

# Inference Attacks Against Graph Generative Diffusion Models

Xiuling Wang  
Hong Kong Baptist University  
xiulingwang@hkbu.edu.hk

Xin Huang  
Hong Kong Baptist University  
xinhuang@comp.hkbu.edu.hk

Guibo Luo  
Peking University  
luogb@pku.edu.cn

Jianliang Xu\*  
Hong Kong Baptist University  
xujl@comp.hkbu.edu.hk

## Abstract

Graph generative diffusion models have recently emerged as a powerful paradigm for generating complex graph structures, effectively capturing intricate dependencies and relationships within graph data. However, the privacy risks associated with these models remain largely unexplored. In this paper, we investigate information leakage in such models through three types of black-box inference attacks. First, we design a graph reconstruction attack, which can reconstruct graphs structurally similar to those training graphs from the generated graphs. Second, we propose a property inference attack to infer the properties of the training graphs, such as the average graph density and the distribution of densities, from the generated graphs. Third, we develop two membership inference attacks to determine whether a given graph is present in the training set. Extensive experiments on three different types of graph generative diffusion models and six real-world graphs demonstrate the effectiveness of these attacks, significantly outperforming the baseline approaches. Finally, we propose two defense mechanisms that mitigate these inference attacks and achieve a better trade-off between defense strength and target model utility than existing methods. Our code is available at <https://zenodo.org/records/17946102>.

## 1 Introduction

Many real-world systems, such as social networks, biological networks, and information networks, can be represented as graphs. Graph learning is crucial for analyzing these systems because of its ability to model and analyze complex relationships and interactions within the data.

Graph generation, a critical task in graph learning, focuses on creating graphs that accurately reflect the underlying structure of graph data. These models have diverse applications, such as recommender systems [32, 68], social network analysis [38, 40], and drug discovery [63, 75]. Diffusion models have recently gained significant attention as a prominent class

of generative models, which work through two interconnected processes. A forward process gradually adds noise to data until it conforms to a predefined prior distribution (e.g., Gaussian). A corresponding reverse process uses a trained neural network to progressively denoise the data, effectively reversing the forward process and reconstructing the original data distribution. Given the success of diffusion models in image generation [26], there has been growing interest in applying these techniques to graph generation [10, 17, 39, 79, 81].

While graph generative diffusion models (GGDMs) are capable of producing various graphs, they often require large training datasets for robust generation. However, these datasets may contain sensitive or confidential information. For example, social graphs can expose private relationships, as seen in cases such as Facebook’s famous Cambridge Analytica scandal; graphs depicting protein-protein interactions or gene regulatory networks may include proprietary research data, any leakage of this information could compromise competitive advantages in biotechnology research; graphs in healthcare datasets can model relationships between patients, treatments, and healthcare providers, where the disclosure of such information could compromise patient confidentiality and violate regulations such as HIPAA [44]. Therefore, this paper investigates the critical question: *how much information about the training data can be inferred from GGDMs?*

**Challenges.** Recent research has uncovered several types of attacks that can infer sensitive information from the training data of graph learning models, including membership inference [16, 23, 24, 55, 70–72, 74], attribute inference [16, 19, 83], property inference [60, 69, 84, 86], and graph reconstruction attacks [55, 85, 87]. Most of these attacks are based on model’s final output (probability vectors) for node/graph classification [19, 23, 24, 60, 69, 72, 84], node/graph embeddings [16, 16, 55, 70, 71, 71, 83, 86, 87], or other features such as model gradients [74, 85]. These approaches are not well-suited for GGDMs for two reasons: *First*, attack features derived from probability vectors or node/graph embeddings are unavailable in GGDMs that generate graphs, making existing attacks infeasible in black-box setting. *Second*, the above at-

\*Jianliang Xu is the corresponding author.

tacks rely on a one-to-one mapping between an input graph and an output to form a (feature, label) pair, where label denotes membership or a specific property. GGDMs lack this mapping as the graphs are generated from a set of training samples, making ground-truth labels unavailable and existing attacks inapplicable. These differences highlight the unique challenge of attacking GGDMs: constructing meaningful features from generated graphs without one-to-one mappings.

On the other hand, recent preliminary studies have begun to investigate the privacy vulnerabilities of image and text-to-image generative diffusion models under various privacy inference attacks. These include membership inference attacks [12, 14, 50], which aim to determine whether a particular data sample is in the training set; data reconstruction attacks, which attempt to recover the training images [12]; and property inference attacks [42], which seek to infer the portion of training images with a specific property. However, all of the above works focus only on image or text-to-image diffusion models and cannot be directly applied to GGDMs due to the unique structural characteristics of graph data. Specifically, graph data exhibit variable sizes and topologies, permutation-invariant, and encode sensitive information through higher-order structural patterns rather than fixed spatial features.

**Our Contributions.** We initiate a systematic investigation into the privacy risks of GGDMs by exploring three types of inference attacks. First, we explore *graph reconstruction attack* (GRA), which attempts to infer graph structures within model’s training set. For example, if the target graph is from a medical database, a reconstructed graph could enable an adversary to gain knowledge of sensitive relationships between patients, health records, and treatments. Second, we introduce *property inference attack* (PIA), which leverages generated graphs to infer statistical properties of the training graphs, such as average graph density or the proportion of graphs within specific density ranges. Revealing these properties may violate the intellectual property (IP) of the data owner, particularly in domains like molecular or protein graphs from biomedical companies. Third, we investigate *membership inference attack* (MIA), which aims to determine whether a given graph is present in GGDM’s training set. For example, in a collaborative training setting where each data owner possesses graphs from different biological companies, with each graph representing a proprietary product, an adversary may attempt to infer whether a specific product is included in another owner’s training set, thereby compromising IP.

To the best of our knowledge, this is the first work to explore the privacy leakage of GGDMs. Overall, we make the following contributions in this paper:

- We develop a novel graph reconstruction attack by aligning each generated graph with its closest counterpart in the generated graph set to identify overlapping edges. These overlapping edges are then considered part of a graph from the target model’s training set. We evaluate the effectiveness of our proposed attack on six real-world graph datasets and

	Attack type			Attack setting		Domain
	DR	PIA	MIA	B-box	W-box	
[12]	✓		✓		✓	Image or text-to-image
[42]		✓		✓		
[36, 49, 76, 82]			✓	✓		
[14, 28, 35, 62, 80]			✓		✓	Graph
[12, 15, 43, 50]						
Ours	✓	✓	✓	✓		

Table 1: Comparison between the existing works on inference attacks against generative diffusion models. "DR", "PIA", and "MIA" refer to data reconstruction, property inference, and membership inference attacks, respectively. "B-box" and "W-box" represent black-box and white-box, respectively.

three state-of-the-art GGDMs. The results show that the attack achieves an F1 score of up to 0.99, with up to 36% of the original training graphs being exactly recovered.

- We launch our property inference attack using a simple yet efficient method that directly calculates property values from the generated graphs. Extensive experiments show that the proposed attack can accurately infer the statistical properties of the training graphs. For example, on the IMDB-MULTI dataset, the difference between the actual and inferred average graph degrees can be as small as 0.005.
- We design our membership inference attacks by employing the shadow-model-training technique, where we train MLP attack models based on two factors: (1) different similarity levels between the generated graphs and their corresponding training graphs for member and non-member graphs; and (2) different similarity levels within the generated graphs from member and non-member graphs. Experimental results show that our attacks can achieve an AUC of up to 0.999 when shadow and target graphs are drawn from the same dataset, and 0.895 when drawn from different datasets.
- To mitigate the inference attacks, we propose two defense mechanisms that introduce noise into either the training or generated graphs of the target GGDM, thereby altering the model’s outputs. Notably, we limit perturbations to the least significant edges and non-edges (by flipping them), to minimize the impact on the target model’s utility. Empirical evaluations show that our approaches render the attacks ineffective while achieving a better defense–utility trade-off than the two existing methods.

## 2 Related Work

**Generative diffusion models for graphs.** Generative diffusion models have recently gained significant attention as a powerful paradigm for deep graph generation, aiming to learn the underlying graph distribution and synthesize novel graphs. The existing graph generative diffusion models (GGDMs) can be broadly categorized into three classes [17, 81]: (1) *Score-based Generative Models (SGM)* [13, 48] that employ a score function to represent the probability distribution of the data; (2) *Denosing Diffusion Probabilistic Models (DDPMs)* [21]

that add discrete Gaussian noise to the graph with Markov transition kernels [9] and then train a neural network to predict the added noise to recover the original graph, [67] adds discrete noise instead of continuous Gaussian, and (3) *Stochastic Differential Equation-based Models (SDEs)* [30, 33, 41] that characterize the development of a system over time under the influence of random noise. SGM and DDPM leverage the score-matching idea and non-equilibrium thermodynamics, respectively, to learn different reverse functions of the diffusion process, while SDE generalizes the discrete diffusion steps into continuous scenarios and further models the diffusion process with stochastic differential equations [17]. We refer the readers to comprehensive surveys on GGDMs [17, 81].

#### Inference attacks against generative diffusion models.

Few studies [12, 42, 58, 64, 65] have investigated the privacy vulnerabilities of generative diffusion models against privacy inference attacks. Van et al. [65] focus on probabilistic deep generative models such as variational autoencoders and formulate the concept of "memorization score" by measuring the impact of removing an observation on a given model. Carlini et al. [12] consider image diffusion models and design three attacks: data extraction, data reconstruction, and membership inference attack. Somepalli et al. [58] focus on text-to-image diffusion models and analyze the data duplication problem in these models. Luo et al. [42] introduce a black-box property inference attack that aims to infer the distribution of specific properties from generated images. Several works focus on membership inference attacks against diffusion models [14, 15, 28, 35, 36, 36, 43, 49, 50, 62, 76, 80, 82]. Specifically, [36, 49, 76, 82] describe black-box attacks by analyzing the generated images; [14, 28, 35, 62, 80] and [12, 15, 43] design white-box attacks that rely on posterior estimation errors and model losses, respectively; and Pang et al. [50] introduce a white-box attack that leverages gradients at each timestep. However, all of these attack models focus on image or text-to-image diffusion models, which cannot be directly applied to GGDMs. Table 1 summarizes the main differences between our work and existing studies on inference attacks against generative diffusion models.

**Inference attacks on graph data.** Recent research has uncovered several types of attacks that can infer sensitive information from the training graph data. These attacks can be categorized into four types: (1) membership inference attacks (MIAs) [16, 23, 24, 55, 70–72, 74], seeking to determine whether a specific graph sample is part of the training dataset; (2) attribute inference attacks (AIAs) [16, 19, 83], aiming to infer the sensitive attributes within the training graphs; (3) property inference attacks (PIAs) [60, 69, 84, 86], trying to infer the sensitive properties of the training graphs; and (4) graph reconstruction attacks [55, 85, 87], attempting to reconstruct the training graphs. However, most of these attacks rely on the model’s final output (probability vectors) for node/graph classification [19, 23, 24, 60, 69, 72, 84], node/graph embeddings [16, 16, 55, 70, 71, 71, 83, 86, 87], or some other

features like model gradients [74, 85]. Consequently, these approaches are not well-suited for GGDMs due to the unique forward and reverse processes in diffusion models as well as the different types of outputs they produce.

### 3 Generative Diffusion Models on Graphs

Graph generation models aim to generate new graph samples resembling a given dataset. Among these, diffusion-based models have become increasingly popular. They gradually introduce noise into data until it conforms to a prior distribution [13, 21, 30, 33, 41, 48, 67].

Generally, existing generative diffusion models on graphs include two processes: (1) the *forward process*, which progressively degrades the original data into Gaussian noise [26], and (2) the *reverse process*, which gradually denoises the noisy data back to its original structure using transition kernels. Based on how the forward and reverse processes are designed, existing generative diffusion models on graphs can be broadly categorized into three classes: *Score-Based Generative Models* (SGM) [13, 48], *Denosing Diffusion Probabilistic Models* (DDPM) [21, 67], and *Stochastic Differential Equations* (SDE) [30, 33, 41]. Next, we briefly describe the forward and reverse processes for the three types of graph generative diffusion models. Note that “synthetic graphs” and “generated graphs” are used interchangeably throughout this paper to refer to the graphs produced by the diffusion models.

**SGM-based graph generation.** The forward process injects Gaussian noise of varying intensity as the perturbation into the original graph. A noise-conditional score network is trained to represent the gradient of the conditional probability density function of the data under varying noise levels. Specifically, given a probability density function  $p(x)$  and the score function  $\nabla_x \log p(x)$ , SGM aims to estimate the data score function in the forward process. The training objective of the score network is given by:

$$\mathcal{L}(\theta) = \min_{\theta} \mathbb{E}_{t \sim \mathcal{U}(1, T), x_0 \sim p(x_0), \epsilon \sim \mathcal{N}(0, \mathbf{I})} \left[ \lambda(t) \left\| \nabla_{x_t} \log p_{0t}(x_t | x_0) - \sigma_t s_{\theta}(x_t, t) \right\|^2 \right], \quad (1)$$

where  $\mathbb{E}$  is the expectation,  $\mathcal{U}(1, T)$  is a uniform distribution over the time set  $\{1, 2, \dots, T\}$ ,  $\epsilon$  is the noise vector, and  $\lambda(t)$  is a positive weighting function. The variable  $x_0$  refers to the original data before noise is added, with its probability density function denoted as  $p(x_0)$ .  $p_{0t}(x_t | x_0)$  is the score function of  $x_t$ .  $\sigma_t$  is the Gaussian noise at  $t$ . The score network  $s_{\theta}$  with parameter  $\theta$ , predicts the noise  $\sigma_t$  based on  $x_t$  and  $t$ , where  $x_t$  is the noisy version of  $x_0$  after adding noise at  $t$ .

In the reverse process, after obtaining the trained conditional score model, synthetic graphs are generated using noise-conditional score networks, such as the Score Matching with Langevin Dynamics model [59] that leverages the learned score models to reconstruct graph data from noise.

**DDPM-based graph generation.** In the forward process, the original data undergoes perturbation with Gaussian




Figure 1: Inference attacks against graph generative diffusion models. The attacker inputs a set of shadow graphs into the target graph generative diffusion model  $\Phi$  or directly executes  $\Phi$  through an API or online marketplaces to obtain a large number of synthetic graphs  $\hat{G}$ , aiming to infer the sensitive information about the training data of  $\Phi$ . In this paper, we investigate three types of inference attacks: (1) reconstructing the graph structures in  $\Phi^{\text{train}}$ ; (2) inferring the properties of  $\Phi^{\text{train}}$ , such as the graph density of  $\Phi^{\text{train}}$ ; and (3) determining the membership of a given target graph.

noise using a fixed number. Specifically, given a probability density function  $x \sim p(x)$ , the forward process generates the noisy  $\hat{x}$  with Markov transition kernels [9, 26] as  $p(x_t|x_{t-1}) = \mathcal{N}(x_t|\sqrt{1-\beta_t}x_{t-1}, \beta_t\mathbf{I})$ , where  $\beta_t$  is a predefined variance schedule at time step  $t$ .

In the reverse process, a neural network is trained to predict the noise added at each step during the forward pass, ultimately recovering the original data. Following the notations in Equation 1, the optimization objective can be expressed as:

$$\mathcal{L}(\theta) = \min_{\theta} \mathbb{E}_{t \sim \mathcal{U}(1, T), x_0 \sim p(x_0), \varepsilon \sim \mathcal{N}(0, \mathbf{I})} [\lambda(t) \|\varepsilon - \varepsilon_{\theta}(x_t, t)\|^2], \quad (2)$$

where  $\varepsilon_{\theta}$  is a deep neural network with parameter  $\theta$  that predicts the noise vector  $\varepsilon$  given  $x_t$  and  $t$ .

**SDE-based models.** The forward process employs a forward SDE [37] to describe the evolution of a state variable over time to generate the noisy graph. It perturbs data to noise with SDE as  $d_x = f(x, t)d_t + g(t)d_w$ , where  $f(x, t)$  and  $g(t)$  are diffusion and drift functions of the SDE, and  $w$  is a standard Wiener process.

In the reverse process, a reverse SDE is utilized to gradually convert noise to data. This is achieved by estimating the score functions of the noisy data distributions. Using similar notations as in Equation 1, the objective for estimating the score function can be formulated as:

$$\mathcal{L}(\theta) = \min_{\theta} \mathbb{E}_{t \sim \mathcal{U}(1, T), x_0 \sim p(x_0), \varepsilon \sim \mathcal{N}(0, \mathbf{I})} [\lambda(t) \|s_{\theta}(x_t, t) - \nabla_{x_t} \log p_{0t}(x_t|x_0)\|^2], \quad (3)$$

Once the score function at each time step is obtained, the synthetic graphs can be generated with various numerical techniques, such as annealed Langevin dynamics, numerical SDE/ODE solvers, and predictor-corrector methods [79].

## 4 Motivation and Threat Model

In this section, we start by explaining our motivation, and then define the scope and objectives of our problem.

### 4.1 Motivations

Machine learning (ML) models have unlocked a variety of applications, such as data analytics, autonomous systems, and

security diagnostics. However, developing robust models often requires substantial training datasets, extensive computational resources, and significant financial investment, which can be prohibitive for small businesses, developers, and researchers with limited budgets. Consequently, online marketplaces for ML models, such as Amazon Web Services [1], Google AI Hub [3], Modzy [5], Microsoft Azure Cognitive Services [2], and IBM Watson [4], have emerged, facilitating model exchange, customization, and access to various machine-learning-as-a-service (MLaaS) APIs. While beneficial, these platforms raise concerns about the potential exposure of sensitive or proprietary information from the training data. Therefore, in this paper, we investigate the privacy vulnerability of increasingly popular GGDMs.

### 4.2 Threat Model

ML models are susceptible to privacy attacks [29, 53]. These attacks can be categorized into two groups based on the adversary’s target [53]: (1) *Privacy attacks on training data*: the adversary aims to infer sensitive information about the training data. (2) *Privacy attacks on ML models*: the adversary considers the ML models themselves as sensitive, like valuable company assets, and attempts to uncover information regarding the model’s architecture and parameters.

In this paper, we primarily focus on privacy attacks on training data. Our threat model considers an adversary that interacts with a graph generative diffusion model  $\Phi$  to extract the information of the model’s training set  $\Phi^{\text{train}}$ . An overview of the attacks is shown in Figure 1.

**Adversary’s Background Knowledge.** We consider the adversary knowledge  $\mathbb{K}$  along two dimensions:

- *Shadow graph  $G^S$*  (optional): The adversary possesses one or more shadow graphs,  $G^S$ , each with its own structure and node features.  $G^S$  may originate from a different domain than the model’s training set,  $\Phi^{\text{train}}$ , and thus exhibit distinct data distributions. In real-world scenarios, the adversary may have knowledge of a partial graph, which is a subset of the training graphs [16, 23, 69]. We treat this partial graph as a specific instance of a shadow graph.
- *Target model  $\Phi$* : The adversary may have either white-box




Figure 2: Overview of graph reconstruction.

or black-box access to the target model. In the white-box setting, the attacker can access  $\Phi$ 's internal components, such as parameters, gradients, and loss values [12, 16, 50]. In contrast, the black-box setting assumes that the adversary can only interact with  $\Phi$  through its outputs. In this paper, we consider a black-box setting, reflecting real-world scenarios such as public or commercial APIs in MLaaS platforms [23, 54, 56, 86]. Specifically, the adversary can submit their own graphs or random graphs to  $\Phi$  via an API and receive a corresponding set of generated graphs  $\hat{G}$  as output. This black-box setting is considered the most challenging setting for the adversary [23, 54, 56].

**Attack Goal.** We consider three types of inference attacks:

- *Graph reconstruction:* The attacker aims to reconstruct the graph structures within target model's training set  $\Phi^{\text{train}}$ .
- *Property inference:* The attacker attempts to infer predefined properties or aggregate characteristics,  $\mathbb{P}$ , of an individual record or a group within the training set  $\Phi^{\text{train}}$ .
- *Membership inference:* The attacker seeks to determine whether a given graph  $G$  is in the training set  $\Phi^{\text{train}}$ .

## 5 Inference Attacks

In this section, we detail the proposed graph reconstruction, property inference, and membership inference attacks.

### 5.1 Graph Reconstruction Attacks

#### 5.1.1 Attack Overview

Given a target GGDM  $\Phi$ , obtained from an online marketplace or via an API on a MLaaS platform, the attacker's goal is to infer the graph structures in  $\Phi^{\text{train}}$ , which is the training set of  $\Phi$ . Figure 2 illustrates the pipeline of our graph reconstruction attack, and the corresponding pseudo-code can be found in Appendix A. The attacker first generates a set of synthetic graphs  $\hat{G}$  and then reconstructs the graphs by aligning each graph in  $\hat{G}$  with its closest counterpart to identify the edges in the original training graphs. Formally, the graph reconstruction attack can be formulated as follows:

$$f: \hat{G} \rightarrow \Phi^{\text{train}}. \quad (4)$$

Next, we detail the attack steps.

#### 5.1.2 Attack Model

Our attack model includes three steps: graph generation, graph alignment, and edge inference.

**Graph generation.** The attacker inputs a set of shadow graphs,  $G^S$ , into  $\Phi$  or executes  $\Phi$  directly to obtain a large number of synthetic graphs  $\hat{G}$ .

**Graph alignment.** With the generated graphs  $\hat{G}$ , for each generated graph  $g_i$  in  $\hat{G}$ , we first identify the most similar graph  $g_j$  to  $g_i$  within  $\hat{G}$  by using graph alignment techniques [25, 78]. This is based on the assumption that if the target model has memorized a particular graph, it will likely produce multiple similar graphs to this graph. This assumption is verified in image-based diffusion models, where diffusion models memorize individual images from their training data and emit them at generation time [12, 65].

Following the same strategy in [25], we use REGAL for graph alignment, which leverages representation learning to effectively map and match nodes between different graphs. Specifically, given a graph  $g_i$ , REGAL aligns it with other graphs in  $\hat{G}$  through the following three steps: (1) node identity extraction that extracts the structure and attribute-related information for all nodes in  $g_i$  based on the degree distributions and node features; (2) similarity-based node representation by using the low-rank matrix factorization-based approach that leverages a combined structural and attribute-based similarity matrix from step (1); and (3) node representation alignment that greedily matches each node in  $g_i$  to its top- $\alpha$  most similar nodes in other graphs in  $\hat{G}$  with k-d trees. The difference between a pair of node representations,  $(Y_i(u), Y_j(v))$ , in  $g_i$  and  $g_j$  is calculated as:

$$Diff(Y_i(u), Y_j(v)) = \exp^{-\|Y_i(u) - Y_j(v)\|^2}. \quad (5)$$

**Edge inference.** After graph alignment, we get the pairwise alignments within  $\hat{G}$ . Then, for each  $g_i$  in  $\hat{G}$ , we identify its most similar counterpart in  $\hat{G}$  by averaging all the node representation differences between the aligned graphs. Specifically, using  $Diff$  from Equation 5, we identify the most similar counterpart  $\hat{g}_j$  for  $g_i$  among its aligned graphs as:

$$\hat{g}_j = \min_{\hat{g}_j} \frac{1}{|V_i|} \sum_{u \in \{V_i\}} Diff(Y_i(u), Y_j(\hat{u})), \quad (6)$$

where  $u$  is a node in the node set of  $g_i$ , denoted as  $\{V_i\}$ ,  $\hat{u}$  is the aligned node of  $u$  in  $g_j$ .

After obtaining the most similar counterpart  $\hat{g}_j$  for each  $g_i$  in  $\hat{G}$ , and the average node representation difference between  $g_i$  and  $\hat{g}_j$ , namely  $D(g_i, \hat{g}_j) = \frac{1}{|V_i|} \sum_{u \in \{V_i\}} Diff(Y_i(u), Y_j(\hat{u}))$ , we pick the graph pairs with top-k% smallest differences to determine the graph structure in  $\Phi^{\text{train}}$ . Empirically, we set  $k$  to 10%. Subsequently, we reconstruct the structure of a training graph from these selected aligned graph pairs. For each selected graph pair  $(g_i, \hat{g}_j)$ , we perform the intersection operation on the edge sets of  $g_i$  and  $\hat{g}_j$ , denoted as  $E_i$  and  $\hat{E}_j$ , to reconstruct the graph  $g_i^{\text{rec}}$  as:

$$\{V_i^{\text{rec}}\} = \{V_i\}, \{E_i^{\text{rec}}\} = \{E_i\} \cap \{\hat{E}_j\}, \quad (7)$$

where  $\{V_i^{\text{rec}}\}$  and  $\{E_i^{\text{rec}}\}$  are the node and edge sets of the reconstructed graph, respectively. We also experimented with

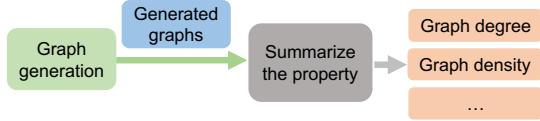


Figure 3: Overview of property inference attack.

using the union operation on the edge sets of  $g_i$  and  $\hat{g}_j$ , but the results showed that the intersection operation performs better than the union operation. Therefore, we use the intersection operation in our attack model. Finally, we get a set of reconstructed graphs from the selected aligned graph pairs.

## 5.2 Property Inference Attacks

### 5.2.1 Attack Overview

Following the same attack setting, where the attacker obtains the target GGDM  $\Phi$  from an online marketplace or a MLaaS API, the attacker’s goal is to infer the graph properties of  $\Phi^{\text{train}}$ . In this paper, we focus on two types of properties: (1) the average statistical properties of the training graphs in  $\Phi^{\text{train}}$ , such as graph density and average node degree; and (2) the distribution of training graphs across different property ranges in  $\Phi^{\text{train}}$ . Figure 3 illustrates the pipeline of our property inference attack, and the corresponding pseudo-code can be found in Appendix A. The attacker first generates a set of synthetic graphs,  $\hat{G}$ , and then directly calculates the property values from  $\hat{G}$ . Formally, the property inference attack can be formulated as follows:

$$f : \hat{G} \rightarrow \mathbb{P}(\Phi^{\text{train}}), \quad (8)$$

where  $\mathbb{P}$  represents the property, including both the average statistical properties and the distribution of training graphs across different property ranges.

### 5.2.2 Attack Model

Our attack model includes two steps: graph generation and property inference.

**Graph generation.** The attacker inputs a set of shadow graphs,  $G^S$ , into  $\Phi$  or executes  $\Phi$  directly to obtain a large number of synthetic graphs  $\hat{G}$ .

**Property inference.** For both property types, the attacker computes the property value over all graphs in  $\hat{G}$  to approximate that of  $\Phi^{\text{train}}$ . The process can be formulated as:

$$\mathbb{P}(\Phi^{\text{train}}) \leftarrow \mathbb{P}(\hat{G}). \quad (9)$$

## 5.3 Membership Inference Attacks

### 5.3.1 Attack Overview

Given the GGDM  $\Phi$  obtained from an online marketplace or through an API, and a target graph  $G$ , the attacker’s goal is to determine whether  $G$  is in  $\Phi$ ’s training set,  $\Phi^{\text{train}}$ . In this attack, we have the following assumptions. First, the

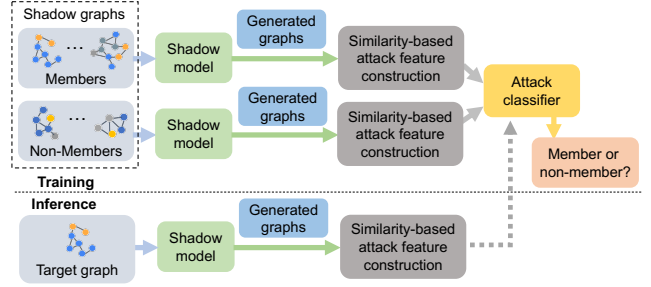


Figure 4: Overview of membership inference attack.

attacker has background knowledge of shadow graphs  $G^S$ , which contain their own nodes and features. These shadow graphs may come from a different domain than  $\Phi^{\text{train}}$ . A special case occurs when  $G^S$  contains partial subgraphs of  $\Phi^{\text{train}}$ , which is plausible in real-world applications. For example, an online marketplace may release a portion of the training set. Second, following [23, 56, 69], we assume the attacker can train a shadow model  $\Phi^S$  using the same service (e.g., Google AI Hub) employed for the target model. Figure 4 illustrates the pipeline of our membership inference attacks, and the corresponding pseudo-code can be found in Appendix A. The attacker first trains  $\Phi^S$  using a subset of graphs from  $G^S$ . Then,  $G^S$  is fed into  $\Phi^S$  to generate a set of synthetic graphs  $\hat{G}$ . An attack classifier is subsequently trained by analyzing the similarity between  $\hat{G}$  and  $G^S$  or the similarity within all the graphs in  $\hat{G}$ . Formally, given a target graph  $G$ , the membership inference attack can be formulated as follows:

$$f : (G, G^S) \rightarrow y_G \in \{0, 1\}, \quad (10)$$

where  $y_G$  is the membership label of  $G$ , with 0 (or 1) indicating the absence (or presence) in  $\Phi^{\text{train}}$ . We next detail the attack.

### 5.3.2 Attack Model

Our attack model includes four steps: shadow model training and graph generation, attack feature construction, attack model training, and membership inference.

**Shadow model training and graph generation.** First, the attacker trains a shadow model  $\Phi^S$  on a subset of graphs from  $G^S$ , referred to as member shadow graphs, to mimic the behavior of the target model. The remaining graphs in  $G^S$  are treated as non-member shadow graphs. To balance member and non-member shadow graphs, the attacker can randomly generate non-member graphs, with negligible overlap due to large variations in node and edge sizes. We denote the member and non-member shadow graphs in  $G^S$  as  $G_{mem}^S$  and  $G_{non-mem}^S$ , respectively. Second, for each graph in  $G_i^S$  in  $G_{mem}^S$  and  $G_{non-mem}^S$ , the attacker inputs it into  $\Phi^S$  to obtain a corresponding set of synthetic graphs  $\hat{G}^S$ . We denote this as a triplet  $(G_i^S, \hat{G}_i^S, y)$ , where  $y$  is  $G_i^S$ ’s membership label.

**Attack feature construction.** After obtaining the synthetic graphs  $\hat{G}_i^S$  for each graph  $G_i^S$  in  $G^S$ , we construct the attack features  $\mathcal{A}^{\text{train}}$  for attack training in two ways.

$\mathcal{A}^{train-1}$ : the design of this attack feature is based on the principle that the generated graphs  $\hat{G}_i^S$  should be much similar to its corresponding shadow graph  $G_i^S$  when  $G_i^S$  is a member shadow graph in  $G^S$ , compared to when it is a non-member. **First**, for each shadow graph  $G_i^S$ , the adversary measures the pairwise similarity between  $G_i^S$  and each generated graph  $\hat{g}_i^S$  in  $\hat{G}_i^S$ . Specifically, we first convert both  $\hat{g}_i^S$  and  $G_i^S$  into embedding vectors using Anonymous Walk Embeddings (AWEs) [31], denoted as  $emb(\hat{g}_i^S)$  and  $emb(G_i^S)$ , respectively. The AWEs method represents graphs by sampling random walks and anonymizing node identities based on their first appearance indices. The frequency distribution of these anonymous walk patterns is then embedded into a lower-dimensional space using neural networks like Word2Vec. This approach captures structural similarity between graphs while being scalable and identity-independent. **Second**, we calculate the similarity between the embedding vectors of  $emb(\hat{g}_i^S)$  and  $emb(G_i^S)$  as  $sim_k(emb(\hat{g}_i^S), emb(G_i^S))$ , where  $sim_k$  represents the  $k$ -th similarity function. In this paper,  $k \in \{1, 2, 3, 4\}$  corresponds to four similarity metrics: Dot product, Cosine similarity, Euclidean distance-based difference (as defined in Equation 5), and Jensen-Shannon Diversity (JSD). **Third**, we construct the attack feature of a triplet  $(G_i^S, \hat{G}_i^S, y)$  by stacking the pairwise embedding similarities between  $G_i^S$  and  $\hat{G}_i^S$ , which can be written as:

$$A_i^{train} = \begin{bmatrix} \|\forall k \in \{1,2,3,4\}, sim_k(emb(\hat{g}_{i,0}^S), emb(G_i^S))\| \\ \dots \\ \|\forall k \in \{1,2,3,4\}, sim_k(emb(\hat{g}_{i,N}^S), emb(G_i^S))\| \end{bmatrix}, \quad (11)$$

where  $\|\forall k \in \{1,2,3,4\}, sim_k(emb(\cdot), emb(\cdot))\|$  denotes the concatenation of similarity values computed using the four metrics between two embeddings.  $\hat{g}_i^S, n$  is the  $n$ -th generated graph in  $\hat{G}_i^S$ , with  $n \in 1, \dots, N$  and  $N$  being the total number of generated graphs in  $\hat{G}_i^S$ . Therefore, for each triplet  $(G_i^S, \hat{G}_i^S, y)$ , the dimension of its corresponding attack feature  $A^{train}$  is  $N \times 4$ .

$\mathcal{A}^{train-2}$ : the design of this attack feature is based on the principal that the generated graphs  $\hat{G}_i^S$  should exhibit much higher pairwise similarity within  $\hat{G}_i^S$  when  $G_i^S$  is a member shadow graph than when it is a non-member. **First**, for each shadow graph  $G_i^S$ , the adversary measures the pairwise similarity within  $\hat{G}_i^S$ . Similar to  $\mathcal{A}^{train-1}$ , we first convert each  $\hat{g}_i^S$  in  $\hat{G}_i^S$  into an embedding vector,  $emb(\hat{g}_i^S)$ , using AWEs. **Second**, we pairwise calculate the similarity between the embedding vectors of  $emb(\hat{g}_i^S)$  and  $emb(\hat{g}_j^S)$  as  $sim_k(emb(\hat{g}_i^S), emb(\hat{g}_j^S))$ ,  $i < j$ . The similarity metrics  $sim_k$  are the same as those used in  $\mathcal{A}^{train-1}$ . **Third**, we construct the attack feature of a triplet  $(G_i^S, \hat{G}_i^S, y)$  by stacking the pairwise embedding similarities within  $\hat{G}_i^S$ :

$$A_i^{train} = \begin{bmatrix} \|\forall k \in \{1,2,3,4\}, sim_k(emb(\hat{g}_{i,0}^S), emb(\hat{g}_{i,1}^S))\| \\ \dots \\ \|\forall k \in \{1,2,3,4\}, sim_k(emb(\hat{g}_{i,N-1}^S), emb(\hat{g}_{i,N}^S))\| \end{bmatrix}. \quad (12)$$

The meaning of the notations is consistent with those in Equation 11. Therefore, for each triplet  $(G_i^S, \hat{G}_i^S, y)$ , the dimension of its corresponding attack feature  $A^{train}$  is  $\frac{N*(N-1)}{2} \times 4$ .

After the adversary generates the feature  $A_i^{train}$  of the shadow graph  $G_i^S$ , it associates  $A_i^{train}$  with its ground-truth membership label  $y$ . Finally, the adversary adds the newly formed data sample  $(A_i^{train}, y)$  to  $\mathcal{A}^{train}$ . In our empirical study, we ensure  $\mathcal{A}^{train}$  is balanced, i.e., both the member and non-member classes have the same number of samples.

**Attack model training.** After  $\mathcal{A}^{train}$  is generated, the adversary proceeds to train the attack classifier, such as Multi-layer Perceptron (MLP), Random Forest (RF), and Linear Regression (LR), on  $\mathcal{A}^{train}$ .

**Membership inference.** At inference time, the adversary uses the same method used to generate the training feature  $A^{train}$  to derive the feature  $A^{att}$  for the target graph  $G$ . Specifically, the adversary inputs  $G$  to the target model  $\Phi$  and obtains a set of generated graphs. Then, the adversary calculates the similarity between the generated graphs and  $G$  or within the generated graphs using the same approaches and similarity functions. Finally, the adversary feeds  $A^{att}$  into the trained attack classifier to obtain predictions, just as in the training phase. The label associated with a higher probability will be selected as the inference output.

## 6 Evaluation

This section evaluates the effectiveness of our attacks.

### 6.1 Experimental Setup

All the algorithms are implemented in Python with PyTorch and executed on NVIDIA A100-PCIE-40GB.

**Datasets.** We use six real-world datasets from four domains: two molecule datasets (MUTAG, QM9), one protein dataset (ENZYMES), one citation dataset (Ego-small), and two social networks (IMDB-BINARY, IMDB-MULTI). These datasets each comprise a collection of graphs and serve as benchmarks for evaluating graph-based models across domains [45]. Statistical details are provided in Appendix B. Throughout the paper, we refer to IMDB-BINARY and IMDB-MULTI as IMDB-B and IMDB-M, respectively.

**Target models.** We employ three state-of-the-art GGDMs, namely EDP-GNN [48], GDSS [33], and Digress [67].

- **EDP-GNN** is a score-based generative diffusion model (SGM) that models data distributions via score functions and represents a pioneer effort in deep graph generation.
- **GDSS** is one of the stochastic differential equations (SDE)-based models that capture the joint distribution of nodes and edges through a system of SDE.
- **Digress** is a type of denoising diffusion probabilistic model (DDPM) that incrementally edits graphs by adding or removing edges and altering categories, subsequently reversing these changes with a graph transformer.

Dataset	Attack	EDP-GNN					GDSS					Digress				
		P	R	F1	R1	R2	P	R	F1	R1	R2	P	R	F1	R1	R2
MUTAG	Ours	<b>0.70</b>	<b>0.90</b>	<b>0.78</b>	<b>0.21</b>	<b>0.27</b>	<b>0.87</b>	<b>0.83</b>	<b>0.85</b>	<b>0.22</b>	<b>0.32</b>	<b>1</b>	<b>0.72</b>	<b>0.84</b>	<b>0.11</b>	<b>0.16</b>
	Baseline-1	0.50	0.65	0.57	0	0	0.54	0.67	0.59	0	0	0.10	0.15	0.12	0	0
	Baseline-2	0.45	0.51	0.48	0	0	0.45	0.52	0.48	0	0	<b>1</b>	0.33	0.50	0	0
ENZYMES	Ours	<b>0.85</b>	<b>1</b>	<b>0.92</b>	<b>0.11</b>	<b>0.36</b>	0.83	<b>0.92</b>	<b>0.85</b>	<b>0.09</b>	<b>0.28</b>	<b>1</b>	<b>0.79</b>	<b>0.88</b>	<b>0.03</b>	<b>0.04</b>
	Baseline-1	0.54	0.73	0.62	0	0	0.75	0.75	0.70	0	0	0.01	0.19	0.01	0	0
	Baseline-2	0.66	0.69	0.67	0	0	<b>0.92</b>	0.31	0.46	0	0	<b>1</b>	0.22	0.36	0	0
Ego-small	Ours	<b>0.84</b>	<b>0.88</b>	<b>0.86</b>	<b>0.11</b>	<b>0.24</b>	<b>0.76</b>	<b>0.78</b>	<b>0.77</b>	<b>0.07</b>	<b>0.19</b>	<b>1</b>	<b>0.75</b>	<b>0.85</b>	<b>0.10</b>	<b>0.22</b>
	Baseline-1	0.63	0.76	0.69	0	0.03	0.60	0.72	0.65	0	0	0.09	0.43	0.15	0	0
	Baseline-2	0.60	0.66	0.63	0	0	0.61	0.62	0.62	0	0	<b>1</b>	0.30	0.46	0	0
IMDB-B	Ours	0.85	<b>1</b>	<b>0.92</b>	<b>0.14</b>	<b>0.27</b>	<b>0.88</b>	<b>1</b>	<b>0.94</b>	<b>0.15</b>	<b>0.28</b>	0.85	<b>0.80</b>	<b>0.82</b>	<b>0.09</b>	<b>0.10</b>
	Baseline-1	0.80	0.79	0.80	0.05	0.12	0.78	0.89	0.83	0.02	0.03	0.001	0.20	0.01	0	0
	Baseline-2	<b>0.96</b>	0.31	0.47	0	0	0.41	0.47	0.44	0	0	<b>1</b>	0.21	0.34	0	0
IMDB-M	Ours	<b>0.92</b>	<b>0.96</b>	<b>0.94</b>	<b>0.13</b>	<b>0.19</b>	<b>0.99</b>	<b>1</b>	<b>0.99</b>	<b>0.15</b>	<b>0.20</b>	0.72	<b>0.89</b>	<b>0.79</b>	<b>0.07</b>	<b>0.12</b>
	Baseline-1	0.73	0.83	0.78	0	0.11	0.70	0.89	0.77	0.05	0.06	0.01	0.20	0.013	0	0
	Baseline-2	0.61	0.78	0.68	0	0	0.54	0.58	0.56	0	0	<b>1</b>	0.22	0.36	0	0
QM9	Ours	<b>0.79</b>	<b>0.84</b>	<b>0.81</b>	<b>0.05</b>	<b>0.08</b>	<b>0.80</b>	<b>0.85</b>	<b>0.82</b>	<b>0.06</b>	<b>0.11</b>	<b>0.88</b>	<b>0.73</b>	<b>0.79</b>	<b>0.05</b>	<b>0.07</b>
	Baseline-1	0.34	0.41	0.37	0	0	0.62	0.75	0.68	0	0	0.08	0.43	0.13	0	0
	Baseline-2	0.38	0.53	0.44	0	0	0.37	0.60	0.46	0	0	0.81	0.23	0.35	0	0

Table 2: Performance of graph reconstruction attack. "P", "R", "F1", "R1", and "R2" represent the metrics of precision, recall, F1 score, coverage ratio of the exact-matched graphs, and coverage ratio of graphs for which the attack achieves an F1 score above 0.75, respectively. For each dataset and evaluation metric under each target model, the better results are highlighted in **bold**.

**Implementation settings.** For the target model training, we use the default parameter settings specified in the respective papers. Both target and shadow models used early stopping: when loss fails to improve for 50 consecutive epochs. For each dataset, we randomly divided the graphs into training, validation, and testing sets with a ratio of 0.7/0.1/0.2.

## 6.2 Graph Reconstruction Performance

**Evaluation metrics.** Recall that the attacker’s goal is to accurately uncover the graph structures of training graphs used by the target generative models based on the generated graphs. After obtaining the reconstructed graphs  $G^{rec}$ , we use a graph alignment algorithm REGAL [25] to align each reconstructed graph  $g^{rec}$  in  $G^{rec}$  with each training graph  $g^{train}$  in  $\Phi^{train}$ . Next, we evaluate the F1 score based on all aligned pairs  $(g^{rec}, g^{train})$ . For each  $g^{rec}$ , we regard the  $g^{train}$  with highest F1 score as the graph that  $g^i$  is generated from, labeled as  $g^*$ . Then we calculate the attack performance as the average attack performance on all  $(g^i, g^*)$  pairs.

- **Structure metrics.** We employ three edge-related metrics, namely precision (P), recall (R), and F1 score (F1), to assess the overall effectiveness of our attack in accurately recovering the exact edges and non-edges in graphs.
- **Global Metrics.** To measure the overall graph reconstruction performance, we evaluate two coverage ratios: the proportion of training graphs exactly matched by the attack (R1), and the proportion with an F1 score above 0.75 (R2).

**Attack setup.** For each dataset and each target model, we directly run the target model  $\Phi$  and obtain 1,000 generated graphs, and then apply our GRA on these generated graphs.

**Competitor.** Previous GRAs typically rely on node or graph embeddings [55, 86], or on the gradients of the target model (GNNs) during the training process [85]. However, these approaches are not directly applicable to GGDMs, which are trained via forward and reverse processes and output a set of generated graphs. Therefore, we compare our attack with two baselines without the graph alignment step:

- **Baseline-1:** We treat the generated graphs  $g_i \in \hat{G}$  as replicas of  $\Phi^{train}$ , and evaluate attack performance as the average over all aligned pairs  $(g^i, g^*)$ , where  $g^*$  in  $\Phi^{train}$  yields the highest F1 score with  $g^i$ .
- **Baseline-2:** We reconstruct a graph  $g^{rec}$  by directly computing the overlapped edges between  $g_i$  and  $g_j$  using Equation 7 without graph alignment. Attack performance is evaluated over all aligned pairs  $(g^{rec}, g^*)$ , where  $g^*$  in  $\Phi^{train}$  yields the highest F1 score with  $g^{rec}$ .

**Experimental results.** Table 2 presents the GRA results. We have the following observations. First, our attack achieves outstanding performance against all three GGDMs, with F1 scores ranging from 0.72 to 0.99, and coverage ratios of exact-matched graphs and graphs with F1 score above 0.75 reaching up to 0.21 and 0.36, respectively. Second, our attack model outperforms the baselines without the graph alignment module on structural metrics. Although Baseline-2 achieves high precision in some cases, it suffers from significantly lower recall, resulting in lower overall F1 scores. For example, the precision performance of Baseline-2 against EDP-GNN model on IMDB-B dataset is 0.96, but the recall drops to 0.31, and thus yields the F1 score of 0.47. Third, our model demonstrates substantially better global performance than both baselines, as reflected in higher coverage ratios of exact-matched graphs and graphs where the attack achieves a strong F1 score.

Dataset	EDP-GNN		GDSS		Digress		Orig.	
	$D(\text{degree})\downarrow$	$D(\text{density})\downarrow$	$D(\text{degree})\downarrow$	$D(\text{density})\downarrow$	$D(\text{degree})\downarrow$	$D(\text{density})\downarrow$	Degree	Density
MUTAG	0.009	0.001	0.123	0.026	0.139	0.013	1.093	0.070
ENZYMES	0.059	0.003	0.117	0.014	0.258	0.010	1.939	0.075
Ego-small	0.009	0.008	0.120	0.024	0.162	0.013	2.000	0.487
IMDB-B	0.109	0.008	0.178	0.010	0.190	0.029	4.850	0.219
IMDB-M	0.005	0.007	0.228	0.003	0.249	0.022	5.266	0.309
QM9	0.045	0.036	0.273	0.049	0.312	0.031	4.136	0.260

Table 3: Performance of PIA - the absolute differences in property values between the target model’s training set and the inferred values from generated graphs. A smaller difference represents a better performance. " $D(\text{degree})$ " and " $D(\text{density})$ " represent the absolute differences of average degree and density, respectively. "Orig." means the property of target model’s training set  $\Phi^{\text{train}}$ .




Figure 5: Performance of PIA - the distribution of training graphs in different degree and density ranges. Figure (a) - (c) shows results for graph density, and Figure (d) - (f) for graph degree ( $k = 5$ , IMDB-B dataset).

### 6.3 Property Inference Performance

**Evaluation metrics.** To assess the attack’s effectiveness, we directly compare the property values of the target model’s training set with those inferred from the attack. The experimental results report the absolute differences between the property values of the actual training set and the inferred values from the generated graphs.

**Attack setup.** For each dataset and each target model, we directly run the target model  $\Phi$  and obtain 1,000 generated graphs. We calculate and summarize the property values on these generated graphs as the property of the training set of  $\Phi$ . In our experiments, we consider four different types of graph properties: graph density, average node degree, average number of triangles per node, and graph arboricity (which represents the minimum number of spanning forests into which the edges of a graph can be partitioned). For each graph property, we consider two types of analysis: the average graph property across the training graphs in  $\Phi^{\text{train}}$ , and the distribution of  $\Phi^{\text{train}}$  across different ranges of property values. For the second type, we uniformly bucketize the property values into  $k$  distinct ranges and infer the proportion of graphs that fall into each range. We set  $k = \{5, 10\}$  in the experiments.

**Experimental results.** Table 3 shows the absolute differences of average graph degree and density between the target model’s training set and the inferred values from generated graphs. We observe that the absolute difference in average node degree is no more than 0.312, corresponding to a difference ratio of below 7.5% from the original value of  $\Phi^{\text{train}}$ . Similarly, the absolute difference in average node density is less than 0.049, with a difference ratio less than 18.8%. These

results demonstrate that our simple yet effective attack can accurately infer the graph properties of  $\Phi^{\text{train}}$ .

Figure 5 shows the distribution in different node degree and density ranges with a bucket number of  $k = 5$ . We observe that the inferred distributions closely align with those of  $\Phi^{\text{train}}$  across all settings, with absolute ratio differences for each range class varying from 0.001 to 0.028 for average graph density, and 0 to 0.014 for average graph degree.

We show the results of the properties of average number of triangles per node and graph arboricity, along with the distribution results of  $k = 10$  in Appendix C. The observations are similar to those in Table 3 and Figure 5.

### 6.4 Membership Inference Performance

**Evaluation metrics.** For assessing MIA effectiveness, we employ three metrics: (1) *Attack accuracy*: the ratio of correctly predicted member/non-member graphs among all target graphs; (2) *Area Under the Curve (AUC)*: measured over the true positive rate (TPR) and false positive rate (FPR) at various thresholds of the attack classifier; and (3) *True-Positive Rate at False-Positive Rates (TPR@FPR)* [11]: the TPR at various FPR values. We use TPR@0.1FPR in our experiments.

**Attack setup.** For each dataset and target GDM, we first split the original data into two halves: with 50% used as the training set for the target generative diffusion model, representing the member graphs ( $G_{\text{mem}}$ ), and the remaining 50% serving as non-member graphs ( $G_{\text{non-mem}}$ ). Next, we generate the MIA training dataset by randomly sampling 50% of the graphs from  $G_{\text{mem}}$  as members, while selecting an equal num-

Dataset	Attack	EDP-GNN			GDSS			Digress		
		Accuracy	AUC	TPR@FPR	Accuracy	AUC	TPR@FPR	Accuracy	AUC	TPR@FPR
MUTAG	Ours-1	0.650	0.763	0.225	0.821	0.811	0.467	0.845	0.881	0.327
	Ours-2	<u>0.698</u>	0.750	<u>0.234</u>	<b>0.887</b>	<b>0.936</b>	<b>0.551</b>	<b>0.850</b>	<b>0.951</b>	<b>0.472</b>
	Baseline-1	<b>0.750</b>	<u>0.813</u>	<b>0.290</b>	<u>0.833</u>	0.889	0.519	<u>0.849</u>	0.914	<u>0.442</u>
	Baseline-2	<b>0.750</b>	<b>0.830</b>	<b>0.290</b>	<u>0.833</u>	<u>0.917</u>	<u>0.548</u>	0.833	<u>0.944</u>	0.420
ENZYMES	Ours-1	0.813	0.831	0.427	0.825	0.868	0.467	0.817	0.845	<b>0.344</b>
	Ours-2	<b>0.900</b>	<b>0.913</b>	<b>0.539</b>	<u>0.852</u>	<u>0.880</u>	<u>0.562</u>	<u>0.835</u>	<b>0.869</b>	<u>0.332</u>
	Baseline-1	0.760	0.741	0.276	0.696	0.752	0.292	0.733	0.750	0.288
	Baseline-2	<u>0.833</u>	<u>0.857</u>	<u>0.437</u>	<b>0.855</b>	<b>0.892</b>	<b>0.585</b>	<b>0.846</b>	<u>0.868</u>	0.330
Ego-small	Ours-1	0.874	0.969	<u>0.902</u>	<u>0.917</u>	<u>0.972</u>	0.768	<u>0.854</u>	<u>0.901</u>	<u>0.453</u>
	Ours-2	<b>0.901</b>	<b>0.991</b>	<b>0.913</b>	<b>0.934</b>	<b>0.994</b>	<b>0.821</b>	<b>0.892</b>	<b>0.943</b>	<b>0.487</b>
	Baseline-1	0.703	0.812	0.412	0.715	0.833	0.517	0.671	0.763	0.272
	Baseline-2	<u>0.882</u>	<u>0.974</u>	0.901	0.909	0.968	<u>0.781</u>	0.825	0.872	0.429
IMDB-B	Ours-1	<u>0.955</u>	0.979	<u>0.757</u>	0.908	<b>0.999</b>	<u>0.957</u>	0.918	0.957	0.652
	Ours-2	<b>0.991</b>	<b>0.995</b>	<b>0.878</b>	<b>0.992</b>	<b>0.999</b>	<b>0.970</b>	0.933	<b>0.989</b>	<b>0.727</b>
	Baseline-1	0.731	0.784	0.313	0.667	<u>0.722</u>	0.288	0.640	0.703	0.292
	Baseline-2	0.917	<u>0.986</u>	0.733	<u>0.986</u>	<b>0.999</b>	<b>0.970</b>	<b>0.946</b>	<u>0.969</u>	<u>0.724</u>
IMDB-M	Ours-1	0.902	<b>0.992</b>	<b>0.837</b>	0.942	<u>0.911</u>	0.667	0.850	0.896	0.538
	Ours-2	<u>0.912</u>	<u>0.971</u>	<u>0.789</u>	<b>0.999</b>	<b>0.999</b>	<b>0.985</b>	<b>0.925</b>	<u>0.955</u>	<u>0.633</u>
	Baseline-1	0.739	0.781	0.308	0.724	0.762	0.302	0.645	0.711	0.298
	Baseline-2	<b>0.938</b>	0.969	0.774	<u>0.973</u>	<b>0.999</b>	0.970	<u>0.892</u>	<b>0.961</b>	<b>0.652</b>
QM9	Ours-1	<u>0.752</u>	0.816	<u>0.412</u>	0.766	0.843	<u>0.437</u>	0.701	0.780	<u>0.335</u>
	Ours-2	<b>0.781</b>	<b>0.854</b>	<b>0.456</b>	<b>0.826</b>	<u>0.849</u>	<b>0.472</b>	<b>0.772</b>	<b>0.829</b>	<b>0.407</b>
	Baseline-1	0.638	0.733	0.225	0.682	0.767	0.305	0.599	0.671	0.209
	Baseline-2	0.750	<u>0.821</u>	0.403	<u>0.793</u>	<b>0.850</b>	0.426	<u>0.726</u>	<u>0.788</u>	0.333

Table 4: MIA performance under setting 1 (Non-transfer). "Ours-1" and "Ours-2" denote our attacks with  $\mathcal{A}^{\text{train-1}}$  and  $\mathcal{A}^{\text{train-2}}$ , respectively. For each dataset and each target model, the best result is highlighted in **bold**, and the second-best is underlined.

ber of graphs from  $G_{\text{non-mem}}$  as non-members. The remaining graphs from both  $G_{\text{mem}}$  and  $G_{\text{non-mem}}$  are used to form the MIA testing dataset. This ensures the testing set contains an equal number of member and non-member graphs, with no overlap with the training set. We feed each graph in the MIA training and testing set is then fed into the trained shadow model  $\Phi^S$  and generate 100 graphs per input graph. These generated graphs are subsequently used to construct the attack features. We consider two distinct settings for shadow graph:

- **Setting 1 (Non-transfer setting):** Both the shadow and target graphs are sampled from the same dataset.
- **Setting 2 (Dataset transfer setting):** The shadow and target graphs are sampled from different datasets. Specifically, we replace the MIA testing set in the non-transfer setting with the testing set from a dataset different from that of the shadow model.

**Competitors.** We compare our method with two state-of-the-art white-box baseline attacks, both assuming the attacker has access to the target model. These baselines are regarded as strong attacks and are commonly used in prior MIA studies: one is loss-based, and the other is gradient-based.

- **Baseline-1: Loss-based MIA.** We adapt the white-box loss-based MIA from image-based diffusion models [12, 15, 43] to our setting. First, we compute the loss values of member and non-member samples across different timesteps. Then, we train an attack classifier using the aggregated loss values

from the most effective timestep range.

- **Baseline-2: Gradient-based MIA.** We adapt the white-box gradient-based MIA from image-based diffusion models [50] to our setting. Specifically, we aggregate gradients from member and non-member samples over the most effective timestep range and feed them into the attack classifier.

**Attacker performance under Setting 1 (Non-transfer setting).** Table 4 illustrates the attack performance under Setting 1. We have the following observations. First, both our MIA models using attack feature  $\mathcal{A}^{\text{train-1}}$  and  $\mathcal{A}^{\text{train-2}}$  demonstrate outstanding performance across four datasets and three diffusion models, yielding attack accuracy ranging from 0.650 to 0.955 for Ours-1 and 0.698 to 0.999 for Ours-2, attack AUC ranging from 0.763 to 0.992 for Ours-1 and 0.750 to 0.999 for Ours-2, TPR@FPR ranging from 0.225 to 0.957 for Ours-1 and 0.234 to 0.985 for Ours-2. Second, Ours-2 outperforms Baselines-1 in most settings. Particularly, Ours-2's attack accuracy and AUC are respectively 0.325 and 0.277 higher than those of Baseline-1 when GDSS is the target model and IMDB-B is the target dataset. Third, Baseline-2 achieves performance comparable to ours in some settings, due to its use of the target model's gradients. However, it is a white-box attack requiring access to the model's structure and parameters, which may be difficult to obtain in real-world scenarios. Moreover, its substantially larger feature dimension causes Baseline-2 to overfit in some settings, leading

Shadow dataset	MUTAG	Ego-small	ENZYMES	IMDB-B	IMDB-M	QM9
MUTAG	0.764	0.820	0.743	0.802	0.676	0.724
Ego-small	0.671	0.969	0.727	0.921	0.924	0.697
ENZYMES	0.626	0.857	0.831	0.903	0.956	0.742
IMDB-B	0.640	0.914	0.756	0.979	0.960	0.723
IMDB-M	0.639	0.890	0.756	0.953	0.992	0.721
QM9	0.612	0.825	0.745	0.901	0.869	0.816

(a) Ours-1

Shadow dataset	MUTAG	Ego-small	ENZYMES	IMDB-B	IMDB-M	QM9
MUTAG	0.750	0.901	0.895	0.883	0.754	0.776
Ego-small	0.674	0.991	0.874	0.868	0.837	0.780
ENZYMES	0.662	0.846	0.913	0.750	0.734	0.808
IMDB-B	0.684	0.889	0.869	0.995	0.866	0.790
IMDB-M	0.702	0.850	0.860	0.857	0.971	0.768
QM9	0.712	0.832	0.883	0.811	0.803	0.854

(b) Ours-2

Shadow dataset	MUTAG	Ego-small	ENZYMES	IMDB-B	IMDB-M	QM9
MUTAG	0.813	0.744	0.671	0.705	0.687	0.624
Ego-small	0.580	0.812	0.639	0.659	0.648	0.601
ENZYMES	0.587	0.647	0.741	0.684	0.683	0.656
IMDB-B	0.546	0.678	0.647	0.784	0.754	0.679
IMDB-M	0.575	0.673	0.709	0.756	0.781	0.646
QM9	0.557	0.679	0.623	0.657	0.646	0.733

(c) Baseline-1

Shadow dataset	MUTAG	Ego-small	ENZYMES	IMDB-B	IMDB-M	QM9
MUTAG	0.830	0.777	0.833	0.813	0.748	0.707
Ego-small	0.589	0.974	0.713	0.784	0.754	0.655
ENZYMES	0.580	0.679	0.857	0.765	0.789	0.699
IMDB-B	0.511	0.767	0.667	0.986	0.828	0.635
IMDB-M	0.617	0.746	0.833	0.839	0.969	0.679
QM9	0.543	0.644	0.660	0.687	0.678	0.821

(d) Baseline-2

Figure 6: MIA AUC performance against EDP-GNN target model under setting 2 (dataset transfer setting).

to worse performance than ours in those cases. Additionally, Ours-2 outperforms Ours-1 in most settings, which indicates the attacker feature of  $\mathcal{A}^{\text{train-2}}$  is more effective than  $\mathcal{A}^{\text{train-1}}$ .

**Attacker performance under Setting 2 (Data transfer setting).** Figure 6 shows the attack results under Setting 2. The following observations can be made. First, for each target dataset, both of our attacks achieve the highest attack accuracy when shadow and target graphs originate from the same dataset (i.e., within the same column). Second, even under the transfer setting, the attack remains effective, with attack AUC ranging from 0.61 to 0.96 for Ours-1 and 0.66 to 0.90 for Ours-2. This demonstrates the attack models’ ability to learn knowledge from shadow data and transfer that knowledge to the target graphs. Specifically, for Ours-1, member graphs yield higher similarity scores between generated and input graphs than non-members; and for Ours-2, generated graphs from member graphs exhibit higher similarity than those from non-member graphs. Third, our attack models outperform both baselines in most cases, which indicates the effectiveness of our attack features. This also highlights an advantage of our model over Baseline-2, as our attacks are more applicable to real-world data transfer scenarios.

## 7 Two Possible Defenses

To enhance the privacy of GGDMs, one intuitive approach is to introduce noise into the models to confuse adversaries and provide privacy protection. Existing defense mechanisms for graph learning models mainly target GNNs [22, 47, 66] or random-walk-based models [20, 52]. These strategies typically inject noise into the gradients/parameters during model training, such as equipping the model with differential privacy [6], or adding noise to model’s output, namely the posteriors or final graph embeddings [23, 55, 69]. However, when we apply these defenses to GGDMs, the empirical evaluation (will be present in Sec. 7.3) shows that these methods significantly degrade the quality of the generated graphs. To address this issue, we propose both pre-processing and post-processing defenses that perturb the structures of training graphs before model training or after graph generation, respectively. Specifically, we selectively flip the edges and non-

edges that are deemed least important. This approach minimizes the quality loss of the generated graphs while providing sufficient privacy protection against inference attacks. We next detail the process of estimating the importance of edges and non-edges (Sec. 7.1), describe our defense mechanisms (Sec. 7.2), and present the empirical performance (Sec. 7.3).

### 7.1 Saliency Map-based Importance Measurement

To measure the importance of edges and non-edges, we utilize the saliency maps-based method [7, 8, 46, 57]. Saliency maps are widely used in explainable machine learning and artificial intelligence, particularly in the image domain, to provide insights into the relevance of input features that contribute to a model’s output [27, 46, 61]. The most popular methods generate a gradient saliency map by back-propagating gradients from the end of the neural network and projecting it onto an image plane [46, 51, 57]. These gradients are typically derived from a loss function, layer activation, or class activation. The absolute values of these gradients are often used to determine the importance of each feature, with a larger magnitude indicating a greater influence on the model’s prediction.

In our approach, we adapt this gradient saliency map-based technique to measure the importance of edges and non-edges for node classification within each graph. This process involves three steps: forward pass through the GNNs, gradient calculation on edges and non-edges, and saliency map construction. The following sections detail each step.

**Step 1: Forward pass through the GNNs.** Give a graph  $G(V, A, X, Y)$  where  $V$ ,  $A$ ,  $X$ , and  $Y$  denote the nodes, adjacency matrix, node features, and node labels, respectively. We feed  $G$  to a GNN for node classification. GNNs typically follow a message-passing paradigm to learn node embeddings. In this paper, we use Graph Convolutional Network (GCN) [34], whose layer-wise propagation is:

$$H^{(l)} = \sigma(\hat{A}H^{(l-1)}W^{(l)}), \quad (13)$$

where  $\hat{A} = \tilde{D}^{-1/2}(\tilde{A})\tilde{D}^{-1/2}$  denotes the normalized adjacency matrix, with  $\tilde{A} = A + I$  and  $I$  is the identity matrix.  $\tilde{D}$  is the degree matrix of  $\tilde{A}$ ,  $W^{(l)}$  is the learnable weight matrix at  $l$ -th

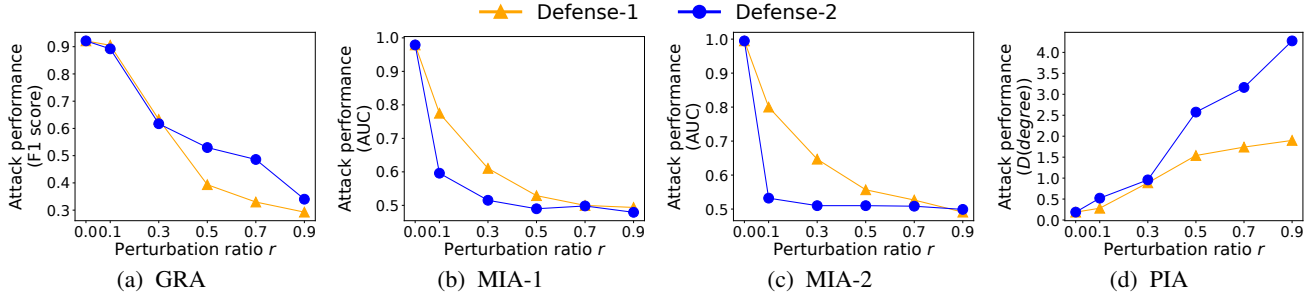


Figure 7: Impact of perturbation ratio  $r$  on defense performance (EDP-GNN model, IMDB-B dataset). "MIA-1" and "MIA-2" represent MIA with attack feature of  $\mathcal{A}^{\text{train-1}}$  and  $\mathcal{A}^{\text{train-2}}$ , respectively.  $D(\text{degree})$  denotes the difference in degree between inferred and original training graphs. For GRA, MIA-1, and MIA-2, a lower attack F1 score or AUC indicates stronger defense, while for PIA, a greater difference between the inferred and original property values represents stronger defense.

layer, and  $\sigma$  denotes a non-linear activation function such as ReLU. The final node embeddings  $H^{(l)}$  are passed through a softmax layer to predict node labels. The model is trained using cross-entropy loss. Note that in real-world applications, some graphs may lack node features  $X$  or node labels  $Y$  or both. We address this by employing common strategies: (1) if  $X$  is missing, we set  $X$  as the identity matrix, and (2) if  $Y$  is missing, we set  $Y$  as the one-hot encoded node degrees.

### Step 2: Gradient calculation on edges and non-edges.

Once the GNN's performance stabilizes after several training epochs, we backpropagate the gradient of the node classification loss to the adjacency matrix. The gradient of the loss  $L$  for each entry  $A_{ij}$  is computed via the chain rule as:

$$\frac{\partial L}{\partial A_{ij}} = \sum_k \left( \frac{\partial L}{\partial H_k} \cdot \frac{\partial H_k}{\partial \hat{A}_k} \cdot \frac{\partial \hat{A}_k}{\partial A_{ij}} \right), \quad (14)$$

where  $A_{ij}$  is an entry in adjacency matrix  $A$ , which represents the link status between node  $v_i$  and node  $v_j$ ,  $\frac{\partial L}{\partial H_k}$  denotes the gradient of loss with respect to embedding  $H_k$  of node  $v_k$ ,  $\frac{\partial H_k}{\partial \hat{A}_k}$  represents the gradient of node embedding  $H_k$  with respect to the normalized adjacency matrix,  $\frac{\partial \hat{A}_k}{\partial A_{ij}}$  denotes the gradient of the normalized adjacency matrix with respect to  $A_{ij}$ .

**Step 3: Saliency map construction.** After obtaining the gradients of all entries in  $A$ , namely the gradients of all edges and non-edges within  $G$ , we construct the saliency map by taking the absolute value of the gradient at each  $A_{ij}$  as:

$$\text{Saliency}(A_{ij}) = \left| \frac{\partial L}{\partial A_{ij}} \right|. \quad (15)$$

A larger gradient magnitude indicates greater importance of the corresponding edge or non-edge to model's performance.

## 7.2 Details of the Defenses

While adding noise to training or generated graphs can reduce attacks' capabilities, random noise often causes substantial quality loss. Therefore, the primary objective of our defense is to provide effective protection against inference attacks while

preserving the model utility. To achieve this objective, we propose a perturbation-based strategy that selectively flips only the least important edges or non-edges or both in the training graphs (pre-processing) or generated graphs (post-processing). Although these targeted modifications may slightly alter the utility of generated graphs, our experiments show that they effectively reduce attack performance with minimal utility loss. Both defenses follow a three-step process. The pseudo-code of the defenses is provided in Appendix D.

### Defense-1: Noisy training graphs (pre-processing).

- Step 1: For each graph  $G$  in target model's training set, we construct a saliency map to measure the importance of each edge and non-edge in  $G$  by using the method in Section 7.1.
- Step 2: We rank all edges and non-edges in a single list based on their importance values in ascending order.
- Step 3: We flip one entry in the adjacency matrix  $A$  with the lowest importance. The perturbed graph is labeled as  $G'$ . We then repeat Steps 1 to 3 on the newly generated graph  $G'$  until the number of flipped entries reaches the pre-defined budget,  $\lfloor r \times |E| \rfloor$ , where  $|E|$  denotes the number of edges in  $G$ , and  $r \in (0, 1]$  is the *perturbation ratio*. A larger  $r$  introduces more structural changes. Finally, we train the target GGDM on the perturbed training graphs.

### Defense-2: Noisy generated graphs (post-processing).

- Step 1: Similar to Defense-1, we first measure the importance of each edge and non-edge for every graph  $\hat{G}$  in the generated graph set using the saliency map-based method.
- Step 2 and Step 3: These steps remain the same as those in Defense-1, involving ranking and flipping the least important edge or non-edge to obtain the perturbed graph  $\hat{G}'$ . The process is repeated until each graph in  $\hat{G}$  reaches its pre-defined flip budget, after which the perturbed graphs form the final output of the target GGDM.

## 7.3 Performance of Defense

In this section, we provide an assessment of the performance of our defense mechanism. We use the same target models and datasets as those for the attack evaluation (Section 6).

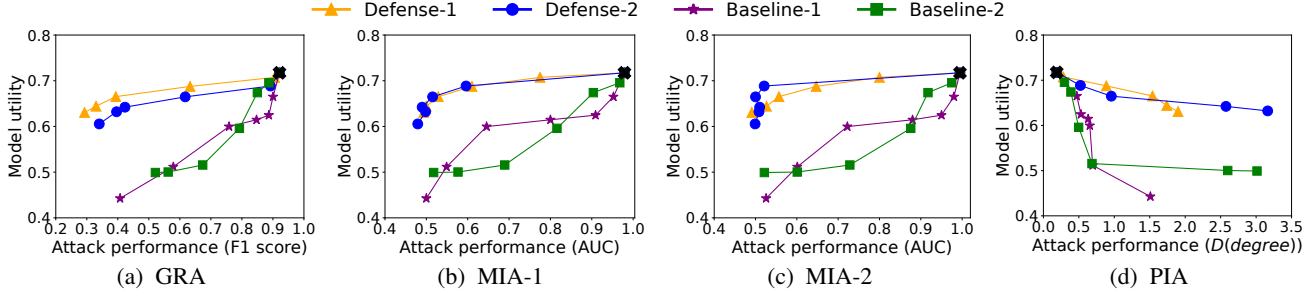


Figure 8: The defense-utility curve (EDP-GNN model, IMDB-B dataset). The  $x$  symbol indicates no defense. For GRA, MIA-1, and MIA-2, a lower attack F1 score or AUC combined with higher model utility indicates a better trade-off. For PIA, a greater difference between the inferred and original property values, along with higher model utility, signifies a better trade-off.

**Setup.** We set the perturbation ratio  $r = \{0, 0.1, 0.3, 0.5, 0.7, 0.9\}$ . Larger values of  $r$  indicate stronger privacy protection.

**Metrics.** We measure *defense effectiveness* as the attack performance after applying the defense, using the same evaluation metrics as in Section 6, with variations specific to attack types. The *utility* of the target model is evaluated as the performance of its generated graphs on a downstream task of graph classification. A higher AUC indicates better model utility.

**Baselines.** For comparative analysis, we compare our defense mechanism with the following two baselines:

- **Baseline-1: Differential privacy (DP).** We adopt DP-SGD, a state-of-the-art DP-based deep learning method [6], as our first baseline. We add Gaussian noise to gradients in each training iteration of target model. We set the privacy budget as  $\epsilon = \{0.001, 0.01, 0.1, 1, 5, 10\}$ , where a smaller  $\epsilon$  indicates stronger privacy guarantees.
- **Baseline-2: Randomly-flipped generated graphs.** This baseline applies random flips where each edge and non-edge in the generated graphs is flipped with a probability  $p$ . We set  $p = \{0.1, 0.2, 0.3, 0.4, 0.5, 0.6\}$  in experiments.

**Defense effectiveness.** Figure 7 presents the attack performance after deploying our defense mechanisms with various perturbation ratios when EDP-GNN is used as the target model. We have the following observations. First, our perturbation methods are highly effective against all three types of attacks, notably reducing their performance. Second, the defense strength increases with a higher perturbation ratio  $r$ . Even with a modest perturbation ratio of 0.1 (perturbing only 10% of graph edges), the attack AUC drops to around 0.6 for MIA-1 and 0.53 for MIA-2 with Defense-2.

**Defense-Utility tradeoff.** To illustrate the defense-utility trade-off, we plot the defense-utility curve of our defenses and the two baselines in Figure 8. We establish a set of configurations by varying the perturbation ratio for our defenses, the privacy budget for Baseline-1, and the flipped probability for Baseline-2, as described in the Setup. For each configuration, we evaluate defense effectiveness and target model utility, with each point on the curve representing a defense-utility

pair. We observe that our defenses outperform the baselines in this trade-off. When defense performance drops below 0.6 for GRA, MIA-1, and MIA-2, which indicates ineffective attacks, our defenses maintain higher model utility compared to baselines, with degradation less than 0.05. Similarly, at the same utility level, our methods provide stronger protection across all four attacks. The superiority stems from selectively adding noise only to the least important edges and non-edges.

## 8 Conclusion

In this paper, we investigated the privacy leakage of graph generative diffusion models (GGDMs) and proposed three black-box attacks that extract sensitive information from a target model using only its generated graphs. Specifically, we introduced (i) a graph reconstruction attack that successfully infers training graph structures of the target GGDM, (ii) a property inference attack that extracts statistical properties from the generated graphs, which closely resemble those of the training graphs, and (iii) membership inference attacks that determine whether a given graph is included in the training set of the target GGDM. Extensive experiments demonstrate the effectiveness of these attacks against three state-of-the-art GGDMs. Furthermore, we have proposed two defense mechanisms based on perturbing training or generated graphs by flipping the least important edges and non-edges, which have been experimentally shown to be effective.

In future work, we will explore the vulnerability of GGDMs to other types of attacks, such as model inversion [18, 73] and substructure inference attacks [72, 86], including strongly connected components, stars, and specialized substructures.

## 9 Acknowledgments

We thank the anonymous reviewers for their feedback. This work was supported by the Hong Kong RGC Grants RGC grants C2003-23Y and C1043-24GF.

## 10 Ethical Considerations

**Stakeholder impact analysis.** The goal of our work is to help researchers and developers identify and reduce GGDM vulnerabilities, improving privacy posture and trustworthiness in real deployments. These vulnerabilities can affect multiple stakeholders, including GGDM deployers, data owners, and clients/users of GGDMs through the platform/API. Model deployers (e.g., cloud providers, AI marketplaces, or enterprises hosting GGDMs) may risk contract breaches, take-downs or patches, reputational harm, and disclosure obligations. Data owners may face re-identification risks and IP leakage; even without exact copies, adversaries can reverse-engineer datasets by stitching reconstructed fragments and statistics. Users may propagate sensitive priors into downstream products, and competitors can de-bias outputs to extract non-public signals. However, we have proposed two efficient defenses (Section 7) to mitigate these threats with moderate utility loss, offering practical tools for responsible GGDM deployment.

**Mitigation strategies.** To minimize potential negative consequences of this research, we release defensive methods alongside the attacks so that practitioners can immediately apply appropriate mitigations. We also gate high-risk components, such as detailed attack code or sensitive hyperparameters, and provide clear guidelines for their safe and responsible use. In addition, we offer deployment recommendations, including API-level safeguards (e.g., rate limiting, authentication, anomaly detection, and detailed auditing to prevent iterative probing attacks), as well as the structure-aware output sanitization introduced in our defense section, to help strengthen real-world GGDM deployments against misuse.

**Safety measures in our research process.** During research, we used public datasets and public models, no human subjects/PII, and isomorphism-invariant aggregates, avoiding reconstruction or release of any identifiable data. Our results are reported as means over multiple seeds/samples to minimize pinpointing any singular graph or any identifiable data.

**Long-term commitment.** We will maintain our public repository, respond to community feedback, and update threat assessments and defense recommendations as GGDM technologies evolve. This ensures that our work continues to support safe and responsible adoption of GGDMs.

## 11 Open Science

In line with open science policies, we have made the code publicly accessible to facilitate reproducibility and foster research. As mentioned in Section 9, we release our attacks together with defenses, while withholding high-risk components (the AWE module in MIA and the graph match module in GRA) to reduce misuse. Our code is available at <https://doi.org/10.5281/zenodo.17946102>. All datasets used in this paper are publicly available at PyTorch\_Geometric

(TUDataset).

## References

- [1] Amazon aws. <https://aws.amazon.com/marketplace/solutions/machinelearning>.
- [2] Azure. <https://azure.microsoft.com/en-us/products/ai-services>.
- [3] Google ai hub. <https://aihub.cloud.google.com/>.
- [4] Ibm. <https://www.ibm.com/products/watsonx-ai>.
- [5] Modzy: Ai model marketplace. <https://www.modzy.com/marketplace/>.
- [6] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, 2016.
- [7] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. Sanity checks for saliency maps. *Advances in neural information processing systems*, 31, 2018.
- [8] Ahmed Alqaraawi, Martin Schuessler, Philipp Weiß, Enrico Costanza, and Nadia Berthouze. Evaluating saliency map explanations for convolutional neural networks: a user study. In *Proceedings of the international conference on intelligent user interfaces*, pages 275–285, 2020.
- [9] Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.
- [10] Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z Li. A survey on generative diffusion models. *IEEE Transactions on Knowledge and Data Engineering*, 2024.
- [11] Nicholas Carlini, Steve Chien, Milad Nasr, Shuang Song, Andreas Terzis, and Florian Tramèr. Membership inference attacks from first principles. In *Proceedings of the Conference on Symposium on Security and Privacy*, pages 1897–1914, 2022.
- [12] Nicolas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models. In *USENIX Security Symposium*, pages 5253–5270, 2023.

- [13] Xiaohui Chen, Yukun Li, Aonan Zhang, and Li-ping Liu. Nvdiff: Graph generation through the diffusion of node vectors. *arXiv preprint arXiv:2211.10794*, 2022.
- [14] Jinhao Duan, Fei Kong, Shiqi Wang, Xiaoshuang Shi, and Kaidi Xu. Are diffusion models vulnerable to membership inference attacks? In *International Conference on Machine Learning*, pages 8717–8730. PMLR, 2023.
- [15] Jan Dubiński, Antoni Kowalczyk, Stanislaw Pawlak, Przemyslaw Rokita, Tomasz Trzcinski, and Pawel Morawiecki. Towards more realistic membership inference attacks on large diffusion models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4860–4869, 2024.
- [16] Vasisht Duddu, Antoine Boutet, and Virat Shejwalkar. Quantifying privacy leakage in graph embedding. In *Proceedings of EAI International Conference on Mobile and Ubiquitous Systems*, pages 76–85, 2020.
- [17] Wenqi Fan, Chengyi Liu, Yunqing Liu, Jiatong Li, Hang Li, Hui Liu, Jiliang Tang, and Qing Li. Generative diffusion models on graphs: Methods and applications. *arXiv preprint arXiv:2302.02591*, 2023.
- [18] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- [19] Neil Zhenqiang Gong and Bin Liu. Attribute inference attacks in online social networks. *ACM Transactions on Privacy and Security*, 21(1):1–30, 2018.
- [20] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proceedings of the ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 855–864, 2016.
- [21] Kilian Konstantin Haefeli, Karolis Martinkus, Nathanaël Perraudin, and Roger Wattenhofer. Diffusion models for graphs benefit from discrete state spaces. *arXiv preprint arXiv:2210.01549*, 2022.
- [22] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- [23] Xinlei He, Jinyuan Jia, Michael Backes, Neil Zhenqiang Gong, and Yang Zhang. Stealing links from graph neural networks. In *USENIX Security Symposium*, pages 2669–2686, 2021.
- [24] Xinlei He, Rui Wen, Yixin Wu, Michael Backes, Yun Shen, and Yang Zhang. Node-level membership inference attacks against graph neural networks. *arXiv preprint arXiv:2102.05429*, 2021.
- [25] Mark Heimann, Haoming Shen, Tara Safavi, and Danai Koutra. Regal: Representation learning-based graph alignment. In *Proceedings of the ACM international conference on information and knowledge management*, pages 117–126, 2018.
- [26] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [27] Chia-Yu Hsu and Wenwen Li. Explainable geoai: can saliency maps help interpret artificial intelligence’s learning process? an empirical study on natural feature detection. *International Journal of Geographical Information Science*, 37(5):963–987, 2023.
- [28] Hailong Hu and Jun Pang. Membership inference of diffusion models. *arXiv preprint arXiv:2301.09956*, 2023.
- [29] Hongsheng Hu, Zoran Salcic, Lichao Sun, Gillian Dobbie, Philip S Yu, and Xuyun Zhang. Membership inference attacks on machine learning: A survey. *ACM Computing Surveys*, 54(11s):1–37, 2022.
- [30] Han Huang, Leilei Sun, Bowen Du, Yanjie Fu, and Weifeng Lv. Graphgdp: Generative diffusion processes for permutation invariant graph generation. In *IEEE International Conference on Data Mining*, pages 201–210, 2022.
- [31] Sergey Ivanov and Evgeny Burnaev. Anonymous walk embeddings. In *International conference on machine learning*, pages 2186–2195, 2018.
- [32] Yangqin Jiang, Yuhao Yang, Lianghao Xia, and Chao Huang. Diffkg: Knowledge graph diffusion model for recommendation. In *Proceedings of the ACM international conference on web search and data mining*, pages 313–321, 2024.
- [33] Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of stochastic differential equations. In *International Conference on Machine Learning*, pages 10362–10383, 2022.
- [34] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [35] Fei Kong, Jinhao Duan, RuiPeng Ma, Hengtao Shen, Xiaofeng Zhu, Xiaoshuang Shi, and Kaidi Xu. An efficient membership inference attack for the diffusion model by proximal initialization. *arXiv preprint arXiv:2305.18355*, 2023.

- [36] Jingwei Li, Jing Dong, Tianxing He, and Jingzhao Zhang. Towards black-box membership inference attack for diffusion models. *arXiv preprint arXiv:2405.20771*, 2024.
- [37] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- [38] Zongwei Li, Lianghao Xia, and Chao Huang. Recdiff: diffusion model for social recommendation. In *Proceedings of the ACM International Conference on Information and Knowledge Management*, pages 1346–1355, 2024.
- [39] Chengyi Liu, Wenqi Fan, Yunqing Liu, Jiatong Li, Hang Li, Hui Liu, Jiliang Tang, and Qing Li. Generative diffusion models on graphs: Methods and applications. *arXiv preprint arXiv:2302.02591*, 2023.
- [40] Chengyi Liu, Jiahao Zhang, Shijie Wang, Wenqi Fan, and Qing Li. Score-based generative diffusion models for social recommendations. *arXiv preprint arXiv:2412.15579*, 2024.
- [41] Tianze Luo, Zhanfeng Mo, and Sinno Jialin Pan. Fast graph generative model via spectral diffusion. *arXiv preprint arXiv:2211.08892*, 2022.
- [42] Xinjian Luo, Yangfan Jiang, Fei Wei, Yuncheng Wu, Xiaokui Xiao, and Beng Chin Ooi. Exploring privacy and fairness risks in sharing diffusion models: An adversarial perspective. *arXiv preprint arXiv:2402.18607*, 2024.
- [43] Tomoya Matsumoto, Takayuki Miura, and Naoto Yanai. Membership inference attacks against diffusion models. In *IEEE Security and Privacy Workshops*, pages 77–83, 2023.
- [44] Wilnellys Moore and Sarah Frye. Review of hipaa, part 1: history, protected health information, and privacy and security rules. *Journal of nuclear medicine technology*, 47(4):269–272, 2019.
- [45] Christopher Morris, Nils M Kriege, Franka Bause, Kristian Kersting, Petra Mutzel, and Marion Neumann. Tugdataset: A collection of benchmark datasets for learning with graphs. *arXiv preprint arXiv:2007.08663*, 2020.
- [46] T Nathan Mundhenk, Barry Y Chen, and Gerald Friedland. Efficient saliency maps for explainable ai. *arXiv preprint arXiv:1911.11293*, 2019.
- [47] Mathias Niepert, Mohamed Ahmed, and Konstantin Kutikov. Learning convolutional neural networks for graphs. In *International conference on machine learning*, pages 2014–2023, 2016.
- [48] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *International Conference on Artificial Intelligence and Statistics*, pages 4474–4484, 2020.
- [49] Yan Pang and Tianhao Wang. Black-box membership inference attacks against fine-tuned diffusion models. *arXiv preprint arXiv:2312.08207*, 2023.
- [50] Yan Pang, Tianhao Wang, Xuhui Kang, Mengdi Huai, and Yang Zhang. White-box membership inference attacks against diffusion models. *arXiv preprint arXiv:2308.06405*, 2023.
- [51] Badri N Patro, Mayank Lunayach, Shivansh Patel, and Vinay P Namboodiri. U-cam: Visual explanation using uncertainty based class activation maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7444–7453, 2019.
- [52] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710, 2014.
- [53] Maria Rigaki and Sebastian Garcia. A survey of privacy attacks in machine learning. *ACM Computing Surveys*, 2020.
- [54] Ahmed Salem, Yang Zhang, Mathias Humbert, Mario Fritz, and Michael Backes. MI-leaks: Model and data independent membership inference attacks and defenses on machine learning models. *arXiv preprint arXiv:1806.01246*, 2018.
- [55] Yun Shen, Yufei Han, Zhikun Zhang, Min Chen, Ting Yu, Michael Backes, Yang Zhang, and Gianluca Stringhini. Finding mnemon: Reviving memories of node embeddings. *arXiv preprint arXiv:2204.06963*, 2022.
- [56] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *IEEE symposium on security and privacy*, pages 3–18, 2017.
- [57] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [58] Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Understanding and mitigating copying in diffusion models. *arXiv preprint arXiv:2305.20086*, 2023.

- [59] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 2019.
- [60] Anshuman Suri and David Evans. Formalizing and estimating distribution inference risks. *arXiv preprint arXiv:2109.06024*, 2021.
- [61] Karolina Szczepankiewicz, Adam Popowicz, Kamil Charkiewicz, Katarzyna Nalkecz-Charkiewicz, Michal Szczepankiewicz, Sławomir Lasota, Pawel Zawistowski, and Krystian Radlak. Ground truth based comparison of saliency maps algorithms. *Scientific Reports*, 13(1):16887, 2023.
- [62] Shuai Tang, Zhiwei Steven Wu, Sergul Aydore, Michael Kearns, and Aaron Roth. Membership inference attacks on diffusion models via quantile regression. *arXiv preprint arXiv:2312.05140*, 2023.
- [63] Jos Torge, Charles Harris, Simon V Mathis, and Pietro Lio. Diffhopp: A graph diffusion model for novel drug design via scaffold hopping. *arXiv preprint arXiv:2308.07416*, 2023.
- [64] Vu Tuan Truong, Luan Ba Dang, and Long Bao Le. Attacks and defenses for generative diffusion models: A comprehensive survey. *arXiv preprint arXiv:2408.03400*, 2024.
- [65] Gerrit van den Burg and Chris Williams. On memorization in probabilistic deep generative models. *Advances in Neural Information Processing Systems*, 34:27916–27928, 2021.
- [66] Petar Velivcković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *Proceedings of International Conference on Learning Representations*, 2018.
- [67] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. *arXiv preprint arXiv:2209.14734*, 2022.
- [68] Wenjie Wang, Yiyan Xu, Fuli Feng, Xinyu Lin, Xiangnan He, and Tat-Seng Chua. Diffusion recommender model. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 832–841, 2023.
- [69] Xiuling Wang and Wendy Hui Wang. Group property inference attacks against graph neural networks. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, pages 2871–2884, 2022.
- [70] Xiuling Wang and Wendy Hui Wang. Link membership inference attacks against unsupervised graph representation learning. In *Proceedings of the Annual Computer Security Applications Conference*, pages 477–491, 2023.
- [71] Xiuling Wang and Wendy Hui Wang. Gcl-leak: Link membership inference attacks against graph contrastive learning. *Proceedings on Privacy Enhancing Technologies*, 2024.
- [72] Xiuling Wang and Wendy Hui Wang. Subgraph structure membership inference attacks against graph neural networks. *Proceedings on Privacy Enhancing Technologies*, 2024.
- [73] Bang Wu, Xiangwen Yang, Shirui Pan, and Xingliang Yuan. Model extraction attacks on graph neural networks: Taxonomy and realisation. In *Proceedings of the Asia Conference on Computer and Communications Security*, pages 337–350, 2022.
- [74] Fan Wu, Yunhui Long, Ce Zhang, and Bo Li. Linkteller: Recovering private edges from graph neural networks via influence analysis. In *Proceedings of IEEE Symposium on Security and Privacy*, 2022.
- [75] Jiayang Wu, Wensheng Gan, and Philip S Yu. Graph diffusion network for drug-gene prediction. *arXiv preprint arXiv:2502.09335*, 2025.
- [76] Yixin Wu, Ning Yu, Zheng Li, Michael Backes, and Yang Zhang. Membership inference attacks against text-to-image generation models. *arXiv preprint arXiv:2210.00968*, 2022.
- [77] Han Xie, Jing Ma, Li Xiong, and Carl Yang. Federated graph classification over non-iid graphs. *Advances in Neural Information Processing Systems*, 34:18839–18852, 2021.
- [78] Junchi Yan, Xu-Cheng Yin, Weiyao Lin, Cheng Deng, Hongyuan Zha, and Xiaokang Yang. A short survey of recent advances in graph matching. In *Proceedings of the ACM conference on multimedia retrieval*, pages 167–174, 2016.
- [79] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 56(4):1–39, 2023.
- [80] Shengfang Zhai, Huanran Chen, Yinpeng Dong, Jiajun Li, Qingni Shen, Yansong Gao, Hang Su, and Yang Liu. Membership inference on text-to-image diffusion models via conditional likelihood discrepancy. *Advances in Neural Information Processing Systems*, 37:74122–74146, 2025.

---

**Algorithm 1:** Graph reconstruction attack

---

**Input:**  $\Phi$  from an online marketplace or an API, shadow graphs  $G^S$  (optional)

**Output:** The reconstructed graphs  $G^{rec}$

- 1 Execute  $\Phi$  or feed  $G^S$  to  $\Phi$  to obtain the synthetic graphs  $\hat{G}$ ;
  - 2 Create the set of representation differences between the graphs and aligned graphs:  $\mathbf{D} = \{\}$ ;
  - 3 Create the set of reconstructed graphs:  $G^{rec} = \{\}$ ;
  - 4 **for** each  $g_i$  in  $\hat{G}$ ,  $i \in [1, \dots, N]$  **do**
  - 5     **for** each  $g_j$  in  $\hat{G}$ ,  $j \in [1, \dots, N], j \neq i$  **do**
  - 6         Do graph alignment on the graph pair  $(g_i, g_j)$  with REGAL;
  - 7         Find the most similar counterpart  $\hat{g}_j$  for  $g_i$  with Eq. 6 ;
  - 8         Add the average node representation difference between  $g_i$  and  $\hat{g}_j$  to  $\mathbf{D}$ , namely  $\mathbf{D} = \mathbf{D} \cup D(g_i, \hat{g}_j)$ ;
  - 9 Pick the graph pairs with top-k% least differences from  $\mathbf{D}$ , label the picked graph pairs as  $\mathbf{D}'$ ;
  - 10 **for** each graph pair  $(g_i, \hat{g}_j)$  in  $\mathbf{D}'$  **do**
  - 11     Do edge inference with Eq. 7;
  - 12     Add the inferred graph  $g_i^{rec}$  to  $G^{rec}$ , namely  $G^{rec} = G^{rec} \cup g_i^{rec}$
  - 13 Return  $G^{rec}$
- 

---

**Algorithm 2:** Property inference attack

---

**Input:**  $\Phi$  from an online marketplace or an API, shadow graphs  $G^S$  (optional), the property  $\mathbb{P}$  to be inferred

**Output:** The inferred property of the training graphs  $\mathbb{P}(\Phi^{\text{train}})$

- 1 Execute  $\Phi$  or feed  $G^S$  to  $\Phi$  to obtain the synthetic graphs  $\hat{G}$ ;
  - 2 Calculate the property of  $\Phi^{\text{train}}$  based on the graphs in  $\hat{G}$  with formula 9;
  - 3 Return  $\mathbb{P}(\Phi^{\text{train}})$
- 

[81] Mengchun Zhang, Maryam Qamar, Taegoo Kang, Yuna Jung, Chenshuang Zhang, Sung-Ho Bae, and Chaoning Zhang. A survey on graph diffusion models: Generative ai in science for molecule, protein and material. *arXiv preprint arXiv:2304.01565*, 2023.

[82] Minxing Zhang, Ning Yu, Rui Wen, Michael Backes, and Yang Zhang. Generated distributions are all you need for membership inference attacks against generative models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4839–4849, 2024.

[83] Shijie Zhang, Hongzhi Yin, Tong Chen, Zi Huang, Lizhen Cui, and Xiangliang Zhang. Graph embedding for recommendation against attribute inference attacks. In *Proceedings of the Web Conference 2021*, pages 3002–3014, 2021.

[84] Wanrong Zhang, Shruti Tople, and Olga Ohrimenko. Leakage of dataset properties in {Multi-Party} machine learning. In *Proceedings of the USENIX Security Symposium*, pages 2687–2704, 2021.

[85] Zaixi Zhang, Qi Liu, Zhenya Huang, Hao Wang, Chengqiang Lu, Chuanren Liu, and Enhong Chen. Graphmi: Extracting private graph data from graph neural networks. *Proceedings of the International Joint Conference on Artificial Intelligence*, 2021.

[86] Zhikun Zhang, Min Chen, Michael Backes, Yun Shen, and Yang Zhang. Inference attacks against graph neural networks. In *Proceedings of the USENIX Security Symposium*, pages 1–18, 2022.

[87] Zhanke Zhou, Chenyu Zhou, Xuan Li, Jiangchao Yao, Quanming Yao, and Bo Han. On strengthening and defending graph reconstruction attack with markov chain approximation. *arXiv preprint arXiv:2306.09104*, 2023.

## Appendix

### A Pseudo code of our attacks

Algorithm 1, Algorithm 2, and Algorithm 3 show the pseudo code of our graph reconstruction, property inference, and membership inference attack, respectively.

### B Details of Datasets

We use six real-world datasets for attack evaluation in our paper, namely MUTAG, ENZYMES, Ego-small, IMDB-BINARY, IMDB-MULTI, and QM9, which serve as benchmarks for evaluating graph-based models in various domains [45]. Table 5 provides details of these datasets. We obtain Ego-small and QM9 from the implementation provided by [33], while all other datasets are downloaded from the torch\_geometric package. For IMDB-BINARY and IMDB-MULTI datasets that do not have node features, we follow [77] and use one-hot degree features as the node features.

### C More results on property inference attack

**PIA performance on the properties of average number of triangles per node and graph arboricity.** Table 6 presents the absolute differences in the average number of triangles

---

**Algorithm 3: Membership inference attack**

---

**Input:**  $\Phi$  from an online marketplace or an API, shadow graphs  $G^S$ , a target graph  $G$

**Output:** The membership label  $y_G$  of  $G$

- 1 Train a shadow model  $\Phi^S$  using a subset of graphs from  $G^S$ ;
- 2 Attack feature matrix  $\mathcal{A}^{\text{train}} = []$ ;
- 3 **for** each  $G_i^S$  in  $G^S$ ,  $i \in [1, \dots, N]$  **do**
- 4     Feed  $G_i^S$  to  $\Phi^S$  to obtain a set of synthetic graphs  $\hat{G}_i^S$ ;
- 5     Learn the graph embedding,  $emb(G_i^S)$ , of the shadow graph  $G_i^S$  by using AWEs;
- 6     **for** each  $\hat{g}_i^S$  in  $\hat{G}_i^S$  **do**
- 7         Learn the graph embedding,  $emb(\hat{g}_i^S)$ , of  $\hat{g}_i^S$  by using AWEs;
- 8     **if** Attack with  $\mathcal{A}^{\text{train-1}}$  **then**
- 9         **for** each embedding vector  $emb(\hat{g}_i^S)$  **do**
- 10             Calculate the similarity between  $emb(\hat{g}_i^S)$  and  $emb(G_i^S)$  as  $sim_k(emb(\hat{g}_i^S), emb(G_i^S))$ ;
- 11     **if** Attack with  $\mathcal{A}^{\text{train-2}}$  **then**
- 12         **for** each graph pair  $(\hat{g}_i^S, \hat{g}_j^S)$  within  $\hat{G}_i^S$  **do**
- 13             Calculate the similarity between the embedding vectors of  $emb(\hat{g}_i^S)$  and  $emb(\hat{g}_j^S)$  as  $sim_k(emb(\hat{g}_i^S), emb(\hat{g}_j^S))$ ;
- 14     Construct the attack feature,  $A_i^{\text{train}}$ , of the triplet  $(G_i^S, \hat{g}_i^S, y)$  with Eq. 11 if  $\mathcal{A}^{\text{train-1}}$ , otherwise, use Eq. 12, where  $y$  is the known membership label of  $G_i^S$ ;
- 15     Append  $(A_i^{\text{train}}, y)$  to  $\mathcal{A}^{\text{train}}$ ;
- 16 Train the attack classifier with  $\mathcal{A}^{\text{train}}$ ;
- 17 Feed  $G$  to  $\Phi$  to obtain a set of synthetic graphs  $\hat{G}$ ;
- 18 **for** each  $\hat{g}$  in  $\hat{G}$  **do**
- 19     Learn the graph embedding,  $emb(\hat{g})$ , of  $\hat{g}$  by using AWEs;
- 20 **if** Attack with  $\mathcal{A}^{\text{train-1}}$  **then**
- 21     **for** each  $emb(\hat{g})$  **do**
- 22         Calculate the similarity between the embedding vectors of  $emb(\hat{g})$  and  $emb(G)$  as  $sim_k(emb(\hat{g}), emb(G))$ ;
- 23 **if** Attack with  $\mathcal{A}^{\text{train-2}}$  **then**
- 24     **for** each graph pair  $(\hat{g}_i, \hat{g}_j)$  within  $\hat{G}$  **do**
- 25         Calculate the similarity between the embedding vectors of  $emb(\hat{g}_i)$  and  $emb(\hat{g}_j)$  as  $sim_k(emb(\hat{g}_i), emb(\hat{g}_j))$ ;
- 26 Construct the attack feature,  $A^{\text{att}}$  of  $G$  using Eq. 11 when adopting  $\mathcal{A}^{\text{train-1}}$ , otherwise, use Eq. 12;
- 27 Feed  $A^{\text{att}}$  to the trained attack classifier to obtain the predicted membership  $y_G$ ;
- 28 Return  $y_G$

---

per node and the average graph arboricity between the target model’s training set,  $\Phi^{\text{train}}$ , and the inferred values from the generated graphs. We observe that the absolute differ-

ence in the average number of triangles per node does not exceed 1.005 (less than 3.92% of its original value in  $\Phi^{\text{train}}$ ). Similarly, the absolute difference in average graph arboricity remains below 0.476 (less than 9.05% of its original value in  $\Phi^{\text{train}}$ ). These results demonstrate that our simple yet effective attack can accurately infer the graph properties of  $\Phi^{\text{train}}$ .

---

**Algorithm 4: Graph perturbation in the defense**

---

**Input:** The graphs to be perturbed, namely, the training graphs of GGDM in Defense-1 or the synthetic graphs in Defense-2, are denoted as  $\mathbf{G}$ . The perturbation ratio  $r$ . Note that we perturb only the graph structures of the graphs in  $\mathbf{G}$ , while other properties unrelated to the structure, such as node features, are kept unchanged. Therefore, we omit these properties in the following code.

**Output:** The perturbed graphs  $\mathbf{G}'$

- 1 **for** each graph  $g$  in  $\mathbf{G}$  **do**
- 2     Calculate the perturbation budget  $b = \lfloor r \times |E| \rfloor$ , where  $|E|$  is the number of edges in  $g$ ;
- 3      $\mathbf{P} = \{\}$ , used to label the perturbed entries in the adjacency matrix  $A$  of  $g$ ;
- 4      $\mathbf{G}' = \{\}$ ;
- 5     **while**  $b > 0$  **do**
- 6         Perform a forward pass of  $g$  through GCN for 10 epochs;
- 7         Calculate the node classification loss  $L$  and backpropagate the gradient to each entry in  $A$  using Eq. 14;
- 8         Construct the saliency map  $Saliency(A)$  with Eq. 15;
- 9         Rank the entries in  $A$  based on their importance values in  $Saliency(A)$  in ascending order, labeled as  $A_{\text{ranked}}$ ;
- 10         **for**  $A_{ij}$  in  $A_{\text{ranked}}$  **do**
- 11             **if**  $A_{ij} \notin \mathbf{P}$  **then**
- 12                 Flip the  $A_{ij}$  in  $A$ ;
- 13                  $\mathbf{P} = \mathbf{P} \cup A_{ij}$ ;
- 14                 Continue;
- 15          $b = b - 1$ ;
- 16      $\mathbf{G}' = \mathbf{G}' \cup A$ ;
- 17 Return  $\mathbf{G}'$

---

**PIA performance on frequency ratio distribution across different property ranges.** Figure 9 shows the frequency distribution in different property ranges with a bucket number of  $k = 10$ . We observe that the inferred frequency distribution closely aligns with the original distribution of  $\Phi^{\text{train}}$  across all settings, with absolute ratio differences for each range class varying from 0.0003 to 0.0235 for average graph density, 0 to 0.0401 for average graph degree, 0 to 0.0589 for average

Dataset	Domain	# Graphs	Avg. # Nodes	Avg. # Edges	Avg. # Degrees	# Node features	# Classes
MUTAG	Molecules	188	17.93	19.79	2.11	7	2
Ego-small	Citation network	200	6.41	8.69	2.00	3,703	-
ENZYMES	Proteins	600	32.60	124.3	7.63	3	6
IMDB-BINARY	Social network	1,000	19.77	96.53	9.77	135	2
IMDB-MULTI	Social network	1,500	13.00	65.94	10.14	88	3
QM9	Molecules	133,855	18.03	37.30	4.14	11	-

Table 5: Statistics of datasets

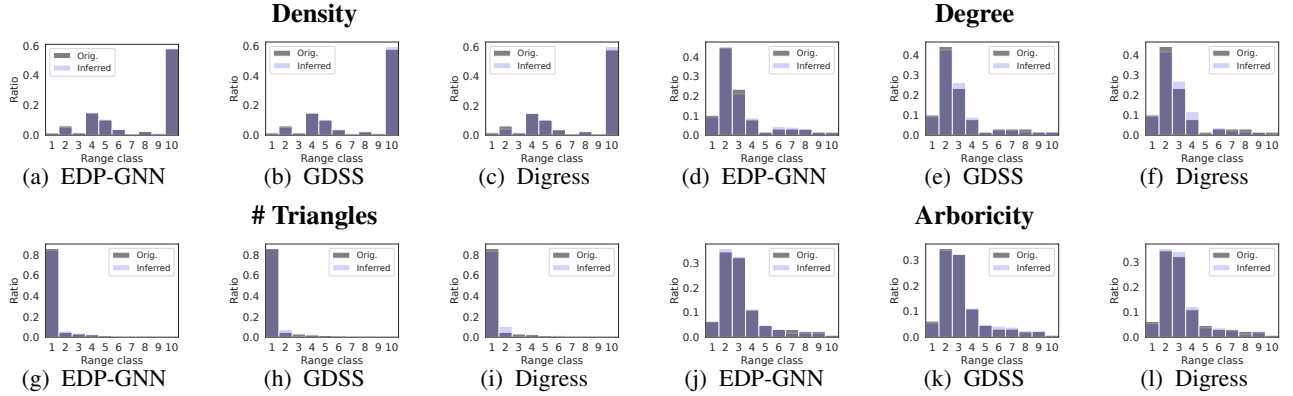


Figure 9: Performance of PIA - the portion of training graphs in different property ranges. Figures (a)–(c) report attack performance on graph density, Figures (d)–(f) on graph degree, Figures (g)–(i) on the average number of triangles per node, and Figures (j)–(l) on graph arboricity ( $k = 10$ , IMDB-M dataset).

Dataset	EDP-GNN		GDSS		Digress		Orig.	
	$D(\text{triangle})\downarrow$	$D(\text{arborivity})\downarrow$	$D(\text{triangle})\downarrow$	$D(\text{arborivity})\downarrow$	$D(\text{triangle})\downarrow$	$D(\text{arborivity})\downarrow$	#Triangles	Arborivity
MUTAG	0.034	0.0001	0.292	0.522	0.398	0.044	0	1
ENZYMES	0.143	0.041	0.212	0.158	0.135	0.245	1.035	1.607
Ego-small	0.203	0.037	0.257	0.174	0.286	0.226	4.605	1.175
IMDB-B	0.914	0.103	0.828	0.351	0.881	0.405	20.167	4.611
IMDB-M	0.448	0.048	0.917	0.351	1.005	0.467	25.627	5.161
QM9	0.374	0.055	0.413	0.219	0.580	0.324	11.185	2.225

Table 6: Performance of property inference attack, namely the absolute differences in property values between the target model’s training set and the inferred values from the generated graphs. A lower difference represents a better performance. " $D(\text{triangle})$ " and " $D(\text{arborivity})$ " represent the absolute difference of average number of triangles per node and graph arborivity, respectively. "Orig." means the property of target model’s training set  $\Phi^{\text{train}}$ .

number of triangles per each node, and 0.0003 to 0.0202 for average graph arboricity.

## D Pseudo code of our defense mechanism

Algorithm 4 shows the pseudo code for the graph perturbation part of our two defense mechanisms.