude*, a new privacy-preserving learning framework based on graph neural networks (GNNs), with formal privacy guarantees based on edge local differential privacy. The crux of *Solitude* is a set of new delicate mechanisms that can calibrate the introduced noise in the decentralized graph collected from the users. The principle behind the calibration is the intrinsic properties shared by many real-world graphs, such as sparsity. Unlike existing work on locally private GNNs, our new framework can simultaneously protect node feature privacy and edge privacy, and can seamlessly incorporate with any GNN with privacy-utility guarantees. Extensive experiments on benchmarking datasets show that *Solitude* can retain the generalization capability of the learned GNN while preserving the users' data privacy under given privacy budgets.

Index Terms—Privacy-Preserving Graph Learning, Graph Neural Networks, Differential Privacy, Decentralized Network Graph



1 INTRODUCTION

Many problems in scientific domains, ranging from computer networks [1] and social networks to biomedicine and healthcare [2], [3], can be naturally cast as problems of property learning on graphs. Typically, these graph domains contain sensitive information. In computer networks, for example, the goal of botnets detection is to isolate the botnet nodes, where the problem can be formulated as binary node classification task on massive background Internet communication graphs [4]. However, the Internet Service Providers (ISPs) may be reluctant to share traffic observations (modeled as edges in the communication graphs [4]). Likewise, in social networks, users' contact lists, profile information, likes or comments, etc., should be kept private, as most users are not willing to release their contact lists to strangers. Inevitably, the use of sensitive and private graph data requires principled and rigorous privacy guarantees.

On the other hand, among various graph learning algorithms, graph neural networks (GNNs) have exhibited superior performance [4], due to their efficiency and inductive learning capability [5]. A recent work shows that the edge information of graphs can be recovered from GNNs through influence analysis [6]. Therefore, it is appealing to tailor the GNNs to perform graph learning tasks while still preserving user data privacy. A plausible approach is local differential privacy (LDP) [7], where each user locally

obfuscates their share of data before sending them to a data curator (who may be malicious). As the data curator performs model learning over the obfuscated data, data privacy may be preserved under given privacy budgets. However, most existing LDP techniques for graphs mainly focused on graph statistics analysis while protecting the privacy of edge/link information. Typical graph statistics include subgraph counting [8] (e.g., triangles and k -stars counting) and graph metric estimations [9] (e.g., clustering coefficient, modularity, or centrality estimation), which are not designed for GNNs.

Learning over the obfuscated graphs with GNNs is quite challenging. In general, training deep learning models with strong differential privacy guarantees comes at a significant cost in utility [10], [11]. Specifically, in the context of graph learning, graph data usually contains *node feature* information and *graph structure* information. The combinatorial nature of graph structures makes the private learning problem more complex than other domains. These can be explained as follows. In GNNs, the node representations — can be used for various downstream tasks — are learned in a way that node information is aggregated and propagated via links through the message passing framework during training [5], [12]. In node classification tasks or link prediction tasks, the training samples (nodes or links in the graph) of the learning model are interdependent. In contrast, the existing work on LDP for tabular data assumes that each user's data is independently and identically drawn from an underlying distribution [13], [14], which is problematic in the context of graphs.

While private graph learning with GNNs is still a nascent research topic, a recent proposal attempts to preserve the privacy of the node features [15]. However, this model raises several privacy and security issues as it assumes that the data curator holds the global graph structure.

- *W. Lin is with the Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China.
E-mail: wanyulin@comp.polyu.edu.hk.*
- *B. Li is with the Department of Electrical and Computer Engineering, University of Toronto, Ontario, M5S 1A1, Canada.
Email: bli@ece.toronto.edu.*
- *C. Wang is with the Department of Computer Science, City University of Hong Kong, Hong Kong, China.
E-mail: congwang@cityu.edu.hk.*

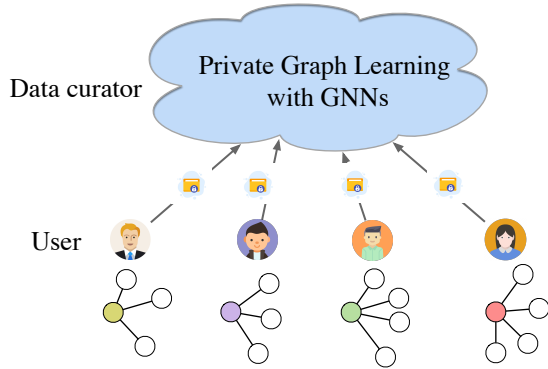


Fig. 1: The illustration of private learning over decentralized graph network: 1) each user holds a data share including a local neighbour list and the profile information; 2) the data curator collects the obfuscated data and conducts learning over the collected noisy graphs.

If the topological features contain sensitive information, this approach may incur information leakage as the data curator can directly access the global topology for the message passing process. Therefore, in this paper, we study the problem of differentially private graph learning with GNNs on a decentralized network graph, as shown in Fig. 1. In particular, the data curator cannot directly access the global structure of the graph. In other words, both the node feature and graph structure information should be protected against the data curator.

Privacy-preserving learning under the setting of decentralized network graphs has various applications, including but not limited to social network analysis and mobile computing. In principle, some social networks are inherently decentralized and distributed, such as Synereo [16]. Though in some other social networks, e.g., Facebook, there exists a centralized party that holds the knowledge of the global network, that party may choose not to share it with a third-party (corresponding to the data curator) for analyzing, due to legal issues or other business concerns.

With the prevalence of decentralized network graphs, we propose a new framework consisting of a set of mechanisms, called *Solitude*, to tailor the GNNs for decentralized graph analysis under local differential privacy. Our framework has provable privacy guarantees based on local differential privacy. Specifically, we leverage the notion of edge local differential privacy proposed in [17]. To protect the privacy of neighbor lists, each user applies Warner’s randomized response mechanism [18] to obfuscate their neighbor lists before sending them to the data curator. For protecting node/user¹ feature privacy, we further incorporate a multi-bit mechanism for multi-dimensional feature perturbation [15], [19], which can ensure that the high-dimensional feature vector of every user can be protected. However, learning over obfuscated graphs introduces challenges, as it could significantly degrade the generalization capability of the GNN. Precisely, the learned GNN may overfit the noisy graphs and generalize poorly to unseen nodes.

To this end, we propose new mechanisms for graph structure calibration and feature vector calibration, respec-

1. In many applications, such as social networks, each node represents a user; we use “node” and “user” interchangeably.

TABLE 1: Notations

Notation	Descriptions
\mathbf{a}_i	the adjacency list of user i
$\tilde{\mathbf{a}}_i$	randomized adjacency list of user i
\mathbf{x}_i	the feature vector of user i
$\tilde{\mathbf{x}}_i$	the encoded feature vector of user i
$\hat{\mathbf{x}}_i$	the rectified feature vector of user i
\mathcal{M}_a	the randomized mechanism for adjacency lists
\mathcal{M}_x	the randomized mechanism for node features
ϵ_x, ϵ_a	privacy budgets for features and adjacency lists
$ \mathcal{N}(\cdot) $	the node degree of user i
\mathbf{A}^c	the adjacency matrix after calibration
\mathbf{X}^c	the node feature matrix after calibration
$\ \cdot\ _{\mathbf{F}}^2$	the Frobenius norm of a matrix
$\ \cdot\ _1$	l_1 norm operator
θ	the model parameters
λ_1, λ_2	the regularization coefficients

tively. Essentially, obfuscating graph structure via a randomized response mechanism introduces edge deletion, addition, and rewiring. Though the obfuscating process satisfies our privacy goal, we theoretically and empirically analyze that this process tends to return a much denser graph than the original one, violating the sparsity property — an inherent property exhibited in many real-world graphs [20]. Therefore, we propose encouraging the sparseness of the graph structure during training as a calibration step to reduce the effect of the noise. In addition, a node might be linked to nodes with task-specific “noise,” leading to the aggregation of non-smooth features in GNNs. Inspired by the assumption of *feature smoothness* in GNNs, we leverage a feature smoothing component before training to reduce the effect of the feature noise. To further boost the prediction accuracy, a label smoothness component is adopted, which is inspired by the principles of label propagation algorithms [21], [22].

In a nutshell, our original contributions are listed as follows. We propose a new privacy-preserving learning framework for decentralized network graphs based on graph neural networks. It consists of a set of mechanisms that can provide local differential privacy for every user yet maintaining the generalization capability of the learned GNN. We formally analyze that our mechanisms preserve local differential privacy for every user, particularly on the notion of edge differential privacy. Towards our goal, we theoretically and empirically analyze the overfitting problem while learning over the noisy graphs. Different from the literature on locally private GNNs, *Solitude* can preserve edge privacy and node feature privacy for every user simultaneously. Our extensive array of experiments on benchmarking datasets demonstrated that *Solitude* can significantly improve the privacy-utility guarantees on canonical graph learning benchmarks. We also empirically show that our framework can seamlessly integrate with any GNN architectures, such as GCN [12] and GraphSage [5], with privacy-utility guarantees.

2 PROBLEM SETUP

2.1 Notations and Problem Definition

Notations. We consider a network graph, denoted as

$\mathcal{G} = (\mathcal{V}, \mathbf{A}, \mathbf{X})$. Specifically, we consider the decentralized setting, where the entire graph is decentralized over users/nodes $\mathcal{V} = \{v_1, \dots, v_{|\mathcal{V}|}\}$. Precisely, each user v_i holds locally a neighbor list depicted as \mathbf{a}_i , which can be modeled as an $|\mathcal{V}|$ -binary vector, and a node attribute/feature vector $\mathbf{x}_i \in \mathbb{R}^{|D|}$, where $|D|$ is the dimension of the user/node feature vector. The corresponding adjacency matrix of the entire graph can be represented as $\mathbf{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_{|\mathcal{V}|}\}$, where $\mathbf{A} \in \mathbb{R}^{|\mathcal{V}| \times |\mathcal{V}|}$, and the node attribute matrix $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_{|\mathcal{V}|}\}$, where $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times |D|}$. In the context of social network graphs, for example, each user locally holds their friend list and the profile information (e.g., WeChat [23]). Without loss of generality, we assume \mathcal{G} is a directed graph; this reflects the real-world application domains, e.g, the follower-followee relationships.

Consistent with prior work [15], we focus on the node classification task with graph neural networks. Specifically, we follow the standard node classification setting, which is commonly employed in various literature [5], [12]. Given a set of labeled nodes $\mathcal{V}_l \subset \mathcal{V}$, with class labels from $\mathcal{Y} = \{y_1, y_2, y_3, \dots, y_K\}$ and a set of unlabeled nodes $\mathcal{V}_u \subset \mathcal{V} \setminus \mathcal{V}_l$, the goal of node classification is to map each node $v \in \mathcal{V}$ to one class in \mathcal{Y} .

Problem Definition. We assume that the data curator is an untrusted party, and it can access the set/index of nodes/users $\mathcal{V} = \{v_1, \dots, v_{|\mathcal{V}|}\}$ and the node labels of the training set \mathcal{V}_l . However, the data curator could not directly access the feature matrices \mathbf{X} and \mathbf{A} , which are decentralized among the users and private to the users. The data curator is allowed to collect information from the users, and performs node classification with GNNs over the noisy graphs.

Therefore, our ultimate goal is to obtain a set of mechanisms that 1) the data curator can collect data portions from each user while preserving the user data privacy; 2) the data curator can train a GNN for node classification over the collected noisy graph with the best possible generalization capability. Without loss of generality, model generalization capability is measured by the prediction accuracy on the held-out test set that has not been seen during training. Note that, different from the state-of-the-art on locally private GNN [15] — LPGNN, we consider a more advanced setting that the global topology of the graph is not accessible to the data curator. Concretely, user v_i 's private data includes the neighbor list, or called *adjacency list*, \mathbf{a}_i and the node attribute/feature vector \mathbf{x}_i . In what follows, we first briefly introduce some necessary background on graph learning with GNNs and the notion of local differential privacy on graphs to facilitate a better understanding of our solution. The mathematical notations used in this paper are summarized in Table 1.

2.2 Message Passing Graph Neural Networks

Graph neural networks (GNNs) are tailored to learn and model information structured as graph data. It is a family of graph message passing architectures that incorporate graph structure and node feature vectors to learn a dense representation of a node or the entire graph. In principle, GNNs share a neighborhood aggregation strategy, where the node representations are refined via iteratively aggregating the representations from its neighboring nodes in

the graph. Representative GNNs are graph convolutional networks, which use mean pooling for aggregation [12], and GraphSage that aggregates the node features via mean/max/LSTM pooling [5].

Taking GCNs as an example, the basic operator for the neighborhood information aggregation is the element-wise **mean**. After L iterations of aggregation, a node's representation can capture the structural information within its L -hop graph neighborhood, which can be formulated as:

$$h_v^{l+1} \leftarrow \sigma \left(W^l \cdot \mathbf{mean} \left(\{h_v^l\} \cup \{h_u^l, \forall u \in \mathcal{N}(v)\} \right) \right) \quad (1)$$

where W^l is a trainable matrix for layer l , σ denotes a nonlinear activation function. Note that, except for **mean** operator, there are many other operators for neighborhood information aggregation, such as **max** pooling. Due to their superior performance, these operators have been the core of many graph neural networks. For more details on other GNN variations, we refer the interested readers to the existing survey [24].

2.3 Local Differential Privacy

Local differential privacy (LDP) has emerged as the *de facto* solution for collecting private data and performing statistical queries, such as mean, counting, etc. In principle, it is a privacy metric to protect the sensitive information of individuals from the data curator [7], [25]. In the setting of LDP, each user does not trust the data curator. Before sending her share of data to the data curator, each user locally perturbs the data portion with a differentially private mechanism. The perturbed data is not meaningful individually but can be used for data analytics when aggregated.

In the context of graph data, a differentially private mechanism can be designed for edge differential privacy [26], or node differential privacy [27]. In essence, edge differential privacy ensures that a randomized mechanism does not reveal the addition or deletion of an edge in the neighbor list of an individual. In contrast, a randomized mechanism for node differential privacy hides the deletion or addition of a node along with its link list. Consistent with prior works [9], [17], we adopt the notion of edge local differential privacy (LDP) based on a user's neighbor list.

Formally, let $\mathbf{a}_i = (a_{i,1}, \dots, a_{i,|\mathcal{V}|}) \in \{0, 1\}^{|\mathcal{V}|}$ be the neighbor list of the user v_i , where a_i is the i -th row of the adjacency matrix \mathbf{A} of the global network graph \mathcal{G} . Stated differently, the adjacency matrix of \mathcal{G} can be represented as $\mathbf{A} = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_{|\mathcal{V}|}\}$. Then edge LDP can be defined as follows [8], [17]:

Definition 1 (Edge local differential privacy). A randomized mechanism \mathcal{M} satisfies ϵ -edge local differential privacy (ϵ -edge LDP) if and only if for any two neighbor lists \mathbf{a} and $\tilde{\mathbf{a}}$, such that \mathbf{a} and $\tilde{\mathbf{a}}$ only differ in one bit, and any $s \subseteq \text{range}(\mathcal{M})$, we have

$$\frac{\Pr[\mathcal{M}(\mathbf{a}) = s]}{\Pr[\mathcal{M}(\tilde{\mathbf{a}}) = s]} \leq e^\epsilon, \quad (2)$$

where ϵ is the privacy budget.

Note that, ϵ -edge LDP in Definition 1 protects one single bit in a neighbor list with privacy budget ϵ . By adopting the notion of group privacy [18], ϵ -edge LDP can be used to

protect $k \in \mathbb{N}$ edges. Concretely, if \mathcal{M} provides ϵ -edge LDP, then for any two neighbor lists \mathbf{a} and $\tilde{\mathbf{a}}$ that differ in k edges and any $s \subseteq \text{range}(\mathcal{M})$, we have

$$\frac{\Pr[\mathcal{M}(\mathbf{a}) = s]}{\Pr[\mathcal{M}(\tilde{\mathbf{a}}) = s]} \leq e^{k\epsilon}, \quad (3)$$

where k edges are protected with privacy budget $k\epsilon$.

Properties. LDP satisfies composition property and transformation invariance [27]. Specifically, composition property enables the modular design of mechanisms: if all the components of a mechanism are differentially private, so is their composition. The transformation invariance demonstrates that performing post-processing on the output of an algorithm that satisfies LDP does not affect the privacy guarantee.

3 Solitude: OUR PROPOSED FRAMEWORK

This section describes the main components of our proposed framework, called *Solitude*, toward differentially private training of GNN over private graph data, including the node feature vectors and the neighbor lists that are decentralized across all users. Specifically, there are two stages in our framework. In the first stage, the data curator sends a query to each user v_i once, and then each user v_i independently sends an answer – the obfuscated data share, including the obfuscated version of node feature vector $\tilde{\mathbf{x}}_i$ and the obfuscated adjacency list $\tilde{\mathbf{a}}_i$ (shown in Fig. 1). In the second stage, the data curator performs training of the GNN over the noisy graph composed of the obfuscated data shares from all users. The second stage of our *Solitude* is illustrated in Fig. 2.

In what follows, we first introduce the technical details of the randomized mechanisms used for preserving data privacy satisfying local differential privacy. Then we theoretically and empirically analyze the overfitting issue caused by the randomized mechanisms. We show how to calibrate the introduced noise such that the generalization capability of the trained GNN can be retained.

3.1 Obfuscating Local Data under Local Differential Privacy

In the setting of the decentralized network graph, each user v_i holds a data portion of the entire graph, including their feature vector \mathbf{x}_i and the adjacency list \mathbf{a}_i , both of which are private to the user. In the subsequent paragraphs, we describe the used mechanisms for protecting the privacy of the feature vector and the adjacency list, respectively.

Randomized Adjacency List. Intrinsically, an adjacency list is a binary bit vector, denoted as $\mathbf{a}_i = \{a_{i,1}, \dots, a_{i,|\mathcal{V}|}\}$, in which $a_{i,j} = 1$ indicates the link between v_i and v_j . Without loss of generality, we leverage of a common methodology, called *randomized response* [17], [18], to impel local differential privacy. Specifically, each user flips each bit of her adjacency list with a probability p constrained by a given privacy budget. More formally, given a privacy budget ϵ_a , the randomized adjacency list, denoted as $\tilde{\mathbf{a}}_i = \{\tilde{a}_{i,1}, \dots, \tilde{a}_{i,|\mathcal{V}|}\}$, is obtained as follows:

$$\tilde{a}_{i,j} = \begin{cases} a_{i,j}, & q = \frac{e^{\epsilon_a}}{1+e^{\epsilon_a}} \\ 1 - a_{i,j}, & p = \frac{1}{1+e^{\epsilon_a}} \end{cases}, \quad (4)$$

where $q = 1 - p$ is the probability of retaining a particular bit in the adjacency list.

Theorem 3.1. The randomized adjacency list mechanism \mathcal{M}_a satisfies ϵ_a -edge local differential privacy.

Proof. Let us consider the case that \mathbf{a}_i and $\tilde{\mathbf{a}}_i$ only differ in one bit. Concretely, we assume that $a_{i,j} \neq \tilde{a}_{i,j}$, and given any output $s = (s_1, \dots, s_n)$ from \mathcal{M} , we have

$$\begin{aligned} \frac{\Pr[\mathcal{M}_a(\mathbf{a}_i) = s]}{\Pr[\mathcal{M}_a(\tilde{\mathbf{a}}_i) = s]} &= \frac{\Pr[a_{i,1} \rightarrow s_1] \cdots \Pr[a_{i,n} \rightarrow s_n]}{\Pr[\tilde{a}_{i,1} \rightarrow s_1] \cdots \Pr[\tilde{a}_{i,n} \rightarrow s_n]} \quad (5) \\ &= \frac{\Pr[a_{i,j} \rightarrow s_j]}{\Pr[\tilde{a}_{i,j} \rightarrow s_j]} = \frac{q}{p} \leq e^{\epsilon_a}. \quad (6) \end{aligned}$$

Randomized Feature Vector. If the privacy of the node feature is also the concern, we further leverage a multi-bit mechanism to protect its privacy for every user. The randomized mechanism for feature vector in our framework follows the similar outline as the mechanism used in [15], [19]. Specifically, the multi-bit mechanism has two components: an encoder and a rectifier, as shown in Fig. 3. To randomize a feature vector $\mathbf{x}_i \in \mathcal{R}^{|D|}$, where each element $x_{i,j}$ falls into the range $[x_{\min}, x_{\max}]$, the encoder first uniformly samples m features out of the D dimensions. Each of the selected features is encoded into -1 or 1 , with a probability formulated as

$$\frac{1}{e^{\epsilon_x/m} + 1} + \frac{x_{i,j} - x_{\min}}{x_{\max} - x_{\min}} \cdot \frac{e^{\epsilon_x/m} - 1}{e^{\epsilon_x/m} + 1}. \quad (7)$$

Correspondingly, the rest of the $d - m$ features are mapped to 0s. The rectifier is to calibrate the encoded vector $\hat{\mathbf{x}}$ to ensure the outcome of the randomized mechanism $\tilde{\mathbf{x}}$ is statistically unbiased. Formally, the rectifier is instantiated as

$$\begin{aligned} \text{Rec}(\hat{x}_{i,j}) &= \frac{|D| \cdot (x_{\max} - x_{\min})}{2m} \cdot \frac{e^{\epsilon_x/m} + 1}{e^{\epsilon_x/m} - 1} \cdot \hat{x}_{i,j} \quad (8) \\ &\quad + \frac{x_{\max} + x_{\min}}{2}. \quad (9) \end{aligned}$$

Theorem 3.2. The randomized mechanism for node feature vector \mathcal{M}_x preserves ϵ_x differential privacy of every user.

Due to the page limitation, we refer to the detailed proof of Theorem 3.2 to [15], [19]. Note that the multi-bit mechanism for feature vector is different from the notion of node local differential privacy (node-LDP). Specifically, node-LDP hides the deletion or addition of a node along with its neighbor list; it is out of the scope in this paper. With the composition property of local differential privacy as described in Sec. 2.3, we arrive at the following corollary:

Corollary 3.1. The randomized mechanisms (\mathcal{M}_a and \mathcal{M}_x) together satisfy $\epsilon_a + \epsilon_x$ differential privacy of every user.

3.2 Private GNN Training with Calibration

Now with the collected graph data from all users, the data curator can reconstruct the global graph, and it is ready to perform graph analytics with message passing GNNs; the task is instantiated with node classification in this work. Yet it can be generalized to other graph learning tasks, such as link prediction, clustering coefficient prediction [28],

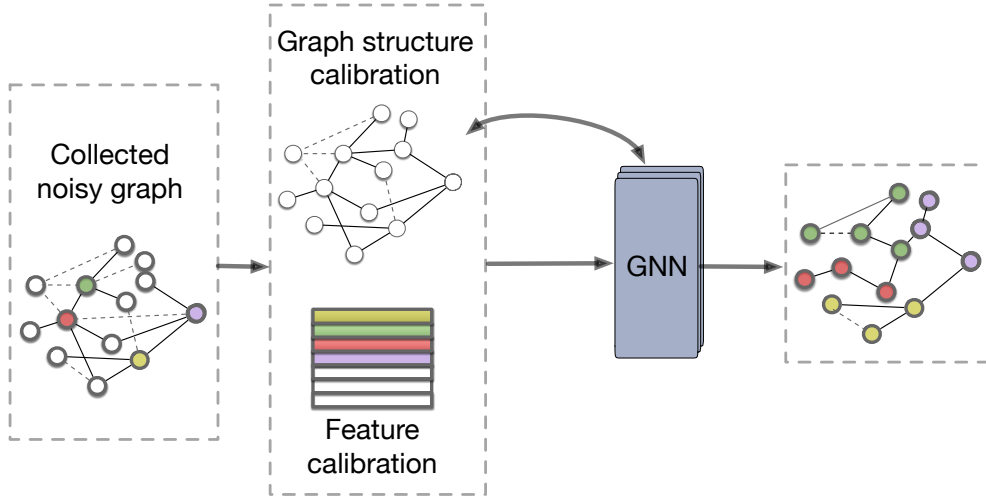


Fig. 2: Illustration of the private GNN training at the data curator side. It consists of two components: 1) a feature denoising component to reduce the effect of non-smooth aggregation; it is carried out before the training process; 2) a graph structure denoising component to enforce the graph structure’s sparseness; this process is jointly optimized with the model training.




Fig. 3: Illustration of multi-bit mechanism. The encoder encodes the multi-dimensional features $x_i \in \mathcal{R}^{|D|}$ into $\hat{x}_i \in \{-1, 0, 1\}^{|D|}$, with a probability defined by the privacy budget. The rectifier is to calibrate the encoded vector \hat{x}_i to ensure the outcome \tilde{x}_i is statistically unbiased.

as these tasks all share the same graph message passing architecture. Unfortunately, the collected graph is noisy; aggregating and propagating the noisy information leads to over-fitting which degenerates the generalization ability.

Why does the performance degrade? Although the randomized collection of the adjacency lists and feature vectors satisfies our privacy goal, the reconstructed graph does not reflect the original decentralized social graph well. From the perspective of the graph structure, it tends to return a much denser graph than the original one. For example, a citation network Cora [29] has $|\mathcal{V}| = 2708$ nodes, and the average node degree is 3.89. Mathematically, the randomized mechanism \mathcal{M}_a introduces $p \times |\mathcal{V}|^2 = 6681$ edge flipping with $\epsilon_a = 7$ in expectation. We empirically analyze the average node degree of the reconstructed graph in the setting of $\epsilon_a = 7$; it is 6.35 by averaging the results of 5 repeats. Concretely, it introduces 6634 edges that may be task-irrelevant, averaged over 5 repeats. A node might be linked to nodes with task-specific “noisy” edges due to the randomized flipping. Aggregating messages from these nodes would compromise the quality of the node embedding and lead to undesirable predictions in the downstream tasks.

To visualize the performance degradation, we evaluated the task of document classification with different values of privacy budgets on Cora [29] and CiteSeer [30] respectively. As observations are similar in other values of ϵ_x , we present

the results of $\epsilon_x = 1$ while changing the value of ϵ_a in Table 2. We observed that under the randomized mechanisms, the performance of GNNs is non-significant. In what follows, we are interested to exploit the intrinsic properties of the data to boost the classification accuracy with further gains when the graph is collected satisfying given privacy goals.

TABLE 2: Classification accuracy (%) on Cora and CiteSeer, with $\epsilon_x = 1$. The performance are insignificant in various edge privacy budgets.

ϵ_a	7.0	7.3	7.5	7.7	7.9
CORA	50.9 ± 3.8	54.9 ± 3.7	57.2 ± 2.5	61.4 ± 2.6	63.7 ± 2.4
CITESEER	35.0 ± 1.8	39.1 ± 1.3	42.3 ± 1.5	44.5 ± 1.2	45.7 ± 1.8

Graph Structure Denoising. The randomized flipping introduces “noisy” edges that can degrade the generalization performance of the GNNs. These edges tend to connect nodes within different communities or with different labels, and they should be pruned for better learning performance. According to the analysis mentioned above, the randomized flipping increases the density of the graph. Therefore, we propose to calibrate the collected noisy graph by encouraging the sparseness of the graph structure. In particular, we calibrate the graph structure by minimizing the l_1 norm of the calibrated adjacency matrix, denoted as $\|\mathbf{A}^c\|_1$. Therefore, the graph structure calibration process can be formulated as an optimization problem, shown in Eq. 10.

$$\min_{\mathbf{A}^c} \|\tilde{\mathbf{A}} - \mathbf{A}^c\|_{\mathbf{F}}^2 + \lambda \|\mathbf{A}^c\|_1, \quad (10)$$

where \mathbf{A}^c represented the adjacency matrix after calibration, and λ control the associated calibration level. The first term is to ensure the calibrated matrix to be close to the collected graph topology. Concretely, the distance of the calibrated matrix and the collected matrix is measured by the Frobenius norm. The Frobenius norm of a matrix \mathbf{A} is defined by $\|\mathbf{A}\|_{\mathbf{F}}^2 = \sum a_{i,j}^2$. We are aware of the drawbacks of reusing notations. \mathbf{A} in the definition of the Frobenius norm represents any matrix for simplicity.

Node Feature Vector Denoising. In GNNs, it computes the new representation of a node by aggregating and propagating information from its neighbors [5]. The learned representations of connected nodes tend to be similar. Stated differently, for message passing graph neural network to work, a certain assumption, called the *smoothness assumption*, has to hold. This reflects the real-world phenomenon on graphs from various domains. For example, two connected users in a social graph are likely to share similar features, fulfilling the property of *feature smoothness*. However, in the noisy graphs, the GNNs may result in degenerated node representations due to the non-smooth features aggregating from the neighbors, leading to performance degradation of the learned GNN.

To address the above issue, we leverage a feature smoothing component that applies a mean aggregator [15] to enhance the node features – denoising via mean aggregation of the node features from the neighbors. Specifically, instead of using the rectified features $\tilde{\mathbf{X}}$, we refine the features of each node by averaging the feature vector from their neighbors within l -hop. Formally, the process of feature smoothing within 1-hop can be formulated as:

$$\mathbf{x}_i^c = \sum_{v_j \in \mathcal{N}(v_i)} \frac{\tilde{\mathbf{x}}_j}{|\mathcal{N}(v_i)||\mathcal{N}(v_j)|}. \quad (11)$$

Our feature smoothing component is executed with l_x times, which is data-driven and needs to be tuned to avoid over-smoothing problem.

These denoising processes, including graph-structure denoising and feature vector denoising, are designed as calibration steps for better learning of the GNNs. According to the transformation invariance of LDP [27], these processes would not affect the privacy guarantee. Stated differently, the output of these processes is still noisy and does not reflect the private data portion of each user. Nevertheless, the feature smoothing operation is designed to reduce the noise effect that may degrade the generalization capability of the GNN.

Model Training. To evaluate the effectiveness of *Solitude* for privacy-preserving graph learning tasks, we instantiate the graph learning task with node classification. We denote the target classifier as $f(\tilde{\mathbf{x}}) = \arg \max_y p(y|\tilde{\mathbf{x}})$, where $p(y|\tilde{\mathbf{x}}) = g(\mathbf{A}^c; \mathbf{X}^c, \theta)$. More specifically, g represents the GNN model and θ are the learned model parameters given the obfuscated graph $\tilde{\mathcal{G}} = (\tilde{\mathbf{A}}, \tilde{\mathbf{X}})$ and the denoising mechanisms. Note that the generalization capability of the learned model is measured by the prediction accuracy over an obfuscated/noisy test set that is not seen during training. To further boost the prediction accuracy, we further incorporate a label smoothing component, which is inspired by the principles of label propagation algorithms [21], [22] — node labels are propagated and aggregated along edges in the graph. Formally, the label smoothing component within 1-hop is formulated as:

$$p(y^c|\tilde{\mathbf{x}}) = \sum_{v_j \in \mathcal{N}(v_i)} \frac{p(y|\tilde{\mathbf{x}})}{|\mathcal{N}(v_i)||\mathcal{N}(v_j)|}. \quad (12)$$

The label smoothing component can be carried out with l_y times, similar to the feature smoothing component. In Sec. 4, we will show that for node classification tasks, improving

privacy-utility guarantees needs more labeled samples for training.

In general, the model parameters are obtained by minimizing the cross-entropy error over all labeled samples:

$$\mathcal{L}_{\text{GNN}}(\mathbf{A}^c, \mathbf{X}^c, \theta) = - \sum_{v_i \in \mathcal{V}_l} \sum_{k=1}^K y_{ik} \ln p(y_k^c|\tilde{\mathbf{x}}), \quad (13)$$

where \mathcal{V}_l is the set of node indices that have labels, and K is the number of classes/labels.

Intuitively, the denoising processes can be regarded as data preprocessing for model training. In other words, at the data curator side, *Solitude* can first proceed with the mechanism for feature vector denoising and then optimize Eq. 10 to obtain a “denoised” graph structure. After that, the “denoised” graph is treated as the input of the GNN model. By solving Eq. 13, the model parameters θ can be obtained. Note that our ultimate goal is to obtain a GNN model with optimum generalization capability measured by the prediction accuracy on the test set. From this perspective, we can treat the process of graph structure denoising as a form of regularization, which is aligned with the regularization technique that prevents neural networks from overfitting [31]. Accordingly, the loss function for model training can be reformulated as:

$$\min_{\mathbf{A}^c, \theta} \mathcal{L}_{\text{GNN}}(\mathbf{A}^c, \mathbf{X}^c, \theta) + \lambda_1 \|\tilde{\mathbf{A}} - \mathbf{A}^c\|_{\mathbf{F}}^2 + \lambda_2 \|\mathbf{A}^c\|_1, \quad (14)$$

where λ_1 and λ_2 control the regularization ratio. To solve Eq. 14, an alternating optimization scheme with Adam is used to iteratively update θ and \mathbf{A}^c . As there is no parameter learning during feature denoising, the feature smoothing component will proceed before model training.

Discussions. In addition to introducing strong baselines for evaluating future improvements to private learning on graphs, our work suggests several open problems and directions for future work:

In this work, the notion of edge-LDP is built upon the randomized mechanism for the neighbor lists represented by a binary vector. This assumption implies that every user and the data curator know how many users and their indexes. We leave the problem of designing new mechanisms that satisfy edge-LDP when the users and the data curator may not know the entire set of the users as future work.

As illustrated in Sec. 2.3, there are two variants of LDP when applying LDP to graph data: node-LDP and edge-LDP, both of which offer different kinds of privacy protection. In this work, we only consider edge-LDP and LDP for node feature vectors. We believe that private learning on graphs under the notion of node-LDP is a promising direction as well.

Differentially private transfer learning has been studied in prior work and has shown to be a natural candidate for privacy-preserving machine learning in various domains [10]. As illustrated in the work of *Florian et al.*, [11], the heuristic rule “better models transfer better” also holds with differential privacy. Therefore, differentially private graph learning with access to public data from a similar domain may be a feasible solution to improve the privacy-utility guarantees.

TABLE 3: Classification accuracy (%) with $\epsilon_x = 1$, on various values of ϵ_a .

DATASET		ϵ_a	7.0	7.3	7.5	7.7	7.9	8.0
CORA	GRAPHSAGE	BASE	50.9 ± 3.8	54.9 ± 3.7	57.2 ± 2.5	61.4 ± 2.6	63.7 ± 2.4	63.2 ± 3.2
		LPGNN	62.5 ± 1.6	69.2 ± 1.6	72.9 ± 2.3	74.7 ± 0.9	75.3 ± 1.2	76.4 ± 0.8
		<i>Solitude</i>	66.4 ± 1.7	72.4 ± 1.2	75.8 ± 0.9	76.7 ± 1.1	77.9 ± 0.6	79.0 ± 0.6
	GCN	BASE	60.8 ± 2.0	63.1 ± 3.4	64.8 ± 3.7	66.9 ± 4.0	68.1 ± 4.0	68.6 ± 3.8
		LPGNN	67.9 ± 1.2	69.4 ± 1.4	73.7 ± 1.8	74.4 ± 0.5	75.7 ± 1.0	76.3 ± 0.9
		<i>Solitude</i>	68.4 ± 2.3	72.6 ± 1.4	75.2 ± 1.2	76.1 ± 0.5	77.3 ± 0.9	77.8 ± 0.5
CITeseER	GRAPHSAGE	BASE	35.0 ± 1.8	39.1 ± 1.3	42.3 ± 1.5	44.5 ± 1.2	45.7 ± 1.8	44.5 ± 1.9
		LPGNN	46.8 ± 2.0	49.7 ± 0.7	50.7 ± 1.4	51.2 ± 1.4	53.3 ± 0.9	54.5 ± 1.2
		<i>Solitude</i>	50.1 ± 1.3	52.7 ± 1.0	54.0 ± 1.0	55.6 ± 0.8	56.3 ± 1.0	57.0 ± 1.4
	GCN	BASE	39.7 ± 3.0	45.3 ± 1.8	48.3 ± 2.9	49.6 ± 1.3	52.5 ± 2.1	52.8 ± 2.5
		LPGNN	47.8 ± 1.1	50.1 ± 1.8	52.3 ± 1.2	53.7 ± 1.3	55.6 ± 1.3	55.6 ± 1.5
		<i>Solitude</i>	49.8 ± 1.4	53.4 ± 1.0	55.1 ± 0.8	56.4 ± 1.5	58.2 ± 0.7	58.0 ± 2.5
LASTFM		ϵ_a	7.7	7.9	8.0	8.3	8.5	8.7
	GRAPHSAGE	BASE	62.6 ± 2.4	64.4 ± 1.8	66.5 ± 1.9	71.2 ± 1.9	72.0 ± 2.0	75.5 ± 1.4
		LPGNN	62.6 ± 2.4	65.3 ± 5.8	68.3 ± 2.4	73.3 ± 3.2	74.6 ± 3.0	76.5 ± 3.1
		<i>Solitude</i>	65.5 ± 7.1	69.4 ± 5.4	71.6 ± 2.4	77.2 ± 1.7	77.2 ± 1.4	79.0 ± 1.5
	GCN	BASE	60.5 ± 2.5	63.2 ± 2.1	63.8 ± 2.0	66.9 ± 1.1	68.5 ± 1.5	69.6 ± 1.6
		LPGNN	62.3 ± 7.2	68.2 ± 0.8	68.3 ± 3.0	73.3 ± 2.4	76.9 ± 2.9	79.5 ± 0.8
<i>Solitude</i>		60.9 ± 2.6	68.9 ± 2.1	68.0 ± 4.7	75.1 ± 1.7	77.5 ± 1.9	78.9 ± 1.3	

4 EVALUATING PRIVACY-PRESERVING GNNS

4.1 Datasets and Experimental Setup

TABLE 4: Statistical Description of Used Datasets.

DATASET	#NODES	#EDGES	#CLASSES	#FEATURES	AVG. DEGREE
CORA	2, 708	5, 278	7	1, 433	3.90
CITeseER	2, 110	3, 668	6	3, 703	2.74
LASTFM	7, 083	25, 814	10	7, 842	7.29

Datasets. We conducted experiments over 3 datasets, falling into two categories: citation networks and social networks. For citation networks, we use Cora [29] and CiteSeer [30], both of which are benchmarking datasets for node classification. In these two datasets, nodes represent scientific publications, and edges correspond to the citation links. These two datasets contain bag-of-words feature vectors for each publication; each publication has a class label. For the real-world social networks, we use LastFM [32], which is a dataset collected from a music streaming service, in which nodes are users from Asian countries, and links represent friendships. Its task is to predict the home country of a user given the artists liked by them. As this dataset was highly imbalanced, for fair comparisons, we limit the classes to the top-10 ones with the most samples, as was done in [15]. The detailed statistics of the used datasets are described in Table 4.

Baselines. Our framework is the first effort towards private learning over decentralized network graphs to the best of our knowledge. We compare *Solitude* against the **Base** methods including GCN [12] and GraphSage [5]. GCN and GraphSage are two representative graph neural networks. By evaluating our approach with various GNN backbones,

we can show that *Solitude* is able to preserve edge privacy and node feature privacy by incorporating with any GNNs. In particular, we use the same randomized mechanisms to obfuscate the feature vectors and adjacency lists. The main difference is that the base methods are directly trained over the noisy graphs without any calibration process. For better effectiveness demonstration, we also use an alternative baseline — LPGNN [15]. This method assumes that the data curator can access the global topology, and it was proposed to preserve the privacy of the node features. Though the setting of LPGNN and ours are different, to be comparable, we adapt LPGNN into our setting by randomizing the adjacency lists with our proposed mechanism and do not consider the labels of the training samples as the private information of the users. We argue that once the model parameters are achieved, the data curator can obtain the labels for any node by treating them as the model input as long as the model is well trained.

Experimental Setup. We follow the data preprocessing as reported in [15]. Specifically, for all datasets, we randomly split each dataset into three portions: 50% for training, 25% for validation, and 25% for the test. For the datasets with node features, including Cora, CiteSeer, and LastFM, the randomized mechanisms for LDP are applied to the node features and the adjacency lists of all training, validation, and test sets. All the GNN models, including the base methods and the backbone models of LPGNN and *Solitude*, consist of two graph convolution layers — each of which has a hidden dimension of size 16 — and a SeLU activation function [33] followed by dropout.

As for the evaluation metrics, we employ the standard metric — classification accuracy on the test set (or called prediction accuracy) under various privacy budgets — to evaluate the generalization capability of the learning model.

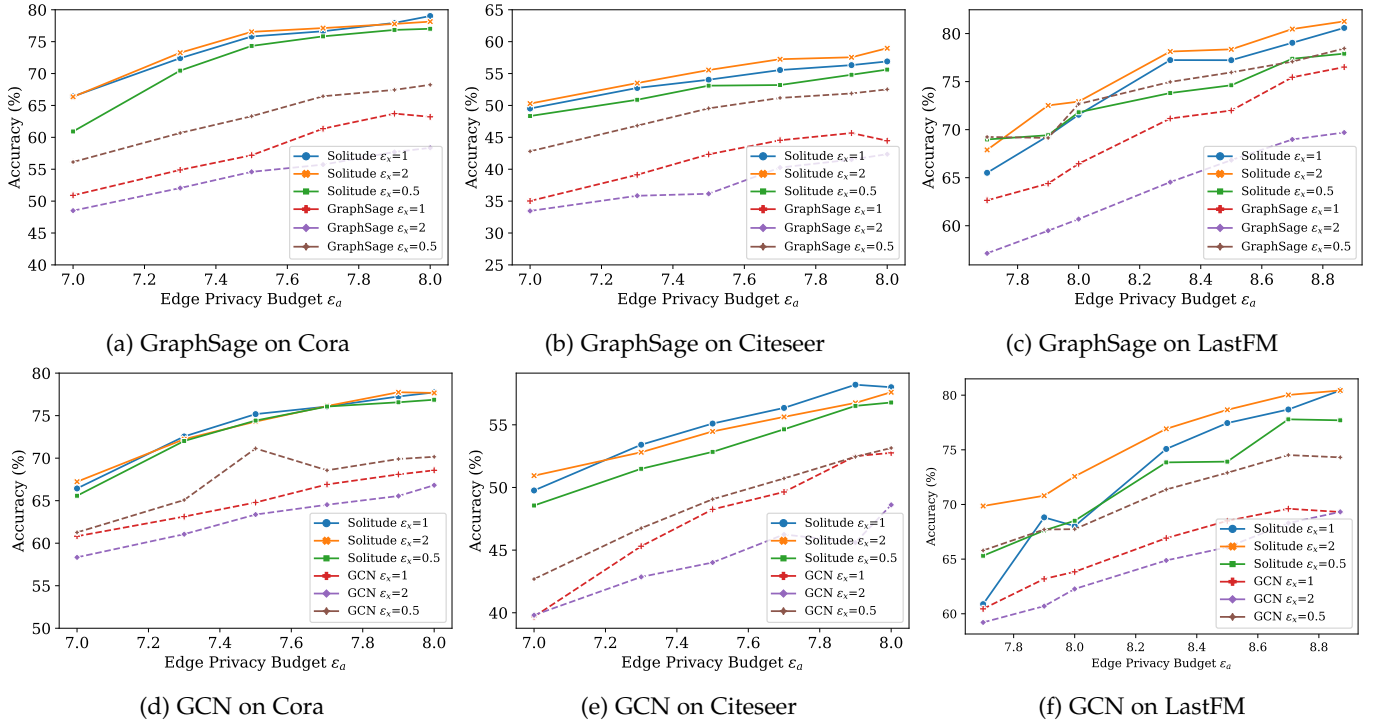


Fig. 4: Classification accuracy with different GNN architectures. The horizontal axes are different values of edge privacy budgets. The vertical axes are the prediction accuracy of the test set. *Solitude* obtains significance gains on Cora, Citeseer, and LastFM on various values of ϵ_x and ϵ_a .

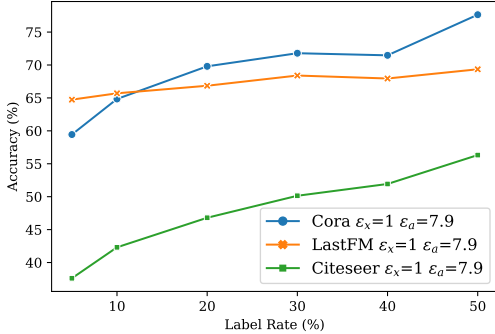


Fig. 5: Classification accuracy under various label rates. It shows that improving privacy-utility guarantees needs more labeled samples.

Unless otherwise stated, all experiments are run 5 times to ensure statistical significance. Specifically, we report the mean and standard deviation values over 5 runs.

Parameter settings. For general hyperparameters, we applied a grid search to find the best choices: the learning rate, dropout rate, and weight decay were tuned among $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}\}$, the feature smoothing steps l_x and label smoothing steps l_y were searched from $\{0, 2, 4, 8\}$, and the coefficients of λ_1 and λ_2 were searched in $\{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\}$. We used the Xavier initializer [34] and the Adam SGD optimizer [35] for all models. In addition, the maximum epoch was set as 500. For all ϵ_x and ϵ_a pairs, we traversed all the parameters to get its optimal performance in the experimental environment. Without specification, we report the results under the hyperparameters with the best performance overall. More specifically,

we use the best learning rate, weight decay, and dropout for every ϵ_x and ϵ_a pairs. The selection rationality of the privacy budgets is discussed below.

4.2 Experimental Results

We first investigate how *Solitude* performs under varying feature and edge privacy budgets. In particular, the privacy budget for the node features varies within $\{0.5, 1, 2\}$. The maximum value of ϵ_x is selected based on the proposition in [15], which indicates that in the high-privacy regime $\epsilon_x \leq 2.18$, the multi-bit mechanism perturbs at least one random dimension. Similarly, for randomized response flipping at least one edge of a node, the probability satisfies $p = \frac{1}{1+e^{\epsilon_a}} \geq \frac{1}{|\mathbf{V}|}$, where $|\mathbf{V}|$ denotes the number of nodes in the graph. Then, we can obtain $\epsilon_a \leq \ln(N - 1)$. Therefore, we varies the edge privacy budgets within $\{7.0, 7.3, 7.5, 7.7, 7.9, 8.0\}$ for two citation networks, and $\{7.7, 7.9, 8.0, 8.3, 8.5, 8.7, 8.9\}$ for LastFM. We first fix the privacy budget for the features and compare the performance under various edge privacy budgets. Table 3 reports the classification accuracy of different methods when $\epsilon_x = 1$.

We can observe that our proposed framework consistently achieves the best performance in most cases. Concretely, *Solitude* has an improvement by up to 9.3% on LastFM as compared to the base methods. As for the two citation networks, our framework achieves significant performance gain ranging from around 7.6% to 18.6% on Cora, and from 5.2% to 15.1% on Citeseer comparing to the base methods (see Table 3). The difference in performance gains between LastFM and two citation networks implies

that lower degree networks enjoy more benefits from our framework. The reason is that the randomized response has a higher impact on the graphs with more sparseness. Though LPGNN achieves comparable performance gains by comparing with the base methods, it still falls behind our framework in most cases. Note that LPGNN was designed specifically to protect the privacy of node features, while our framework can protect the node features and graph structure information simultaneously. The performance gain achieved by *Solitude* upon LPGNN indicates that the denoising component for graph structure is indeed effective in reducing the effects of introduced noise due to the randomized flipping.

To further examine the effects of introduced noise in the feature vectors, we evaluate the performance of our framework with various privacy budgets for the node features and the adjacency lists. Fig. 4 shows the classification accuracy with different GNN architectures across three datasets. We see *Solitude* achieves better performance in the lower-privacy regime (with higher privacy budgets). This can be explained by the fact that the predictions become more accurate due to an increase in the privacy budget. Interestingly, we also observe that *Solitude* achieves better performance gains with higher privacy budgets for node features when the edge privacy budget is fixed. The maximum performance gain is different across datasets and privacy budgets. This result indicates that our denoising mechanisms effectively mitigate the overfitting issue caused by the introduced noise for privacy protection, improving the model utility.

Essentially, for GNNs to work, the *smoothness assumption* has to hold. As the crux of GNNs for node classification is to propagate features and labels throughout the network graph. Inspired by this, we are interested in investigating how the label rates affect the model utility with local differential privacy. In particular, label rate denotes the number of labeled nodes for training divided by the total number of nodes in each dataset. Fig. 5 depicts the classification accuracy under various label rates. It shows that the prediction accuracy consistently increases with the increase of the label rates over all datasets. We can conclude that improving privacy-utility guarantees needs more labeled nodes for training.

5 RELATED WORK

Privacy-Preserving Analysis of Graph Statistics. Differential privacy (DP) has emerged as a de-facto standard for privacy guarantees [7], [25]. Most existing work focuses on centralized differential privacy [36], [37], and they are specialized for analyzing the graph statistics, such as degree distribution estimation [27], subgraph counts [36]. LDP is a special case of differential privacy in the local model. Each participant obfuscates their data portion locally and then sends the obfuscated data to a data curator (possibly a malicious third-party). Prior works on graph data with LDP mainly focus on the estimations of graph statistics, such as clustering coefficient estimation [9], heavy hitter estimation [13], and frequency estimation [14].

Among others, some works are specialized for subgraph counts [8], [38], such as triangles, 3-hop paths, and k -

cliques. Specifically, Sun *et al.* [38] proposes a multi-phase framework under decentralized differential privacy, which assumes that each user/data holder is aware of not only her connections but also a broader subgraph in her local neighborhood. Subsequently, Imola *et al.* [8] propose new algorithms for triangle and k -star counts, assuming that each user can only access her connections. LDPGen [17] was proposed to generate synthetic graphs in the setting where a data curator collects subgraphs from the participants under the notion of edge differential privacy.

Privacy Leakage on GNNs. There exists some prior work on the privacy implication of training GNNs on graphs [39]–[41]. Duddu *et al.* [40] analyzed the privacy risks of graph embedding with various attacks, including membership inference attacks, graph reconstruction attacks, and attribute inference attacks. Later, Olatunji *et al.* [39] proposed membership inference attacks under restricted scenarios, by 0-hop query and 2-hop query to the victim models. He *et al.* [41] showed that given various background knowledge, the outputs of a GNN indeed could be utilized to infer the rich information about the graph structure used to train the model. They coined the attack “link stealing attack.”

Privacy-Preserving Graph Learning with GNNs. There is a research line that attempts to address privacy in graph learning with federated learning [42], [43] and split learning, which is orthogonal to our setting. The closest to ours is LPGNN [15], wherein a set of mechanisms are proposed to protect the privacy of node features. Specifically, LPGNN assumes that a central server holds the global graph topology, and the server is allowed to collect the node features satisfying local differential privacy. It strives to preserve the privacy of node features while maintaining the generalization capability of the learned GNNs. However, the assumption of holding the global topology in the central server hinders its practicality in many real-world applications, where the privacy of the graph topology is of paramount importance. Here we fill this gap and show the effectiveness of our privacy-preserving learning framework based on graph neural networks, with local differential privacy guarantees.

6 CONCLUDING REMARKS

In this paper, we propose a new privacy-preserving learning framework for decentralized network graphs based on graph neural networks, called *Solitude*. It can simultaneously preserve edge privacy and node feature privacy for every user and seamlessly incorporate with any GNN architectures, such as GCN and GraphSage, with privacy-utility guarantees. The key of *Solitude* is a set of new mechanisms that can calibrate the introduced noise in the decentralized graph to mitigate the overfitting problem while learning over noisy graphs. Theoretical analysis and extensive experiments on benchmarks have demonstrated the rationality and effectiveness of our proposed mechanisms as compared to the alternative baselines.

Due to the inherent limitations of current differential privacy formalisms for graph learning, our design needs a certain amount of privacy budget to have a comparative learning performance. We hope our study pushes forward

the progress on privacy-preserving GNNs under local differential privacy. There are many possible directions for future work, such as establishing novel formalisms of local differential privacy, i.e., edge differential privacy, specifically for graph learning. We leave this in our future work.

7 ACKNOWLEDGMENTS

Wanyu Lin has been supported by the Research Fund at The Hong Kong Polytechnic University P0035763 and P0042687. Cong Wang has been supported by the Research Grants Council of Hong Kong under Grant CityU 11217620, N_CityU139/21, and RFS2122-1S04. The assistance provided by Yuteng Sun from the Chinese University of Hong Kong was greatly appreciated.

REFERENCES

- [1] K. Rusek, J. Suárez-Varela, P. Almasan, P. Barlet-Ros, and A. Cabellos-Aparicio, "RouteNet: Leveraging Graph Neural Networks for Network Modeling and Optimization in SDN," *IEEE Journal on Selected Areas in Communications*, vol. 38, no. 10, pp. 2260–2270, 2020.
- [2] M. Zitnik, M. Agrawal, and J. Leskovec, "Modeling Polypharmacy Side Effects with Graph Convolutional Networks," *Bioinformatics*, vol. 34, no. 13, pp. i457–i466, 2018.
- [3] M. Zitnik and J. Leskovec, "Predicting Multicellular Function Through Multi-Layer Tissue Networks," *Bioinformatics*, vol. 33, no. 14, pp. i190–i198, 2017.
- [4] J. Zhou, Z. Xu, A. M. Rush, and M. Yu, "Automating Botnet Detection with Graph Neural Networks," in *AutoML for Networking and Systems Workshop of MLSys 2020 Conference*, 2020.
- [5] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive Representation Learning on Large Graphs," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [6] F. Wu, Y. Long, C. Zhang, and B. Li, "LINKTELLER: Recovering Private Edges from Graph Neural Networks via Influence Analysis," in *Proc. of the Symposium on Security and Privacy*, 2021.
- [7] J. C. Duchi, M. I. Jordan, and M. J. Wainwright, "Local Privacy and Statistical Minimax Rates," in *Annual Symposium on Foundations of Computer Science*. IEEE, 2013.
- [8] J. Imola, T. Murakami, and K. Chaudhuri, "Locally Differentially Private Analysis of Graph Statistics," in *Proc. {USENIX} Security Symposium ({USENIX} Security)*, 2021.
- [9] Q. Ye, H. Hu, M. H. Au, X. Meng, and X. Xiao, "LF-GDPR: A Framework for Estimating Graph Metrics with Local Differential Privacy," *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- [10] M. Abadi, A. Chu, I. Goodfellow, H. B. McMahan, I. Mironov, K. Talwar, and L. Zhang, "Deep Learning with Differential Privacy," in *Proc. the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2016.
- [11] T. Florian and B. Dan, "Differentially Private Learning Needs Better Features (or Much More Data)," in *Proc. International Conference on Learning Representations (ICLR)*, 2021.
- [12] T. N. Kipf and M. Welling, "Semi-Supervised Classification with Graph Convolutional Networks," in *Proc. International Conference on Machine Learning (ICML)*, 2017.
- [13] Z. Qin, Y. Yang, T. Yu, I. Khalil, X. Xiao, and K. Ren, "Heavy Hitter Estimation over Set-Valued Data with Local Differential Privacy," in *Proc. the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2016.
- [14] T. Wang, J. Blocki, N. Li, and S. Jha, "Locally Differentially Private Protocols for Frequency Estimation," in *Proc. {USENIX} Security Symposium ({USENIX} Security)*, 2017.
- [15] S. Sajadmanesh and D. Gatica-Perez, "Locally Private Graph Neural Networks," in *Proc. the ACM Conference on Computer and Communications Security (CCS)*, 2021.
- [16] D. Konforty, Y. Adam, D. Estrada, and L. G. Meredith, "Synereo: the Decentralized and Distributed Social Network," 2015.
- [17] Z. Qin, T. Yu, Y. Yang, I. Khalil, X. Xiao, and K. Ren, "Generating Synthetic Decentralized Social Graphs with Local Differential Privacy," in *Proc. the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2017.
- [18] C. Dwork, A. Roth *et al.*, "The Algorithmic Foundations of Differential Privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, no. 3-4, pp. 211–407, 2014.
- [19] B. Ding, J. Kulkarni, and S. Yekhanin, "Collecting Telemetry Data Privately," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [20] D. Zhu, Z. Zhang, P. Cui, and W. Zhu, "Robust Graph Convolutional Networks against Adversarial Attacks," in *Proc. the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, 2019.
- [21] Y. Bengio, O. Delalleau, and N. Le Roux, "Label Propagation and Quadratic Criterion," 2006.
- [22] D. Zhou, O. Bousquet, T. N. Lal, J. Weston, and B. Schölkopf, "Learning with Local and Global Consistency," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2004.
- [23] Tencent, "Wechat: Connecting a Billion People with Calls, Chats, and More," <https://www.wechat.com>.
- [24] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and S. Y. Philip, "A Comprehensive Survey on Graph Neural Networks," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 1, pp. 4–24, 2020.
- [25] S. Raskhodnikova, A. Smith, H. K. Lee, K. Nissim, and S. P. Kasiviswanathan, "What Can We Learn Privately," in *Proc. the Symposium on Foundations of Computer Science*. IEEE, 2008.
- [26] J. Blocki, A. Blum, A. Datta, and O. Sheffet, "The Johnson-Lindenstrauss Transform Itself Preserves Differential Privacy," in *Proc. the Symposium on Foundations of Computer Science*. IEEE, 2012.
- [27] W.-Y. Day, N. Li, and M. Lyu, "Publishing Graph Degree Distribution with Node Differential Privacy," in *Proc. the International Conference on Management of Data (SIGMOD)*, 2016.
- [28] J. You, J. M. Gomes-Selman, R. Ying, and J. Leskovec, "Identity-Aware Graph Neural Networks," in *Proc. the AAAI Conference on Artificial Intelligence*, 2021.
- [29] A. K. McCallum, K. Nigam, J. Rennie, and K. Seymore, "Automating the Construction of Internet Portals with Machine Learning," *Information Retrieval*, vol. 3, no. 2, pp. 127–163, 2000.
- [30] C. L. Giles, K. D. Bollacker, and S. Lawrence, "Citeseer: An Automatic Citation Indexing System," in *Proc. the ACM Conference on Digital Libraries*, 1998.
- [31] J. Kukavcka, V. Golkov, and D. Cremers, "Regularization for Deep Learning: A taxonomy," *arXiv preprint arXiv:1710.10686*, 2017.
- [32] B. Rozenberczki and R. Sarkar, "Characteristic Functions on Graphs: Birds of a Feather, from Statistical Descriptors to Parametric Models," in *Proc. the ACM International Conference on Information & Knowledge Management*, 2020.
- [33] G. Klambauer, T. Unterthiner, A. Mayr, and S. Hochreiter, "Self-Normalizing Neural Networks," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [34] X. Glorot and Y. Bengio, "Understanding the Difficulty of Training Deep Feedforward Neural Networks," in *Proc. International Conference on Artificial Intelligence and Statistics*, 2010, pp. 249–256.
- [35] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," in *Proc. International Conference on Learning Representations (ICLR)*, 2015.
- [36] V. Karwa, S. Raskhodnikova, A. Smith, and G. Yaroslavtsev, "Private Analysis of Graph Structure," *Proc. the VLDB Endowment*, vol. 4, no. 11, pp. 1146–1157, 2011.
- [37] Y. Wang, X. Wu, and L. Wu, "Differential Privacy Preserving Spectral Graph Analysis," in *Proc. Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 2013.
- [38] H. Sun, X. Xiao, I. Khalil, Y. Yang, Z. Qin, H. Wang, and T. Yu, "Analyzing Subgraph Statistics from Extended Local Views with Decentralized Differential Privacy," in *Proc. the ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2019.
- [39] I. E. Olatunji, W. Nejdl, and M. Khosla, "Membership Inference Attack on Graph Neural Networks," in *2021 Third IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications (TPS-ISA)*. IEEE, 2021, pp. 11–20.
- [40] V. Duddu, A. Boutet, and V. Shejwalkar, "Quantifying Privacy Leakage in Graph Embedding," in *Mobiquitous 2020-17th EAI International Conference on Mobile and Ubiquitous Systems: Computing, Networking and Services*, 2020, pp. 1–11.
- [41] X. He, J. Jia, M. Backes, N. Z. Gong, and Y. Zhang, "Stealing Links from Graph Neural Networks," in *30th {USENIX} Security Symposium ({USENIX} Security 21)*, 2021.
- [42] C. Meng, S. Rambhatla, and Y. Liu, "Cross-Node Federated Graph Neural Network for Spatio-Temporal Data Modeling," in *Proc. the*

ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD), 2021.

- [43] C. He, K. Balasubramanian, E. Ceyani, Y. Rong, P. Zhao, J. Huang, M. Annavaram, and S. Avestimehr, "FedGraphNN: A Federated Learning System and Benchmark for Graph Neural Networks," in *Workshop on Graph Neural Networks and Systems*, 2021.



Wanyu Lin received her Ph.D. degree from the Department of Electrical and Computer Engineering at the University of Toronto. She received her B.Engr. degree from the School of Electronic Information and Communications, Huazhong University of Science and Technology, China and her MPhil. degree from the Department of Computing, The Hong Kong Polytechnic University. Her research interests include graph machine learning, trustworthy machine learning, data privacy, and model interpretability.

She has served as associate editor for IEEE Transactions on Neural Networks and Learning Systems (TNNLS). She is a member of IEEE.



Baochun Li received his B.Engr. degree from the Department of Computer Science and Technology, Tsinghua University, China, in 1995 and his M.S. and Ph.D. degrees from the Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, in 1997 and 2000. Since 2000, he has been with the Department of Electrical and Computer Engineering at the University of Toronto, where he is currently a Professor. He holds the Bell Canada Endowed Chair in Computer Engineering since August 2005.

His research interests include cloud computing, distributed systems, datacenter networking, and wireless systems. He was the recipient of the IEEE Communications Society Leonard G. Abraham Award in the Field of Communications Systems in 2000. In 2009, he was a recipient of the Multimedia Communications Best Paper Award from the IEEE Communications Society, and a recipient of the University of Toronto McLean Award. He is a member of ACM and a Fellow of IEEE.



Cong Wang is a Professor in the Department of Computer Science, City University of Hong Kong. His research interests include data and network security, blockchain and decentralized applications, and privacy-enhancing technologies. He has been one of the Founding Members of the Young Academy of Sciences of Hong Kong since 2017, and has been conferred the RGC Research Fellow in 2021. He received the Outstanding Researcher Award in 2019, the Outstanding Supervisor Award in 2017 and the

President's Awards in 2019 and 2016, all from City University of Hong Kong. He is a co-recipient of the Best Paper Award of IEEE ICDCS 2020, ICPADS 2018, MSN 2015, the Best Student Paper Award of IEEE ICDCS 2017, and the IEEE INFOCOM Test of Time Paper Award 2020. His research has been supported by multiple government research fund agencies, including National Natural Science Foundation of China, Hong Kong Research Grants Council, and Hong Kong Innovation and Technology Commission. He has served as an associate editor for IEEE Transactions on Dependable and Secure Computing (TDSC), IEEE Transactions on Services Computing (TSC), IEEE Internet of Things Journal (IoT-J), IEEE Networking Letters, and The Journal of Blockchain Research, and TPC co-chairs for a number of IEEE conferences and workshops. He is a Fellow of IEEE and a member of ACM.