

Devil's Hand: Data Poisoning Attacks to Locally Private Graph Learning Protocols

Longzhu He
Beijing University of Posts and
Telecommunications
Beijing, China
helongzhu@bupt.edu.cn

Chaozhuo Li
Beijing University of Posts and
Telecommunications
Beijing, China
lichaozhuo@bupt.edu.cn

Peng Tang
Shandong University
Qingdao, China
tangpeng@sdu.edu.cn

Litian Zhang
Beihang University
Beijing, China
litianzhang@buaa.edu.cn

Sen Su*
Beijing University of Posts and
Telecommunications
Beijing, China
susen@bupt.edu.cn

ABSTRACT

Graph neural networks (GNNs) have achieved significant success in graph representation learning and have been applied to various domains. However, many real-world graphs contain sensitive personal information, such as user profiles in social networks, raising serious privacy concerns when graph learning is performed using GNNs. To address this issue, locally private graph learning protocols have gained considerable attention. These protocols leverage the privacy advantages of local differential privacy (LDP) and the effectiveness of GNN's message-passing in calibrating noisy data, offering strict privacy guarantees for users' local data while maintaining high utility (e.g., node classification accuracy) for graph learning. Despite these advantages, such protocols may be vulnerable to *data poisoning attacks*, a threat that has not been considered in previous research. Identifying and addressing these threats is crucial for ensuring the robustness and security of privacy-preserving graph learning frameworks. This work introduces the first data poisoning attack targeting locally private graph learning protocols. The attacker injects fake users into the protocol, manipulates these fake users to establish links with genuine users, and sends carefully crafted data to the server, ultimately compromising the utility of private graph learning. The effectiveness of the attack is demonstrated both theoretically and empirically. In addition, several defense strategies have also been explored, but their limited effectiveness highlights the need for more robust defenses.

KEYWORDS

Local Differential Privacy, Graph Learning, Data Poisoning Attack

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).
Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00
<https://doi.org/XXXXXXX.XXXXXXX>

ACM Reference Format:

Longzhu He, Chaozhuo Li, Peng Tang, Litian Zhang, and Sen Su. 2018. Devil's Hand: Data Poisoning Attacks to Locally Private Graph Learning Protocols. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 13 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Graph Neural Networks (GNNs) have achieved remarkable success in graph representation learning [14] and have been widely applied in various domains, including graph mining [27, 52], recommendation systems [53, 56], and bioinformatics [11, 61]. However, most real-world graphs contain sensitive personal information, such as users' profiles and comments on social networks. Training GNNs on such graph data can raise serious privacy concerns, as malicious attackers can potentially infer sensitive information, such as user attributes, by interacting with the GNN model [34, 50, 63]. Therefore, it is of particular significance to develop privacy-preserving GNN algorithms to safeguard user privacy.

Recently, locally private graph learning protocols have garnered significant attention from the security research community [15, 28, 31, 38, 40, 43, 58, 66], as illustrated in Fig. 1(a). In these protocols, each user first perturbs their origin data locally using local differential privacy (LDP) [54], typically through noise injection [8, 43, 48], and then transmits the perturbed data to an untrusted third-party server that could potentially compromise the user's privacy. The server performs private graph learning on the collected noisy data. Given the superiority of LDP in protecting data privacy and the effectiveness of GNN's multi-hop message passing [1, 10] in calibrating noisy data, locally private graph learning protocols offer strict privacy guarantees for the user data while maintaining high utility for graph learning tasks.

However, despite these advantages, these protocols can be vulnerable to *data poisoning attacks* due to their distributed setup, which has not been considered in previous studies on locally private graph learning. In distributed privacy scenarios, although an attacker cannot take control of genuine users or directly compromise the graph database, it is easy to register fake accounts and create social identities with profiles. Consider a threat scenario as shown in Fig. 1(b), where an attacker aims to degrade the graph

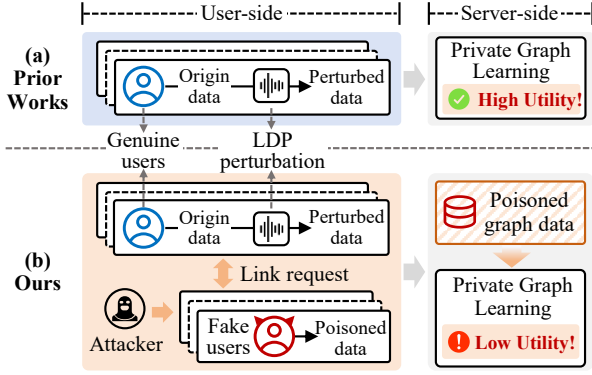


Figure 1: Comparison of prior works with ours. (a) Prior work involves a cloud server and multiple genuine users, where users’ sensitive data is perturbed with LDP to ensure privacy before being sent to the server for private graph learning. (b) In contrast, our work targets poisoning the private graph learning process. An attacker manipulates fake users, links them with genuine ones, and sends poisoned data to the server to compromise the utility of the learning process.

learning utility (e.g., node classification accuracy [20] or link prediction accuracy [59]) of the original graph by injecting fake users. Specifically, the attacker first manipulates multiple fake users to be injected into the protocol, then sends friendship requests to genuine users, and finally transmits carefully crafted poisoned data to the server. Social users seeking social influence [3, 36] are generally more inclined to accept friendship requests. As a result, some genuine users accept requests from attackers, allowing fake accounts to connect to genuine users and poison the original graph. When the GNN is trained on the poisoned graph data, malicious information spreads throughout the graph through message passing [1, 10], significantly damaging the utility of privacy-preserving graph learning. In summary, this paper proposes a data poisoning attack to achieve the aforementioned objectives.

Unlike previous data poisoning attacks against GNNs [45, 67–69], achieving the above attack faces the following unique challenges. ① *Strict data constraints*. Unlike traditional poisoning attacks that manipulate data without considering privacy concerns, this attack must be conducted within the LDP framework, which enforces specific rules for perturbing sensitive information. The challenge lies in balancing, on the one hand, the attacker’s need to carefully calibrate privacy-preserving perturbations to avoid detection and, on the other, the desire to maximize the compromise of private graph learning’s utility. This delicate balance makes the attack more difficult. ② *Limited background knowledge*. In this attack scenario, a realistic black-box setup is considered, where the attacker has no access to information about the target node or its neighbors, nor any details about the architecture or parameters of the GNN model within the protocol. ③ *Complexity of the protocol*. Poisoning locally private graph learning protocols requires considering various LDP mechanisms (e.g., the piecewise mechanism [38, 48], the multi-bit mechanism [31, 43], and the square wave mechanism [28, 29]) and different GNN models (e.g., graph convolutional networks [20], GraphSAGE [13], and graph attention networks [47]). This complexity adds another layer of difficulty in ensuring the effectiveness

and universality of the attack. In general, given these constraints, designing the data to be sent by the fake user to the server and determining how to connect with genuine nodes to maximize the attack’s effectiveness is a highly challenging task.

To address these challenges, we leverage a general framework for locally private graph learning protocols [28, 31, 38, 43] and identify potential attack surfaces by analyzing the phases most susceptible to manipulation. Focusing on the realism and universality of the attacks, we target the *data perturbation* phase of the protocol. When carrying out the attack, for each genuine target node, we model the attack to maximally deviate from its original embedding (server-side aggregated vectors) by crafting the features of the fake nodes connected to it. To maintain the stealthiness of the attack, it is crucial that the crafted malicious node features adhere to the specific output rules of the LDP mechanism [28, 38, 43]. With this limitation, we achieve the goal by selecting *extreme feature values*. Additionally, we further enhance the attack’s effectiveness by adjusting the links between the fake nodes to circumvent detection and increase disruption to the original graph.

The effectiveness of the attack is evaluated both theoretically and empirically. Theoretically, the error introduced by the poisoning attack on the embedding of target nodes is analyzed, and its impact on the entire graph is further assessed. Empirically, evaluations are conducted on six real-world datasets, applying three de facto LDP mechanisms [28, 38, 43] and three GNN models [13, 20, 47]. Our empirical results confirm the effectiveness of our attacks. In addition, several potential defenses are explored, including graph homophily analysis [4, 60], anomaly node detection [32], and others. However, the limited effectiveness of these defenses highlights the urgent need for more robust countermeasures against our attack. Our contributions are summarized as follows.

- To the best of our knowledge, this is the first study of data poisoning attacks on locally private graph learning protocols.
- We focus on the perturbation phase of the protocol and achieve a data poisoning attack, which aims to damage the learning utility.
- Through both theoretical analysis and empirical evaluation, it is demonstrated that our attacks can effectively degrade the utility of privacy-preserving graph learning, including tasks such as node classification and link prediction.
- Several defense strategies are also explored to mitigate these attacks, with empirical results emphasizing the urgent need for more robust defensive countermeasures.

Organization. The rest of this paper is organized as follows. Sec. 2 introduces the problem definition and provides essential background knowledge. Sec. 3 presents the threat model and a detailed explanation of our proposed attack approach. Theoretical analysis is provided in Sec. 4, while Sec. 5 empirically evaluates the effectiveness of the attack. Sec. 6 explores various defense countermeasures. Sec. 7 reviews related literature, and Sec. 8 concludes the paper.

2 PRELIMINARIES

Before introducing our poisoning algorithm, we first define the problem (Sec. 2.1) and outline the essential background on LDP (Sec. 2.2) and locally private graph learning protocols (Sec. 2.3).

2.1 Problem Definition

The node data privacy problem is considered in the context of GNN-based graph learning. Formally, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ represent a graph, where \mathcal{V} is the set of nodes and \mathcal{E} is the set of edges. $\mathbf{X} \in \mathbb{R}^{|\mathcal{V}| \times d}$ is the set of node feature vectors. Each node/user¹ $v \in \mathcal{V}$ has a d -dimensional feature vector $\mathbf{x}_v \in [\alpha, \beta]^d$, which contains sensitive information about the user. Uploading it directly to an untrusted server for GNN-based graph learning poses a significant privacy risk [34, 62, 63]. Locally private graph learning protocols [28, 31, 38, 43] aim to ensure node privacy while enabling effective graph learning through LDP [54]. This paper proposes poisoning a series of locally private graph learning protocols to expose their potential security vulnerabilities, thereby promoting more secure and robust privacy-preserving graph learning frameworks.

2.2 Local Differential Privacy

LDP is widely adopted in decentralized data collection and distribution scenarios [6, 16, 17, 48, 49, 54] due to its ability to protect individual privacy while enabling meaningful data analysis. It has seen widespread application in industry, including but not limited to tech giants like Google [9], Microsoft [7], and Apple [2]. By introducing random noise into the computation process, LDP provides robust privacy guarantees for the original data. Specifically, LDP employs a randomized perturbation mechanism \mathcal{M} to perturb the user's original data before transmitting the perturbed version to the server. Formally, the definition of ϵ -LDP is presented as follows.

Definition 1 (ϵ -LDP). A random perturbation mechanism \mathcal{M} satisfies ϵ -LDP, where $\epsilon > 0$, if and only if for any user's private data x and x' , and for all possible outputs $y \in \text{Range}(\mathcal{M})$, we have:

$$\Pr[\mathcal{M}(x) = y] \leq e^\epsilon \cdot \Pr[\mathcal{M}(x') = y], \quad (1)$$

where the parameter ϵ , known as the “privacy budget”, serves to balance utility and privacy. A higher value of ϵ corresponds to weaker privacy guarantees but improved learning utility.

2.3 Locally Private Graph Learning Protocol

The locally private graph learning protocol Π comprises three successive phases: *perturbation*, *calibration*, and *GNN-based learning*.

2.3.1 Perturbation. In this phase, each user $v \in \mathcal{V}$ applies the LDP mechanism \mathcal{M} to perturb their sensitive node features $\mathbf{x}_v \in \mathbb{R}^d$ to ensure privacy, formalized as $\mathbf{x}'_v \leftarrow \mathcal{M}(\mathbf{x}_v, \epsilon)$. The primary LDP mechanisms used include the piecewise mechanism (PM) [38, 48], the multi-bit mechanism (MB) [31, 43], and the square wave mechanism (SW) [28, 29]. Given the high-dimensional nature of node features, these LDP mechanisms follow a structured approach to balance privacy protection and data utility. The general perturbation process consists of two steps: ① the $m \in \{1, 2, \dots, d\}$ values are randomly selected from the d -dimensional space. ② An ϵ/m -LDP noise perturbation is applied to each of the sampled dimensions, while the remaining $d - m$ dimensions are set to zero. Specifically,

- In the one-dimensional multi-bit mechanism [31, 43], for any original input data $x \in [\alpha, \beta]$, the perturbed value x' falls within

the range $\{-1, 1\}$ and is sampled from the following distribution:

$$\Pr[x' = c|x] = \begin{cases} \frac{1}{e^\epsilon + 1} + \frac{x - \alpha}{\beta - \alpha} \cdot \frac{e^\epsilon - 1}{e^\epsilon + 1}, & \text{if } c = 1 \\ \frac{e^\epsilon}{e^\epsilon + 1} - \frac{x - \alpha}{\beta - \alpha} \cdot \frac{e^\epsilon - 1}{e^\epsilon + 1}, & \text{if } c = -1 \end{cases}. \quad (2)$$

- In the one-dimensional piecewise mechanism [38, 48], the input domain is $[\alpha, \beta]$, and the perturbed data range is $[-s, s]$, where $s = \frac{e^{\epsilon/2} + 1}{e^{\epsilon/2} - 1}$. Given an original value x , the perturbed value x' is sampled from the following probability density function:

$$\Pr[x' = c|x] = \begin{cases} p, & \text{if } c \in [l(x), r(x)] \\ p/e^\epsilon, & \text{if } c \in [-s, l(x)) \cup (r(x), s] \end{cases}, \quad (3)$$

where $p = \frac{e^\epsilon - e^{\epsilon/2}}{2e^{\epsilon/2} + 2}$, $l(x) = \frac{s+1}{2} \cdot x - \frac{s-1}{2}$, and $r(x) = l(x) + s - 1$.

- In the one-dimensional square wave mechanism [28, 29], the perturbed data range is $[-b - 1, b + 1]$, where $b = \frac{\epsilon e^\epsilon - e^\epsilon + 1}{e^\epsilon(e^\epsilon - e - 1)}$. The noisy output x' is sampled from the following distribution:

$$\Pr[x' = c|x] = \begin{cases} p, & \text{if } c \in [x - b, x + b] \\ p/e^\epsilon, & \text{if } c \in [-b - 1, x - b) \cup (x + b, 1 + b] \end{cases}, \quad (4)$$

where $p = e^\epsilon / (2be^\epsilon + 2)$. More details on this phase, see App. A.

2.3.2 Calibration. In this phase, for node $v \in \mathcal{V}$, the server calibrates the noise by invoking the following $\text{AGGREGATE}(\cdot)$ function iteratively over the features of neighboring nodes up to K steps away, without applying any non-linear transformation in between:

$$\widehat{\mathbf{h}}_v^{(k)} = \text{AGGREGATE}_k(\{\widehat{\mathbf{h}}_u^{(k-1)}, \forall u \in \mathcal{N}(v)\}), \quad (5)$$

where $\mathbf{h}_v^{(k)}$ is the embedding of node v at layer k , $\mathcal{N}(v)$ is the set of neighbors of v , and $\text{AGGREGATE}_k(\cdot)$ is a differentiable, permutation-invariant aggregation function (such as sum or mean) at layer k . When initialized, we have $\widehat{\mathbf{h}}_v^{(0)} = \mathbf{x}'_v$, i.e., the initial embedding of node v is its perturbed feature vector \mathbf{x}'_v .

2.3.3 GNN-based learning. In this phase, the server performs graph learning tasks such as node classification [51], by training GNN models f_θ , such as graph convolutional networks (GCN) [20], GraphSAGE [13], and graph attention networks (GAT) [47], using node embeddings $\widehat{\mathbf{h}}_v^K$ obtained from the previous calibration phase.

In summary, the locally private graph learning protocol is defined as $\Pi = (\mathcal{M}, \mathcal{A}, f_\theta)$, where \mathcal{M} , \mathcal{A} , and f_θ correspond to the three phases described above. Due to the black-box setup considered in this paper, the attacker cannot directly manipulate the \mathcal{A} and f_θ phases on the server side. Therefore, the poisoning attack is carried out by targeting the \mathcal{M} phase, located on the user side.

3 METHODOLOGY

In this section, we first present the threat model in Sec. 3.1, followed by a detailed description of the data poisoning attacks proposed against protocol Π in Sec. 3.2. Fig. 2 illustrates the attack process.

3.1 Threat Model

3.1.1 Attacker's capability. We assume that the malicious attacker can inject multiple fake nodes $v_{\text{atk}} \in \mathcal{V}_{\text{atk}}$ into the original graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, and establish malicious links between the target node $v_t \in \mathcal{V}$ and the fake nodes, as well as between the fake nodes themselves. The sets of fake and target nodes are defined as \mathcal{V}_{atk} and $\mathcal{V}_t \subset \mathcal{V}$, respectively, with $1 < |\mathcal{V}_{\text{atk}}| \leq |\mathcal{V}_t| \ll |\mathcal{V}|$. The

¹Note that in many real-world applications, graph nodes correspond to human users, and therefore, the terms “node” and “user” are used interchangeably in this paper.

set of malicious links injected is defined as \mathcal{E}_{atk} . Considering the realism and subtlety of the attack, *only one* malicious link is allowed to connect to each target node v_t . Furthermore, for each fake node, the attacker has the freedom to specify its graph data.

3.1.2 Attacker's background knowledge. We assume that the attacker knows the predefined parameters of the LDP mechanisms, as these are publicly available to all users. In addition, we consider a practical black-box scenario in which the attacker has *no access* to any information about the target node, including its features, labels, or neighborhood information, as well as any details about the architecture or parameters of the GNN model in protocol II.

3.1.3 Attacker's goal. The attacker's goal is to deceive the GNN model f_θ in protocol II by strategically crafting malicious fake nodes and malicious edges. This attack aims to compromise the utility of all target nodes while also further degrading the overall learning utility of the entire graph, such as reducing node classification accuracy [20], by exploiting the message-passing mechanism of GNNs. Formally, the attacker's objective is to

$$\min \mathcal{L}_{\text{atk}}(f_\theta(\mathcal{G}')), \quad \text{subject to } \|\mathcal{G}' - \mathcal{G}\|_p \leq \Delta, \quad (6)$$

where \mathcal{L}_{atk} can be defined as $-\mathcal{L}_{\text{train}}$. $\mathcal{L}_{\text{train}}$ denotes the classification loss (e.g., cross-entropy), which measures how well the GNN model performs. $\mathcal{G}' = (\mathcal{V} \cup \mathcal{V}_{\text{atk}}, \mathcal{E} \cup \mathcal{E}_{\text{atk}})$ is the poisoned graph with fake node set \mathcal{V}_{atk} and malicious link set \mathcal{E}_{atk} . Δ denotes the allowed perturbation or stealthiness of the attack. $\|\cdot\|_p$ defines the distance measure between the original and modified graphs.

3.2 Attack Details

3.2.1 Crafting malicious local graph data. In implementing the poisoning attack, we link a single fake node v_{atk} to each target node v_t . The attacker's goal is to compromise the post-aggregation embedding of node v_t by crafting the features $\mathbf{x}'_{v_{\text{atk}}}$ of node v_{atk} , which subsequently degrades the utility of privacy-preserving graph learning. However, the design of $\mathbf{x}'_{v_{\text{atk}}}$ must adhere to the output constraints \mathbb{R} of a specific LDP mechanism [28, 43, 48] for stealth purposes. Thus, the challenge lies in refining the design of $\mathbf{x}'_{v_{\text{atk}}}$ to maximally undermine the embedding of the node v_t while satisfying the constraints \mathbb{R} . The attack objectives are as follows.

Definition 2 (Attack Objective). For any target node $v_t \in \mathcal{V}_t$, the attacker aims to maximize the deviation from its embedding $\hat{\mathbf{h}}_{N(v_t)}$ by manipulating the feature $\mathbf{x}'_{v_{\text{atk}}} = [x'_i]_{i=1}^d$ of fake nodes v_{atk} , i.e.,

$$\max_{\mathbf{x}'_{v_{\text{atk}}}} \left\| \hat{\mathbf{h}}_{N(v_t) \cup \{v_{\text{atk}}\}} - \hat{\mathbf{h}}_{N(v_t)} \right\|, \quad \text{s.t. } \mathbb{R} : x'_i \sim \begin{cases} [-B, B], & i \in \mathcal{S} \\ 0, & i \notin \mathcal{S} \end{cases}, \quad (7)$$

where $\hat{\mathbf{h}}_{N(v_t) \cup \{v_{\text{atk}}\}}$ is the embedding of node v_t after the attack. B denotes the perturbation boundary, which takes the values $\frac{(e^{\epsilon/2}+1)d}{e^{\epsilon/2}-1}$, $\frac{d(\beta-\alpha)}{2m} \frac{e^{\epsilon}+1}{e^{\epsilon}-1} + \frac{\alpha+\beta}{2}$ and $\frac{\bar{\epsilon}e^{\bar{\epsilon}}-e^{\bar{\epsilon}+1}}{e^{\bar{\epsilon}}(e^{\bar{\epsilon}}-\bar{\epsilon}-1)}$ (where $\bar{\epsilon} = \epsilon/m$) for PM [48], MB [43] and SW [28], respectively. \mathcal{S} represents a set of m values drawn uniformly at random without replacement from $\{1, 2, \dots, d\}$. To achieve the above objective, the feature vector of the fake node v_{atk} should be chosen to be as dissimilar as possible to the target node $v_t \in \mathcal{V}_t$, thus maximizing the Euclidean distance between the original and modified embeddings. This can be accomplished by selecting *extreme feature values* according to the output of the

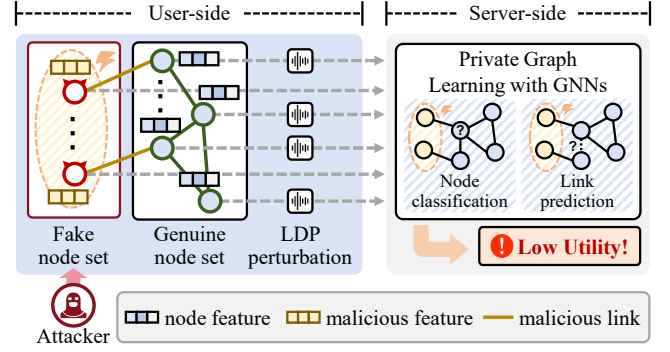


Figure 2: Overview of the proposed data poisoning attacks. The attacker injects multiple fake nodes, establishes links with genuine target nodes, and crafts malicious features for the fake nodes, aiming to maximally degrade the utility of private graph learning, such as node classification accuracy.

private graph learning LDP protocol. Therefore, $\mathbf{x}'_{v_{\text{atk}}}$ is generated as follows: first, randomly select the m dimensions from the d -dimensional space and assign the data of the selected dimension as $B \cdot (-1)^{\mathbf{1}(\mathcal{U} > 0.5)}$, where $\mathcal{U} \sim \text{Uniform}(0, 1)$; then, the remaining $d - m$ dimensions are set to zero. Furthermore, in addition to the node features, we randomly assign a label to each fake node.

3.2.2 Linking fake nodes to target nodes. Consider the set of target nodes \mathcal{V}_t and the set of fake nodes \mathcal{V}_{atk} , where $|\mathcal{V}_{\text{atk}}| \leq |\mathcal{V}_t|$. During the linking process between the target nodes and fake nodes, we propose a cyclic matching strategy \mathbb{M} to ensure that each target node is assigned a fake node. Under this strategy, the fake nodes are recycled, guaranteeing that all target nodes can be successfully matched. For any node $v_t \in \mathcal{V}_t$,

$$\mathbb{M} : \mathcal{V}_t \rightarrow \mathcal{V}_{\text{atk}}, \quad \text{where } \mathbb{M}(v_t) = (v_{\text{atk}})_{(j \bmod |\mathcal{V}_{\text{atk}}| + 1)}, \quad (8)$$

where $\mathbb{M}(v_t)$ denotes the fake node to which the target node v_t is matched, while j is the index of the target node v_t , with values ranging from $1 \leq j \leq |\mathcal{V}_t|$. For each node v_t , we ensure that the indexes of the fake nodes are recycled within the range $[1, |\mathcal{V}_{\text{atk}}|]$ using the modulo operation. This ensures that fake nodes are duplicatedly assigned when there are insufficient fake nodes. This matching mechanism guarantees that all target nodes are matched by effectively utilizing the fake nodes, maximizing their usage when their number is limited, and thereby allowing the attacker to achieve the greatest coverage effect with constrained resources.

3.2.3 Constructing inner links of fake nodes. From the perspective of attack stealthiness, we aim to maintain the average degree of the poisoned graph \mathcal{G}' equal to that of the original graph \mathcal{G} to evade potential homomorphism detection defenses. To achieve this, we connect the $|\mathcal{V}_{\text{atk}}|$ fake nodes with a probability q between them. By controlling the connection probability between the fake nodes, we ensure that the newly added nodes do not significantly alter the global degree distribution of the graph, thereby preserving the original structure of the graph as much as possible in terms of statistical features. Choosing an appropriate connection probability q is crucial, and we define q as follows.

Proposition 1. Let $\langle d_{\text{orig.}} \rangle$ represent the average degree of the original graph \mathcal{G} . After adding $|\mathcal{V}_{\text{atk}}|$ fake nodes, the probability q of linking these fake nodes to each other is

$$q = \frac{\langle d_{\text{orig.}} \rangle - 2 \cdot |\mathcal{V}_t| / |\mathcal{V}_{\text{atk}}|}{|\mathcal{V}_{\text{atk}}| - 1}, \quad (9)$$

where $|\mathcal{V}_{\text{atk}}| > 2|\mathcal{V}_t| / \langle d_{\text{orig.}} \rangle$, $|\mathcal{V}_t| \geq (\langle d_{\text{orig.}} \rangle + 1)^2 / 8$, and $\langle d_{\text{orig.}} \rangle \geq 2$. These conditions are easy to fulfill in realistic scenarios. Please refer to App. B.1 for the proof and further explanation of Eq. (9).

3.2.4 Optimizing the features of all fake nodes. Note that further optimizing the malicious feature vectors of all fake nodes from a global perspective is beneficial to maximize the damage to the target nodes. Specifically, as described in Sec. 3.2.1, the target node $v_t \in \mathcal{V}_t$ is linked to a fake node. Furthermore, as detailed in Sec. 3.2.3, this fake node is connected to multiple other fake nodes with the estimation of $q \cdot (|\mathcal{V}_{\text{atk}}| - 1)$. In this topology, crafting malicious node features in a way that maximizes the impact of the attack is crucial. To this end, we consider three ways of crafting the non-zero bits in the malicious node features: ① *Random*. The non-zero bits in the features of all fake nodes are randomly selected. ② *Diverse*. The non-zero bits in the features of all fake nodes differ from each other, which is formalized as $S_v \cap S_u = \emptyset$, for $v, u \in \mathcal{V}_{\text{atk}}$. ③ *Identical*. The non-zero bits in the features of all fake nodes are identical, formalized as $S_v = S_u$, for $v, u \in \mathcal{V}_{\text{atk}}$. The analysis in Sec. 4.1 demonstrates that the “Identical” feature setting maximizes the attack’s effectiveness, inflicting the most damage on the target node. We adopt this setting in subsequent experiments.

4 THEORETICAL ANALYSIS

In this section, we successively analyze the error propagation affecting the target nodes and the entire graph in Sec. 4.1 and Sec. 4.2, and finally discuss the security-privacy trade-off in Sec. 4.3.

4.1 Error Analysis of Target Nodes

In Π , for each target node v_t , the two properties—*unbiased expectation* (Def. 3) and *variance minimization*—ensure optimal perturbation of the node features, forming the basis for high-utility private graph learning. However, the injection of fake nodes disrupts this balance, may introducing biased expectations and increasing variance, ultimately degrading the utility of the learned representations.

Definition 3 (Unbiased Expectation). In Π , the perturbed node features remain unbiased, i.e., $\mathbb{E}[\mathbf{x}'] = \mathbf{x}$. Furthermore, the subsequent aggregator defined by Eq. (5) also provides an unbiased estimate.

As in Prop. 2, we quantify the bias $\Delta_{\mathbb{E}}$ and Δ_{var} introduced by the expectation and variance of our attack on the target node. Through this analysis, we demonstrate that a *identical* approach to crafting malicious features facilitates maximum damage to the target node.

Proposition 2. Considering two-layer aggregation, our attack does not introduce an expectation bias. However, it causes a variance bias Δ_{var} . We quantify the expected value of the variance of the aggregated embedding before and after the attack in App. B.2.

Table 1: Statistics of datasets.

Type	Dataset	Nodes	Edges	Features	Classes	Avg. Deg.
Citation Network	Cora	2,708	5,278	1,433	7	3.90
	Citeseer	3,327	4,552	3,703	6	2.74
	Pubmed	19,717	44,324	500	3	4.50
Social Network	LastFM	7,624	27,806	7,842	18	7.29
	Twitch	4,648	61,706	128	2	16.19
	Facebook	22,470	170,912	4,714	4	15.21

Avg. Deg. in the table denotes the average node degree.

4.2 Error Analysis of the Entire Graph

After injecting fake nodes and propagating malicious information through the K -layer aggregation process, our attack severely distorts the node representation in the entire graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. To quantify this distortion, we define the global error energy function [37] $\Psi(\mathcal{G})$ before the attack as follows:

$$\Psi(\mathcal{G}) = \sum_{v \in \mathcal{V}} \|\mathbf{h}_v^{(K)} - \mathbf{x}_v\|^2 + \lambda \sum_{(u,v) \in \mathcal{E}} \|\mathbf{h}_u^{(K)} - \mathbf{h}_v^{(K)}\|^2, \quad (10)$$

where λ represents the aggregation balancing parameter. For the graph after the attack, denoted \mathcal{G}' , the global error energy function is defined as $\Psi(\mathcal{G}')$. Furthermore, the global error increment is quantified as $\Delta\Psi = \Psi(\mathcal{G}') - \Psi(\mathcal{G})$ and the expected deviation is

$$\mathbb{E}[\Delta\Psi] = (1 + \lambda q / (1 - \lambda q)) \cdot |\mathcal{V}_{\text{atk}}| \cdot B^2 \cdot K, \quad (11)$$

where $|\mathcal{V}_{\text{atk}}|$ is the number of injected fake nodes. The result indicates that the perturbation impact is linearly related to $|\mathcal{V}_{\text{atk}}|$. Furthermore, higher values of B and K exacerbate the adversarial impact on the graph. This reveals the fundamental vulnerability of GNNs in our poisoning attack. More details on Eq. (11) in App. B.3.

4.3 Security-Privacy Trade-Off

Additionally, we observe a security-privacy trade-off in our attack on protocol Π , highlighting the inherent tension between enhancing privacy protection and maintaining robustness against poisoning attacks. As the privacy budget decreases, the level of noise injection increases, reducing the model’s sensitivity to individual data points. However, this also amplifies the impact of adversarial perturbations. Mathematically, this relationship can be expressed as $\mathbb{E}[\Delta\Psi] \propto e^{-\epsilon}$, where privacy budget ϵ controls the level of privacy, and a lower ϵ exacerbates the adversarial impact. More details on this phenomenon can be found in the experiments in Sec. 5.2.

5 EVALUATION

In this section, we comprehensively evaluate the performance of our attack based on node classification and link prediction tasks.²

5.1 Experimental Settings

5.1.1 Datasets. We conduct experiments on six representative real-world datasets: three citation networks (Cora [55], Citeseer [55], and Pubmed [55]) and three social networks (LastFM [42], Twitch [41], and Facebook [41]). We summarize the dataset statistics in Table 1.

5.1.2 GNN models. We consider three state-of-the-art GNN models: graph convolutional networks (GCN) [20], GraphSAGE (SAGE) [13], and graph attention networks (GAT) [47]. All GNN models consist of two graph convolutional layers, with 64 neurons in each

²The code will be made publicly available upon acceptance of the paper.

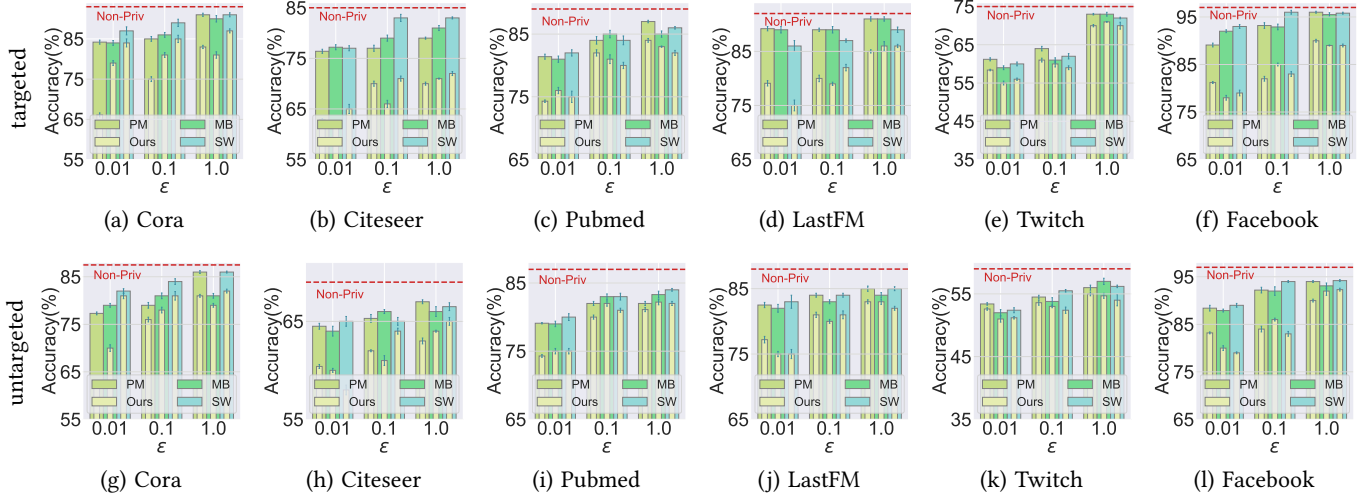


Figure 3: Attack performance under different privacy budgets ϵ and LDP mechanisms for node classification task. The x-axis represents the privacy budget ϵ , while the y-axis shows the accuracy for both the targeted and untargeted attacks. “Non-Priv” denotes the non-perturbed baseline. Our attacks significantly reduce the accuracy of various LDP mechanisms across all cases.

hidden layer and ReLU as the activation function [21], followed by dropout. The GAT model has four attention heads. We implement the GNN models in PyTorch using PyTorch-Geometric (PyG)³. All experiments are carried on a machine running Ubuntu 20.04 LTS, equipped with two Intel® Xeon® Gold 6348 CPUs, 100GB RAM and an NVIDIA® A800 80GB GPU. By default, the model used is GCN.

5.1.3 LDP mechanism. We consider three de facto LDP perturbation mechanisms for protecting node features: the multi-bit mechanism (MB) [31, 43], the piecewise mechanism (PM) [38, 48], and the square wave mechanism (SW) [28, 29]. Both mechanisms perturb the node features based on a privacy budget ϵ , ensuring strict privacy guarantees. By default, the LDP mechanism used is PM.

5.1.4 Parameter settings. For all datasets, we randomly divide the nodes into training, validation, and testing sets with a 50/25/25% split. The privacy budget ϵ ranges from $\{0.01, 0.1, 1.0\}$, the ratio of target nodes to the original graph, $\eta_1 = |\mathcal{V}_t|/|\mathcal{V}|$, ranges from $\{0.01, 0.03, 0.05, 0.07, 0.09\}$, and the ratio of fake nodes to target nodes, $\eta_2 = |\mathcal{V}_{atk}|/|\mathcal{V}_t|$, ranges from $\{0.6, 0.7, 0.8, 0.9, 1.0\}$. In default, ϵ is set to 0.01; η_1 and η_2 are set to 0.09 and 0.8, respectively. The parameter K for the calibration phase takes values from the range $\{0, 2, 4, 8, 16\}$. For more details on hyperparameters, refer to App. C.1.

5.1.5 Evaluation metrics. We conduct experiments based on node classification [28, 31, 38, 43] and link prediction [59, 64] tasks. We use node classification accuracy (or link prediction accuracy) as the evaluation metric to assess the effectiveness of attacks on all target nodes (*targeted attacks*) and on the entire graph (*untargeted attacks*). The default task is node classification. We measure the accuracy of the test set over 10 consecutive runs and report the mean and 95% confidence intervals, calculated via bootstrapping with 1000 samples. More details, refer to App. C.2.

Table 2: Comparison of accuracy before and after the attack across different GNN models. ↓ represents accuracy drop.

GNN	Meth.	Cora	Citeseer	Pubmed	LastFM	Twitch	Facebook
GCN	PM	84.2	76.4	81.4	89.2	61.2	89.1
	PM [†]	66.1 ↓	62.9 ↓	74.3 ↓	79.1 ↓	58.4 ↓	81.2 ↓
SAGE	PM	82.7	71.9	78.5	81.8	61.8	83.3
	PM [†]	62.3 ↓	55.2 ↓	55.0 ↓	54.8 ↓	56.5 ↓	73.9 ↓
GAT	PM	83.1	77.9	77.7	84.3	64.3	83.6
	PM [†]	67.8 ↓	54.9 ↓	72.2 ↓	58.9 ↓	62.3 ↓	73.4 ↓

[†] represents the accuracy after targeted attack.

5.2 Evaluating the Effectiveness of the Attack

5.2.1 Node Classification. In this experiment, we assess the attack performance under different privacy budgets and LDP mechanisms using the node classification task, with the results shown in Fig. 3. As demonstrated in the figure, our attack method significantly reduces the accuracy of the three LDP mechanisms in all cases, validating its effectiveness. Notably, the accuracy reduction is more pronounced under the targeted attack (affecting the target nodes, top row of Fig. 3) than under the untargeted attack (affecting the entire graph, bottom row). This is because our attack specifically targets the set of nodes, which further compromises the utility of the entire graph through malicious message passing, making the impact more significant for the target nodes. Furthermore, our attack method proves to be more effective under higher privacy settings. For instance, in Fig. 3(c), when $\epsilon = 1.0$, there is an average accuracy degradation of 3%, while for $\epsilon = 0.01$, the degradation rises to 8.4%. This is due to the fact that higher privacy settings result in more noise being injected into the node features, reducing the proportion of benign information, thereby amplifying the impact of malicious node features on the private graph learning process.

5.2.2 Link Prediction. In addition to node classification, we also evaluate the attack performance under different privacy budgets $\epsilon \in \{0.01, 0.1, 1.0\}$ and LDP mechanisms (PM, MB, and SW) using

³<https://www.pyg.org>




Figure 4: Variation in target attack accuracy at different η_1 and η_2 values. The impact ratio represents the proportion of accuracy reduction compared to the normal accuracy.

the link prediction task, with the results presented in Fig. 5. As shown in the figure, our attack methods significantly reduce the accuracy of the different LDP mechanisms in this task setting as well, confirming the effectiveness of our attacks. Link prediction works by inferring the likelihood of a link between two nodes based on the similarity of their embeddings. Our attack compromises the embedding of the target node, which then propagates through the graph via malicious message passing, affecting the embeddings of neighboring nodes. This interference disrupts the node-to-node link prediction and results in a substantial decrease in utility. Furthermore, similar experimental trends to those observed in the node classification task are evident here, so these results are not repeated.

5.3 Evaluating the Effectiveness of the Attack Across Different GNN Models

In this experiment, we evaluate the performance of the attack across six representative datasets using three GNN models (GCN, SAGE, GAT). As shown in the targeted attack results in Table 2, our attack significantly reduces the accuracy of the original PM mechanism across different GNN models and datasets, with an average accuracy reduction of over 10% for all three models. These results demonstrate the generalization and effectiveness of our attack method across various GNN models. Additional experimental results and analysis on different GNN models can be found in App. C.3.1.

5.4 Parameter Analysis

5.4.1 Analyzing the effect of the parameter η_1 . As shown in the left panel of Fig. 4, the accuracy before attack remains relatively stable for different values of $\eta_1 \in \{0.01, 0.03, 0.05, 0.07, 0.09\}$. However, as η_1 increases—meaning the number of target nodes the attacker can damage grows—the accuracy after the attack steadily decreases. Correspondingly, the impact ratio (the ratio of reduced accuracy to normal accuracy) increases, highlighting the expanding malicious impact of our attack. This occurs because as η_1 increases, the attacker can influence more genuine nodes, leading to greater disruption in the utility of the privacy-preserving learning process.

5.4.2 Analyzing the effect of the parameter η_2 . As shown in the right panel of Fig. 4, we validate the effect of different η_2 on the attack, and the observed trend is similar to that of η_1 . As η_2 increases, the attacker can inject more fake nodes, thereby introducing more malicious information. This results in greater damage to the private learning utility. For more details on η_1 and η_2 , see App. C.3.2.

5.4.3 Analyzing the effect of the parameter K . As shown in Fig. 6, we validate the change in accuracy before and after the attack for different values of $K \in \{0, 2, 4, 8, 16\}$. As K increases, accuracy follows a trend of first increasing and then decreasing. The initial increase is due to noise calibration through neighbor aggregation, while the subsequent decrease is caused by over-smoothing [18, 25], which leads to the convergence of node embeddings. Notably, our method significantly reduces the accuracy of the original method across all values of K , demonstrating the generalization and effectiveness of our approach. Further analysis can be found in App C.3.3.

6 DEFENSES

In this section, we evaluate the resilience of our attack approach against potential defenses. We sequentially consider two possible defenses: *graph homophily analysis* (Sec. 6.1) and *anomaly node detection* (Sec. 6.2). We also discuss other possible defenses in Sec. 6.3.

6.1 Graph Homophily Analysis

Existing studies [4] have shown that attacks [67, 68] against graphs can significantly disrupt the homophily distribution [33] of the original graph. To address this, defenses based on homophily analysis [4, 45, 60, 69] have been proposed to detect such attacks. In this paper, we conduct graph homophily analysis on all datasets both before and after the attack. Specifically, we use two homophily metrics [4]: *node-centered homophily*, which measures the similarity between a target node and its neighbors, and *edge-centered homophily*, which compares the similarity between nodes connected by a target edge. See App. C.3.4 for more details on these metrics.

Results. We perform homophily analyses using cosine similarity, and Fig. 7 compares the homophily distributions of different datasets before and after the attack. As noted in [4], homophily-based defenses can detect malicious nodes or edges with abnormally low homophily. However, as shown in the figure, our attack method does not cause significant deviations. This is because our carefully designed node features and links perturb the original graph on a much smaller scale than typical adversarial attacks [67]. Additionally, Fig. 7 shows that node-centered homophily demonstrates higher similarity than edge-centered homophily across all datasets. This observation aligns with previous studies [33, 34, 67]. These results demonstrate that our attack method is highly resistant to homophily-based defenses.

6.2 Anomaly Node Detection

We further evaluate the resistance of our attacks to existing anomaly detection methods [32]. Specifically, we employ two widely used techniques: ① *Girvan-Newman Community Discovery*. The Girvan-Newman (GN) algorithm [12] recursively divides communities by removing edges with the highest betweenness centrality. In a typical scenario, normal nodes maintain strong connectivity with others in their community, while malicious attacks can lead to misclassification or isolation of certain nodes, which appear as anomalies in the graph structure. ② *Node Features Based on KMeans Clustering*. KMeans clustering [30] is an unsupervised learning method that partitions nodes into clusters. Anomalous nodes might either form outliers or be misassigned to incorrect clusters, making them detectable via feature-based anomaly detection.




Figure 5: Attack performance under different privacy budgets ϵ and LDP mechanisms for link prediction task. The x-axis represents the privacy budget ϵ , while the y-axis shows the accuracy for both the targeted and untargeted attacks. “Non-Priv” denotes the non-perturbed baseline. Our attacks significantly reduce the accuracy of various LDP mechanisms across all cases.




Figure 6: Variation in accuracy before and after targeted (left plot) or untargeted (right plot) attack for different K values.

Results. As shown in the left panel of Fig. 8, our attack method does not introduce noticeable anomalies in community division, indicating that it maintains strong stealth under GN detection. Meanwhile, as shown in the right panel of Fig. 8, the perturbation introduced by our attack in the feature space is minimal and does not significantly alter the node distribution, making it difficult to detect under KMeans evaluation. These anomaly node detection experiments further validate the stealthiness of our attack method.

6.3 Other Defenses

We also consider other potential defenses. First, this paper aims to compromise locally private graph learning protocols by injecting fake nodes and edges, and authenticated data feed systems may be able to help reestablish trust in the data source. However, current solutions are limited to well-known entities [57], and authenticating data from unknown sources in distributed privacy scenarios remains an open problem. Furthermore, while adversarial defense methods [4, 45, 60, 65, 69] may provide some mitigation, our experiments demonstrate that due to the careful design of node features and links, the perturbation scale of our attack is significantly smaller than that of adversarial attacks [67, 68]. Consequently, our attack exhibits strong resistance against adversarial defenders. Finally, other possible defenses include hardware-assisted trusted computation [5, 22, 24, 26], but this approach may introduce additional costs

for hardware-software co-design and security auditing to prevent extensive side-channel attacks [35, 44, 46].

7 RELATED WORK

7.0.1 Data Poisoning Attacks on GNNs. Data poisoning attacks aim to disrupt the learning process of machine learning models by injecting malicious data into the training set. These attacks [45, 67–69] are especially concerning in the context of GNNs due to the complex dependencies between nodes, edges, and graph structures that GNNs rely on to learn node representations. As discussed in Sec. 1, the data poisoning attack introduced in this paper, specific to locally private graph learning protocols, differs fundamentally from previous work on traditional attacks against GNNs. This distinction primarily arises from the strict data constraints, limited background knowledge, and the complexity of the protocols involved.

7.0.2 Locally Private Graph Learning Protocols. Recently, locally private graph learning protocols have attracted significant attention in the security research community [15, 28, 31, 38, 40, 43, 58, 66]. These approaches leverage the benefits of LDP [8, 54] in safeguarding data privacy and ensuring strict privacy guarantees for user data. Moreover, during server-side privacy graph learning, GNN’s multi-layer message passing mechanism [1, 10] effectively reduces the estimation error in node embeddings after perturbation and aids in the calibration of noisy data, thus maintaining high utility for the local privacy graph learning task. However, despite these advantages, such protocols may be vulnerable to data poisoning attacks. More details on this section are in App D.

8 CONCLUSION

In this paper, we present the first study on data poisoning attacks targeting locally private graph learning protocols, with the aim of compromising the utility of private graph learning within these protocols. Through theoretical analysis and empirical evaluation, we demonstrate that our attacks significantly degrade the utility of




Figure 7: Comparison of homophily distributions before and after the attack. The results show that before and after the attack do not exhibit significant deviations, which indicates that our attack method is highly resistant to homophily-based defense.




Figure 8: Anomaly node detection based on Girvan-Newman community discovery (left) and KMeans clustering (right).

private graph learning. Additionally, we examine several potential defense strategies to mitigate these attacks, but our findings underscore the urgent need for more robust countermeasures. We believe this work is essential for advancing the development of secure and resilient privacy-preserving graph learning frameworks.

REFERENCES

- [1] Sami Abu-El-Hajja, Bryan Perozzi, Amol Kapoor, Nazanin Alipourfard, Kristina Lerman, Hrayr Harutyunyan, Greg Ver Steeg, and Aram Galstyan. 2019. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *International Conference on Machine Learning (ICML 2019)*. PMLR, 21–29.
- [2] D Apple. 2017. Learning with privacy at scale. *Apple Machine Learning Journal* 1, 8 (2017), 71.
- [3] Christopher J Carpenter. 2012. Narcissism on Facebook: Self-promotional and anti-social behavior. *Personality and individual differences* 52, 4 (2012), 482–486.
- [4] Yongqiang Chen, Han Yang, Yonggang Zhang, MA KAILI, Tongliang Liu, Bo Han, and James Cheng. 2022. Understanding and Improving Graph Injection Attack by Promoting Unnoticeability. In *International Conference on Learning Representations (ICLR 2022)*.
- [5] André Cirne, Patricia R Sousa, João S Resende, and Luís Antunes. 2024. Hardware security for Internet of Things identity assurance. *IEEE Communications Surveys & Tutorials* (2024).
- [6] Graham Cormode, Somesh Jha, Tejas Kulkarni, Ninghui Li, Divesh Srivastava, and Tianhao Wang. 2018. Privacy at scale: Local differential privacy in practice. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD 2018)*. 1655–1658.
- [7] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. 2017. Collecting telemetry data privately. *Advances in Neural Information Processing Systems (NeurIPS 2017)* 30 (2017).
- [8] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. 2006. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography: Third Theory of Cryptography Conference (TCC 2006)*. Springer, 265–284.
- [9] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. 2014. Rappor: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security (CCS 2014)*. 1054–1067.
- [10] Jiarui Feng, Yixin Chen, Fuhai Li, Anindya Sarkar, and Muhan Zhang. 2022. How powerful are k-hop message passing graph neural networks. *Advances in Neural Information Processing Systems (NeurIPS 2022)* 35 (2022), 4776–4790.
- [11] Alex Fout, Jonathon Byrd, Basir Shariat, and Asa Ben-Hur. 2017. Protein interface prediction using graph convolutional networks. *Advances in Neural Information Processing Systems (NeurIPS 2017)* 30 (2017).
- [12] Michelle Girvan and Mark EJ Newman. 2002. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* 99, 12 (2002), 7821–7826.
- [13] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).
- [14] William L Hamilton. 2020. *Graph representation learning*. Morgan & Claypool Publishers.
- [15] Seira Hidano and Takao Murakami. 2024. Degree-Preserving Randomized Response for Graph Neural Networks under Local Differential Privacy. *Transactions on Data Privacy* 17, 2 (2024), 89–121.
- [16] Peter Kairouz, Keith Bonawitz, and Daniel Ramage. 2016. Discrete distribution estimation under local privacy. In *International Conference on Machine Learning (ICML 2016)*. PMLR, 2436–2444.
- [17] Peter Kairouz, Sewoong Oh, and Pramod Viswanath. 2014. Extremal mechanisms for local differential privacy. *Advances in Neural Information Processing Systems (NeurIPS 2014)* 27 (2014).
- [18] Nicolas Keriven. 2022. Not too little, not too much: a theoretical analysis of graph (over) smoothing. *Advances in Neural Information Processing Systems (NeurIPS 2022)* 35 (2022), 2268–2281.
- [19] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [20] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations (ICLR 2017)*.
- [21] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. 2017. Self-normalizing neural networks. *Advances in Neural Information Processing Systems (NeurIPS 2017)* 30 (2017).
- [22] Dayeol Lee, David Kohlbrenner, Shweta Shinde, Krste Asanović, and Dawn Song. 2020. Keystone: An open framework for architecting trusted execution environments. In *Proceedings of the Fifteenth European Conference on Computer Systems (EuroSys 2020)*. 1–16.
- [23] Dian Lei, Zijun Song, Yanli Yuan, Chunhai Li, and Liehuang Zhu. 2025. Achieving Personalized Privacy-Preserving Graph Neural Network via Topology Awareness. In *The Web Conference 2025 (WWW 2025)*.
- [24] Mengyuan Li, Yuheng Yang, Guoxing Chen, Mengjia Yan, and Yinqian Zhang. 2024. SoK: Understanding Design Choices and Pitfalls of Trusted Execution Environments. In *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security (ASIACCS 2024)*. 1600–1616.
- [25] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI 2018)*, Vol. 32.
- [26] Xupeng Li, Xuheng Li, Christoffer Dall, Ronghui Gu, Jason Nieh, Yousuf Sait, and Gareth Stockwell. 2022. Design and verification of the arm confidential compute architecture. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2022)*. 465–484.
- [27] Yujia Li, Chenjie Gu, Thomas Dullien, Oriol Vinyals, and Pushmeet Kohli. 2019. Graph matching networks for learning the similarity of graph structured objects. In *International Conference on Machine Learning (ICML 2019)*. PMLR, 3835–3845.
- [28] Zening Li, Rong-Hua Li, Meihao Liao, Fusheng Jin, and Guoren Wang. 2024. Privacy-Preserving Graph Embedding based on Local Differential Privacy. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management (CIKM 2024)*. 1316–1325.
- [29] Zitao Li, Tianhao Wang, Milan Lopuhaä-Zwakenberg, Ninghui Li, and Boris Skoric. 2020. Estimating numerical distributions under local differential privacy. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data (SIGMOD 2020)*. 621–635.
- [30] Aristidis Likas, Nikos Vlassis, and Jakob J Verbeek. 2003. The global k-means clustering algorithm. *Pattern Recognition* 36, 2 (2003), 451–461.

- [31] Wanyu Lin, Baochun Li, and Cong Wang. 2022. Towards private learning on decentralized graphs with local differential privacy. *IEEE Transactions on Information Forensics and Security* 17 (2022), 2936–2946.
- [32] Xiaoxiao Ma, Jia Wu, Shan Xue, Jian Yang, Chuan Zhou, Quan Z Sheng, Hui Xiong, and Leman Akoglu. 2021. A comprehensive survey on graph anomaly detection with deep learning. *IEEE Transactions on Knowledge and Data Engineering* 35, 12 (2021), 12012–12038.
- [33] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27, 1 (2001), 415–444.
- [34] Lingshuo Meng, Yijie Bai, Yanjiao Chen, Yutong Hu, Wenyuan Xu, and Haiqin Weng. 2023. Devil in Disguise: Breaching Graph Neural Networks Privacy through Infiltration. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security (CCS 2023)*. 1153–1167.
- [35] Kit Murdoch, David Oswald, Flavio D Garcia, Jo Van Bulck, Daniel Gruss, and Frank Piessens. 2020. Plundervolt: Software-based fault injection attacks against Intel SGX. In *2020 IEEE Symposium on Security and Privacy (S&P)*. IEEE, 1466–1482.
- [36] Michele Nitti, Luigi Atzori, and Irena Pletikosa Cvijikj. 2014. Friendship selection in the social internet of things: challenges and possible strategies. *IEEE Internet of things journal* 2, 3 (2014), 240–247.
- [37] Antonio Ortega, Pascal Frossard, Jelena Kovačević, José MF Moura, and Pierre Vanderghenst. 2018. Graph signal processing: Overview, challenges, and applications. *Proc. IEEE* 106, 5 (2018), 808–828.
- [38] Xinjun Pei, Xiaoheng Deng, Shengwei Tian, Jianqing Liu, and Kaiping Xue. 2023. Privacy-enhanced graph neural network for decentralized local graphs. *IEEE Transactions on Information Forensics and Security* (2023).
- [39] Yang Pei, Renxin Mao, Yang Liu, Chaoran Chen, Shifeng Xu, Feng Qiang, and Blue Elephant Tech. 2021. Decentralized federated graph neural networks. In *International Workshop on Federated and Transfer Learning for Data Sparsity and Confidentiality in Conjunction with IJCAI*.
- [40] Yuxin Qi, Xi Lin, Ziyao Liu, Gaolei Li, Jingyu Wang, and Jianhua Li. 2024. LinkGuard: Link Locally Privacy-Preserving Graph Neural Networks with Integrated Denoising and Private Learning. In *Companion Proceedings of the ACM on Web Conference 2024 (WWW 2024)*. 593–596.
- [41] Benedek Rozemberczki, Carl Allen, and Rik Sarkar. 2021. Multi-scale attributed node embedding. *Journal of Complex Networks* 9, 2 (2021), cnab014.
- [42] Benedek Rozemberczki and Rik Sarkar. 2020. Characteristic functions on graphs: Birds of a feather, from statistical descriptors to parametric models. In *Proceedings of the 29th ACM international conference on information & knowledge management*. 1325–1334.
- [43] Sina Sajadmanesh and Daniel Gatica-Perez. 2021. Locally private graph neural networks. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS 2021)*. 2130–2145.
- [44] Michael Schwarz, Samuel Weiser, Daniel Gruss, Clémentine Maurice, and Stefan Mangard. 2017. Malware guard extension: Using SGX to conceal cache attacks. In *Detection of Intrusions and Malware, and Vulnerability Assessment: 14th International Conference (DIMVA 2017)*. Springer, 3–24.
- [45] Lichao Sun, Yingdong Dou, Carl Yang, Kai Zhang, Ji Wang, S Yu Philip, Lifang He, and Bo Li. 2022. Adversarial attack and defense on graph data: A survey. *IEEE Transactions on Knowledge and Data Engineering* 35, 8 (2022), 7693–7711.
- [46] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenisch, Yuval Yarom, and Raoul Strackx. 2018. Foresadow: Extracting the keys to the intel {SGX} kingdom with transient {Out-of-Order} execution. In *27th USENIX Security Symposium (USENIX Security 2018)*. 991–1008.
- [47] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *6th International Conference on Learning Representations (ICLR 2018)*.
- [48] Ning Wang, Xiaokui Xiao, Yin Yang, Jun Zhao, Siu Cheung Hui, Hyejin Shin, Junbum Shin, and Ge Yu. 2019. Collecting and analyzing multidimensional data with local differential privacy. In *2019 IEEE 35th International Conference on Data Engineering (ICDE 2019)*. IEEE, 638–649.
- [49] Tianhao Wang, Bolin Ding, Jingren Zhou, Cheng Hong, Zhicong Huang, Ninghui Li, and Somesh Jha. 2019. Answering multi-dimensional analytical queries under local differential privacy. In *Proceedings of the 2019 International Conference on Management of Data (SIGMOD 2019)*. 159–176.
- [50] Xiuling Wang and Wendy Hui Wang. 2022. Group property inference attacks against graph neural networks. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS 2022)*. 2871–2884.
- [51] Jun Wu, Jingrui He, and Jiejun Xu. 2019. Net: Degree-specific graph neural networks for node and graph classification. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2019)*. 406–415.
- [52] Lingfei Wu, Peng Cui, Jian Pei, Liang Zhao, and Xiaojie Guo. 2022. Graph neural networks: foundation, frontiers and applications. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2022)*. 4840–4841.
- [53] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. 2022. Graph neural networks in recommender systems: a survey. *Comput. Surveys* 55, 5 (2022), 1–37.
- [54] Mengmeng Yang, Taolin Guo, Tianqing Zhu, Ivan Tjuawinata, Jun Zhao, and Kwok-Yan Lam. 2024. Local differential privacy and its applications: A comprehensive survey. *Computer Standards & Interfaces* 89 (2024), 103827.
- [55] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International Conference on Machine Learning (ICML 2016)*. PMLR, 40–48.
- [56] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2018)*. 974–983.
- [57] Fan Zhang, Ethan Cecchetti, Kyle Croman, Ari Juels, and Elaine Shi. 2016. Town crier: An authenticated data feed for smart contracts. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS 2016)*. 270–282.
- [58] Guan hong Zhang, Xiang Cheng, Jiaan Pan, Zihan Lin, and Zhaofeng He. 2024. Locally differentially private graph learning on decentralized social graph. *Knowledge-Based Systems* 304 (2024), 112488.
- [59] Muhan Zhang and Yixin Chen. 2018. Link prediction based on graph neural networks. *Advances in Neural Information Processing Systems (NeurIPS 2018)* 31 (2018).
- [60] Xiang Zhang and Marinka Zitnik. 2020. GnnGuard: Defending graph neural networks against adversarial attacks. *Advances in Neural Information Processing Systems (NeurIPS 2020)* 33 (2020), 9263–9275.
- [61] Xiao-Meng Zhang, Li Liang, Lin Liu, and Ming-Jing Tang. 2021. Graph neural networks and their current applications in bioinformatics. *Frontiers in genetics* 12 (2021), 690049.
- [62] Yi Zhang, Yuying Zhao, Zhaoping Li, Xueqi Cheng, Yu Wang, Olivera Kotevska, S Yu Philip, and Tyler Derr. 2024. A survey on privacy in graph neural networks: Attacks, preservation, and applications. *IEEE Transactions on Knowledge and Data Engineering* (2024).
- [63] Zhikun Zhang, Min Chen, Michael Backes, Yun Shen, and Yang Zhang. 2022. Inference attacks against graph neural networks. In *31st USENIX Security Symposium (USENIX Security 2022)*. 4543–4560.
- [64] Tianyi Zhao, Jian Kang, and Lu Cheng. 2024. Conformalized link prediction on graph neural networks. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2024)*. 4490–4499.
- [65] Dingyuan Zhu, Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2019. Robust graph convolutional networks against adversarial attacks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2019)*. 1399–1407.
- [66] Xiaochen Zhu, Vincent YF Tan, and Xiaokui Xiao. 2023. Blink: Link Local Differential Privacy in Graph Neural Networks via Bayesian Estimation. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*. 2651–2664.
- [67] Xu Zou, Qinkai Zheng, Yuxiao Dong, Xinyu Guan, Evgeny Kharlamov, Jialiang Lu, and Jie Tang. 2021. Tdgia: Effective injection attacks on graph neural networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD 2021)*. 2461–2471.
- [68] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD 2018)*. 2847–2856.
- [69] Daniel Zügner, Oliver Borchert, Amir Akbarnejad, and Stephan Günnemann. 2020. Adversarial attacks on graph neural networks: Perturbations and their patterns. *ACM Transactions on Knowledge Discovery from Data* 14, 5 (2020), 1–31.

A LDP MECHANISMS

In Sec. 2.3.1, we introduce three different LDP mechanisms for node feature perturbation. To enhance clarity and provide a structured understanding, we present a generalized algorithmic description using the piecewise mechanism [38, 48] as a representative example, as outlined in Alg. 1. In the one-dimensional piecewise mechanism, the input domain is $[\alpha, \beta]$, and the perturbed data range is $[-s, s]$, where $s = \frac{e^{\epsilon/2} + 1}{e^{\epsilon/2} - 1}$. Given an original value x , the perturbed value x' is sampled from the following probability density function:

$$\Pr[x' = c|x] = \begin{cases} p, & \text{if } c \in [l(x), r(x)] \\ p/e^\epsilon, & \text{if } c \in [-s, l(x)) \cup (r(x), s] \end{cases}, \quad (12)$$

where $p = \frac{e^\epsilon - e^{\epsilon/2}}{2e^{\epsilon/2} + 2}$, $l(x) = \frac{s+1}{2} \cdot x - \frac{s-1}{2}$, and $r(x) = l(x) + s - 1$.

Algorithm 1: Piecewise Mechanism For Node Feature

Input: node feature vector $\mathbf{x} \in [\alpha, \beta]^d$ and privacy budget ϵ .
Output: perturbed node feature vector $\mathbf{x}' \in [-s \cdot d, s \cdot d]^d$.

- 1 Let $\mathbf{x}' = \langle 0, 0, \dots, 0 \rangle$;
- 2 Let $m = \max\{1, \min\{d, \lfloor \epsilon/2.5 \rfloor\}\}$;
- 3 Sample m values uniformly without replacement from $\{1, 2, \dots, d\}$;
- 4 **for** each sampled dimension i **do**
- 5 Feed \mathbf{x}_i and $\frac{\epsilon}{m}$ as input to Eq. (12), and obtain a noisy value t_i ;
- 6 $\mathbf{x}'_i = \frac{d}{m} \cdot t_i$;
- 7 **end**
- 8 **return** \mathbf{x}'

B THEORETICAL PROOF**B.1 Proof of Proposition 1**

PROOF. To ensure that the degree remains constant before and after the attack, we have the following equation for the average degree of the original graph:

$$\langle d_{\text{orig.}} \rangle = 2 \left(|\mathcal{E}| + |\mathcal{V}_t| + \binom{|\mathcal{V}_{\text{atk}}|}{2} \cdot q \right) / (|\mathcal{V}| + |\mathcal{V}_{\text{atk}}|), \quad (13)$$

where $\langle d_{\text{orig.}} \rangle$ represent the average degree of the original graph \mathcal{G} , $|\mathcal{V}|$ and $|\mathcal{E}|$ denote the number of nodes and edges in the graph \mathcal{G} , respectively. $|\mathcal{V}_t|$ denotes the number of target nodes. $|\mathcal{V}_{\text{atk}}|$ denotes the number of fake nodes. Additionally, we have the following relationship for the number of edges in the original graph:

$$|\mathcal{E}| = (\langle d_{\text{orig.}} \rangle \cdot |\mathcal{V}|) / 2. \quad (14)$$

Substituting this into the first equation, we have

$$\langle d_{\text{orig.}} \rangle = \frac{2 \cdot \left(\frac{\langle d_{\text{orig.}} \rangle \cdot |\mathcal{V}|}{2} + |\mathcal{V}_t| + \binom{|\mathcal{V}_{\text{atk}}|}{2} \cdot q \right)}{|\mathcal{V}| + |\mathcal{V}_{\text{atk}}|}. \quad (15)$$

Solving the above for q , we get:

$$q = \frac{\langle d_{\text{orig.}} \rangle - 2 \cdot |\mathcal{V}_t| / |\mathcal{V}_{\text{atk}}|}{|\mathcal{V}_{\text{atk}}| - 1} \quad (16)$$

where $1 < |\mathcal{V}_{\text{atk}}| \leq |\mathcal{V}_t|$. Considering the reasonable range of q , the above equation should satisfy the following range restriction:

$$0 < \frac{\langle d_{\text{orig.}} \rangle - 2 \cdot |\mathcal{V}_t| / |\mathcal{V}_{\text{atk}}|}{|\mathcal{V}_{\text{atk}}| - 1} \leq 1. \quad (17)$$

To satisfy $q \leq 1$, we have

$$\langle d_{\text{orig.}} \rangle - 2 \cdot |\mathcal{V}_t| / |\mathcal{V}_{\text{atk}}| \leq |\mathcal{V}_{\text{atk}}| - 1. \quad (18)$$

Rearranging this inequality:

$$(|\mathcal{V}_{\text{atk}}|)^2 - (1 + \langle d_{\text{orig.}} \rangle) \cdot |\mathcal{V}_{\text{atk}}| + 2 \cdot |\mathcal{V}_t| \geq 0. \quad (19)$$

Providing that $(1 + \langle d_{\text{orig.}} \rangle)^2 - 8|\mathcal{V}_t| \leq 0$, we have the following bound for $|\mathcal{V}_t|$,

$$|\mathcal{V}_t| \geq (\langle d_{\text{orig.}} \rangle + 1)^2 / 8. \quad (20)$$

Furthermore, to satisfy $q > 0$, we have

$$|\mathcal{V}_{\text{atk}}| > 2|\mathcal{V}_t| / \langle d_{\text{orig.}} \rangle. \quad (21)$$

Consider $|\mathcal{V}_{\text{atk}}| \leq |\mathcal{V}_t|$, we have

$$\langle d_{\text{orig.}} \rangle \geq 2. \quad (22)$$

It should be noted that Eq. (20), Eq. (21) and Eq. (22) are easily satisfied in realistic graph learning scenarios, as can be verified in the experimental setup. \square

B.2 Proof of Proposition 2

First, we provide a brief description of Definition 3. Consider the first layer of aggregation operations as follows:

$$\widehat{\mathbf{h}}_v = \text{AGGREGATE}(\{\mathbf{x}'_u, \forall u \in \mathcal{N}(v)\}), \quad (23)$$

so, we have:

$$\mathbb{E}[\widehat{\mathbf{h}}_v] = \mathbb{E}[\text{AGGREGATE}(\{\mathbf{x}'_u, \forall u \in \mathcal{N}(v)\})]. \quad (24)$$

Since $\text{AGGREGATE}(\cdot)$ is linear, we have:

$$\mathbb{E}[\widehat{\mathbf{h}}_v] = \text{AGGREGATE}(\{\mathbb{E}[\mathbf{x}'_u], \forall u \in \mathcal{N}(v)\}). \quad (25)$$

Considering $\mathbb{E}[\mathbf{x}'_u] = \mathbf{x}_u$ and hence have:

$$\mathbb{E}[\widehat{\mathbf{h}}_v] = \text{AGGREGATE}(\{\mathbf{x}_u, \forall u \in \mathcal{N}(v)\}) = \mathbf{h}_v. \quad (26)$$

Therefore, provided that the perturbed node features remain unbiased, the subsequent aggregator also provides an unbiased estimate. Since our feature crafting method follows a uniform distribution, it does not destroy the unbiased property.

Then, we analyze the variance increment caused before and after the attack under the consideration of the two-layer aggregation operation. For the *before attack*, consider the mean aggregation function, the first layer embedding of the target node v_t is:

$$\widehat{\mathbf{h}}_{v_t}^{(1)} = \frac{1}{|\mathcal{N}(v_t)| + 1} (\mathbf{x}'_{v_t} + \sum_{i \in \mathcal{N}(v_t)} \mathbf{x}'_i), \quad (27)$$

where \mathbf{x}'_{v_t} denotes the perturbed node feature vector of node v_t , $\mathcal{N}(v_t)$ denotes the set of genuine neighbor nodes of v_t . Further, consider $\text{Cov}(\mathbf{x}'_{v_t}, \mathbf{x}'_i) = 0$, then the variance of $\widehat{\mathbf{h}}_{v_t}^{(1)}$ is

$$\text{Var}(\widehat{\mathbf{h}}_{v_t}^{(1)}) = \frac{1}{(|\mathcal{N}(v_t)| + 1)^2} (\text{Var}(\mathbf{x}'_{v_t}) + \sum_{i \in \mathcal{N}(v_t)} \text{Var}(\mathbf{x}'_i)) \quad (28)$$

Assuming that the variance of the node features perturbed by the LDP mechanism is σ^2 , then

$$\text{Var}(\widehat{\mathbf{h}}_{v_t}^{(1)}) = \frac{\sigma^2 + |\mathcal{N}(v_t)| \cdot \sigma^2}{(|\mathcal{N}(v_t)| + 1)^2} = \frac{\sigma^2}{|\mathcal{N}(v_t)| + 1}. \quad (29)$$

Further, considering the second layer of aggregation, we have

$$\widehat{\mathbf{h}}_{v_t}^{(2)} = \frac{1}{|\mathcal{N}(v_t)| + 1} (\widehat{\mathbf{h}}_{v_t}^{(1)} + \sum_{i \in \mathcal{N}(v_t)} \widehat{\mathbf{h}}_i^{(1)}). \quad (30)$$

Then, its variance is

$$\text{Var}(\widehat{\mathbf{h}}_{v_t}^{(2)}) = \frac{\text{Var}(\widehat{\mathbf{h}}_{v_t}^{(1)})}{|\mathcal{N}(v_t)| + 1} = \frac{\sigma^2}{(|\mathcal{N}(v_t)| + 1)^2}. \quad (31)$$

For *after attack*, the first layer embedding of the node v_t is:

$$\bar{\mathbf{h}}_{v_t}^{(1)} = \frac{1}{|\mathcal{N}(v_t)| + 2} (\mathbf{x}'_{v_t} + \mathbf{x}'_{v_{\text{atk}}} + \sum_{i \in \mathcal{N}(v)} \mathbf{x}'_i) \quad (32)$$

Assuming that the variance of the malicious node features is σ_{atk}^2 , then

$$\text{Var}(\bar{\mathbf{h}}_{v_t}^{(1)}) = \frac{(|\mathcal{N}(v)| + 1) \cdot \sigma^2 + \sigma_{\text{atk}}^2}{(|\mathcal{N}(v)| + 2)^2}. \quad (33)$$

Table 3: Comparison of accuracy before and after the attack across different GNN models. ↓ represents accuracy drop.

GNN	Meth.	Cora	CiteSeer	Pubmed	LastFM	Twitch	Facebook
GCN	PM	77.3	64.5	79.1	82.5	53.4	88.4
	PM [†]	62.1 ↓	60.4 ↓	74.3 ↓	77.2 ↓	51.6 ↓	83.2 ↓
SAGE	PM	75.3	64.9	76.7	80.5	54.7	83.7
	PM [†]	65.6 ↓	59.1 ↓	70.6 ↓	71.4 ↓	52.8 ↓	76.9 ↓
GAT	PM	78.9	65.3	77.5	81.0	56.2	83.2
	PM [†]	64.4 ↓	60.6 ↓	75.1 ↓	62.9 ↓	54.5 ↓	81.0 ↓

[†] represents the accuracy after untargeted attack.

In addition, the first layer embedding of the node v_{atk} is:

$$\bar{\mathbf{h}}_{v_{\text{atk}}}^{(1)} = \frac{1}{q \cdot (|\mathcal{V}_{\text{atk}}| - 1) + 2} \left(\mathbf{x}'_{v_t} + \mathbf{x}'_{v_{\text{atk}}} + \sum_{i=1}^{q \cdot (|\mathcal{V}_{\text{atk}}| - 1)} \mathbf{x}'_i \right). \quad (34)$$

Let $N_{\text{atk}} = q \cdot (|\mathcal{V}_{\text{atk}}| - 1)$, then, the variance of $\bar{\mathbf{h}}_{v_{\text{atk}}}^{(1)}$ is

$$\text{Var}(\bar{\mathbf{h}}_{v_{\text{atk}}}^{(1)}) = \frac{\sigma^2 + (N_{\text{atk}} + 1) \cdot \sigma_{\text{atk}}^2 + 2 \sum_{i=1}^{N_{\text{atk}}} \text{Cov}(\mathbf{x}'_{v_t}, \mathbf{x}'_i)}{(N_{\text{atk}} + 2)^2}. \quad (35)$$

Further, considering the second layer of aggregation, we have

$$\bar{\mathbf{h}}_{v_t}^{(2)} = \frac{1}{|\mathcal{N}(v_t)| + 2} \left(\bar{\mathbf{h}}_{v_t}^{(1)} + \bar{\mathbf{h}}_{v_{\text{atk}}}^{(1)} + \sum_{i \in |\mathcal{N}(v_t)|} \bar{\mathbf{h}}_i^{(1)} \right) \quad (36)$$

Then, its variance is

$$\text{Var}(\bar{\mathbf{h}}_{v_t}^{(2)}) = \frac{(|\mathcal{N}(v_t)| + 1) \cdot \text{Var}(\bar{\mathbf{h}}_{v_t}^{(1)}) + \bar{\mathbf{h}}_{v_{\text{atk}}}^{(1)}}{(|\mathcal{N}(v_t)| + 2)^2}. \quad (37)$$

To summarize, we have

$$\mathbb{E}[\Delta_{\text{Var}}] = \mathbb{E} \left[\text{Var}(\bar{\mathbf{h}}_{v_t}^{(2)}) - \text{Var}(\bar{\mathbf{h}}_{v_t}^{(2)}) \right] \quad (38)$$

$$= \frac{(|\mathcal{N}(v_t)| + 1)^2 \cdot \sigma^2}{(|\mathcal{N}(v_t)| + 2)^4} + \frac{\sigma^2 + (N_{\text{atk}} + 1) \cdot \sigma_{\text{atk}}^2}{(|\mathcal{N}(v_t)| + 2)^2 \cdot (N_{\text{atk}} + 2)^2} \quad (39)$$

$$- \frac{\sigma^2}{(|\mathcal{N}(v_t)| + 1)^2} > 0. \quad (40)$$

The above process quantifies the error estimation under the consideration of identical feature crafting approach, which is conducive to maximize the damage to the data utility of the target node.

B.3 Global Error Increment

In this section, we derive the global error increment $\Delta\Psi$ that quantifies the distortion of node representations caused by the attack. The global error energy function before the attack is defined as:

$$\Psi(\mathcal{G}) = \sum_{v \in \mathcal{V}} \|\mathbf{h}_v^{(K)} - \mathbf{x}_v\|^2 + \lambda \sum_{(u,v) \in \mathcal{E}} \|\mathbf{h}_u^{(K)} - \mathbf{h}_v^{(K)}\|^2,$$

where $\mathbf{h}_v^{(K)}$ represents the feature of node v after K -layer aggregation, and \mathbf{x}_v is the original feature of node v . The parameter λ controls the balance between node similarity and edge similarity. After the attack, we define the graph as \mathcal{G}' and the corresponding error energy function as $\Psi(\mathcal{G}')$. The global error increment is

$$\Delta\Psi = \Psi(\mathcal{G}') - \Psi(\mathcal{G}).$$

Given the perturbations, the expected deviation of the global error increment is derived as:

$$\mathbb{E}[\Delta\Psi] = \left(1 + \frac{\lambda q}{1 - \lambda q} \right) \cdot |\mathcal{V}_{\text{atk}}| \cdot B^2 \cdot K.$$

This expression shows that the attack's impact on the graph is proportional to the number of fake nodes $|\mathcal{V}_{\text{atk}}|$, the strength of the perturbation B^2 , and the number of aggregation layers K . The term $\left(1 + \frac{\lambda q}{1 - \lambda q} \right)$ reflects the influence of aggregation, where λ controls the trade-off between node and edge feature similarities, and q is the probability of connectivity between fake nodes.

C ADDITIONAL DETAILS OF EXPERIMENTS

C.1 Hyperparameters

For the hyperparameters, we conduct a grid search to identify the optimal choice. The learning rate is selected from the set $\{10^{-1}, 10^{-2}, 10^{-3}\}$, the weight decay is chosen from $\{10^{-3}, 10^{-4}, 10^{-5}, 0\}$, the dropout rate is sampled from $\{10^{-1}, 10^{-2}, 10^{-3}, 0\}$. All models are trained for up to 300 epochs using the Adam optimizer [19], and the best model was selected for testing based on validation loss.

C.2 More Details About Validation

In our experiments, we comprehensively evaluate the effectiveness of our attack methodology by considering two fundamental tasks: node classification and link prediction. Furthermore, we assess our attack under both *targeted* and *untargeted* attack setups.

- *Targeted Attack.* This setup aims to evaluate the attacker's impact on a specific target node $v_t \in \mathcal{V}_t$. The goal is to assess how the attack alters the performance related to this particular node, such as its classification accuracy or predictive capabilities.
- *Untargeted Attack.* In contrast, the untargeted attack measures the overall effectiveness of the attack across the entire graph, without focusing on any specific node. This setup aims to explore the broader effects of the attack on the graph as a whole.

C.3 More Experimental Results and Analysis

C.3.1 Attack effectiveness under different GNN models. In Sec. 5.3, we analyze the attack performance under different GNN models. Tables 2 and 3 present the experimental results under the targeted and untargeted attack settings, respectively. The results from the untargeted attack setting provide further support to the analysis in Section 5.3, demonstrating the consistency and effectiveness of our attack across different GNN models.

C.3.2 On the analysis of η_1 and η_2 . In Sections 5.4.1 and 5.4.2, we analyze the effects of the parameters η_1 and η_2 on the performance of our attack. Fig. 5 presents the experimental results under the targeted attack setting, while Fig. 9 shows the experimental results for the untargeted attack setting. The experiments conducted under the untargeted attack setting further corroborate the analysis of the results presented in Sec 5.4.1 and 5.4.2, providing additional evidence for the effectiveness of our attack under different configurations.

C.3.3 On the analysis of the parameter K . In Sec. 5.4.3, we validate the change in accuracy before and after the attack with different values of $K \in \{0, 2, 4, 8, 16\}$. The results are shown in Fig. 6. The parameter K represents the number of steps in the second phase (calibration) of the locally private graph learning protocol II. As seen in Fig. 6, the accuracy before and after the attack exhibits an *increasing and then decreasing* trend with the increasing K .

The initial increase in accuracy can be attributed to the calibration of noise by neighbor aggregation. As defined in Def. 4 [43], the error deviation is measured by the embedding disparity of v before and after the perturbation, which is influenced by the neighborhood size $|\mathcal{N}(v)|$ as given in Eq. (41). The estimation error decreases at a rate proportional to the square root of the node degree. Therefore, increasing K expands the neighborhood that can be aggregated, facilitating better noise calibration, which improves the accuracy.

Definition 4 (Error Deviation). Given the aggregator function for the first layer and $\delta > 0$, with probability at least $1 - \delta$ for any node v , we have the following bound on the aggregation result:

$$\max_{i \in \{1, 2, \dots, d\}} \left| (\hat{\mathbf{h}}_v^{(1)})_i - (\mathbf{h}_v^{(1)})_i \right| = O \left(\frac{\sqrt{d \log(d/\delta)}}{\epsilon \sqrt{|\mathcal{N}(v)|}} \right), \quad (41)$$

where $\hat{\mathbf{h}}_v^{(1)}$ is the perturbed aggregated embedding of node v at the first layer, and $\mathbf{h}_v^{(1)}$ is the true aggregation without noise. $\mathcal{N}(v)$ is the set of neighbors of node v .

However, the later decrease in accuracy when K becomes large is due to over-smoothing [25]. Over-smoothing occurs as node features tend to converge towards a common value with an increasing number of aggregation layers, resulting in diminished utility for graph learning tasks. This is a well-known issue in GNNs [18]. Importantly, our attack method significantly reduces the accuracy of the original method across all values of K . This validates the generalization and effectiveness of our attack approach, showing that it is robust to different values of K and can significantly degrade the performance of the locally private graph learning protocol.

C.3.4 Metrics for graph homophily analysis. In the defense experiments in Sec. 6.1, we consider two homology metrics [4]: *node-centered homology* and *edge-centered homology*. The former evaluates the similarity between a target node and its neighboring nodes, while the latter assesses the similarity between nodes connected by a target edge. They are defined as follows:

Definition 5 (Node-Centric Homophily). The homophily of a node v is quantified by measuring the similarity between its feature vector and the aggregated embedding of its neighboring nodes:

$$h_v = \text{sim}(r_v, \mathbf{x}_v), \quad r_v = \sum_{j \in \mathcal{N}(v)} \frac{1}{\sqrt{d_j} \sqrt{d_v}} \mathbf{x}_j, \quad (42)$$

where \mathbf{x}_v denotes the feature vector of node v , d_v represents the degree of v , and $\text{sim}(\cdot)$ is a similarity metric, e.g., cosine similarity.

Definition 6 (Edge-Centric Homophily). The homophily for an edge (v, u) can be defined as follows:

$$h_e = \text{sim}(\mathbf{x}_v, \mathbf{x}_u), \quad (43)$$

where \mathbf{x}_v denotes the feature vector of node v and $\text{sim}(\cdot)$ represents a distance metric, e.g., cosine similarity.

D MORE DETAILS ON RELATED WORK

Recently, locally private graph learning protocols have attracted significant attention in the security research community [15, 28, 31, 38, 40, 43, 58, 66]. Sajadmanesh and Gatica-Perez [43] extended the one-bit mechanism [7] to multidimensional space and introduced




Figure 9: Variation in untargeted attack accuracy at different η_1 and η_2 values. The impact ratio represents the proportion of accuracy reduction compared to the normal accuracy.

the multi-bit mechanism (MB) for protecting node features, which serves as a foundation for private graph learning. Lin et al. [31] further refined the locally private graph learning protocol based on the MB mechanism [43]. Additionally, Pei et al. [39] investigated privacy-preserving learning for decentralized graphs using the piecewise mechanism (PM) [48], while Li et al. [28] extended the original square wave mechanism (SW) [29] to enhance the utility of locally private graph learning. Some studies [15, 23, 40, 58, 66] have also explored link protection. These approaches leverage the benefits of LDP [8, 54] in safeguarding data privacy and ensuring strict privacy guarantees for user data. Moreover, during server-side privacy graph learning, GNN's multi-layer message passing mechanism [1, 10] effectively reduces the estimation error in node embeddings after perturbation and aids in the calibration of noisy data, thus maintaining high utility for the local privacy graph learning task. However, despite these advantages, such protocols may be vulnerable to data poisoning attacks, a threat that has not been considered in previous research. Identifying and addressing these threats is crucial for ensuring the robustness and security of private graph learning frameworks. This work introduces the first data poisoning attack targeting locally private graph learning protocols.