Private Graph Extraction via Feature Explanations

Iyiola E. Olatunji
L3S Research Center
Hannover, Germany
iyiola@l3s.de

Mandeep Rathee
L3S Research Center
Hannover, Germany
rathee@l3s.de

Thorben Funke
L3S Research Center
Hannover, Germany
tfunke@l3s.de

Megha Khosla
TU Delft
Delft, Netherlands
m.khosla@tudelft.nl

ABSTRACT

Privacy and interpretability are two of the important ingredients for achieving trustworthy machine learning. We study the interplay of these two aspects in graph machine learning through graph reconstruction attacks. The goal of the adversary here is to reconstruct the graph structure of the training data given access to model explanations. Based on the different kinds of auxiliary information available to the adversary, we propose several graph reconstruction attacks. We show that additional knowledge of post-hoc feature explanations substantially increases the success rate of these attacks. Further, we investigate in detail the differences between attack performance with respect to three different classes of explanation methods for graph neural networks: gradient-based, perturbation-based, and surrogate model-based methods. While gradient-based explanations reveal the most in terms of the graph structure, we find that these explanations do not always score high in utility. For the other two classes of explanations, privacy leakage increases with an increase in explanation utility. Finally, we propose a defense based on a randomized response mechanism for releasing the explanations which substantially reduces the attack success rate. Our anonymized code is available at xxxxxxxxx.

KEYWORDS

graph reconstruction attacks, graph neural networks, gnn, explanations

1 INTRODUCTION

Graphs are highly informative, flexible, and natural way of representing data in various real-world domains. Graph neural networks (GNNs) [15, 19, 34] have emerged as the standard tool to analyze graph data which is non-euclidean and irregular in nature. GNNs have gained state-of-the-art results in various graph analytical tasks ranging from applications in biology and healthcare [1, 5] to recommending friends in a social network [9]. GNNs' success can be attributed to their ability to extract powerful latent features via complex aggregation of neighborhood aggregations [13, 42].

However, these models are inherently black-box and complex, making it extremely difficult to understand the underlying reasoning for their predictions. With the growing adoption of these models in various sensitive domains, efforts have been made to explain their decisions in terms of feature as well as neighborhood
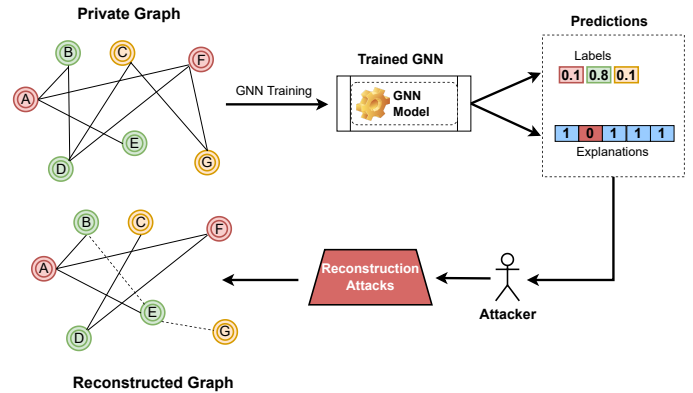
Figure 1: The importance scores for the features as provided by an explanation can be exploited to infer the graph structure. Here, we provide an example of binary explanation where a score of 1 indicates that the corresponding feature is part of the explanation.

attributions. Model explanations can offer insights into the internal decision-making process of the model which builds the trust of the users. Moreover, owing to the current regulations [24] and guidelines for building trustworthy AI systems, several proposals advocate for deploying (automated) model explanations [14, 27].

Nevertheless, releasing additional information such as explanations can have adverse affects on the privacy of the training data. While the risk to privacy due to model explanations exists for machine learning models in general [30], it can have more severe implications for graph neural networks. For instance, several works [6, 23] have established the increased vulnerability of GNNs to privacy attacks due to the additional encoding of graph structure in the model itself. We initiate the *first investigation* of the effect of releasing feature explanations for graph neural networks on *the leakage of private information* in the training data.

To analyze the information leakage due to explanations, we take the perspective of an adversary whose goal is to infer the hidden connections among the training nodes. Consider a setting where the user has access to node features and labels but not the graph structure among the nodes. For example, node features could be part of the public profile of various individuals. The graph structure among the nodes could be some kind of social network which is private. Now, let's say that the user wants to obtain a trained GNN

on her data while providing the node features and labels to a central authority that has access to the private graph structure. We ask here the question: *how much do the feature-based explanations leak information about the graph structure used to train the GNN model?* We quantify this information leakage via several graph reconstruction attacks. A visual illustration is shown in Figure 1. Specifically, our threat model consists of two main settings: first, where the adversary has access only to feature explanations, and second, where the adversary has additional information on node features/labels. Note that we only focus on feature-based explanations for GNNs in this work. For other explanation types such as node or edge explanations, as the returned nodes and edges belong to the node's neighborhood, it becomes trivial to reconstruct the graph structure.

1.1 Our Contributions and Findings

Ours is the first work to analyze the risks of releasing feature explanations in GNNs on the privacy of the relationships/connections in the training nodes. We quantify the information leakage via explanations by the success rate of several graph reconstruction attacks. Our attacks range from simple explanation similarity-based attack to more complex attacks exploiting graph structure learning techniques. Besides, we provide a thorough analysis of information leakage via feature-based explanations produced by three classes of GNN post-hoc explanation methods including *gradient-based*, *perturbation-based* and *surrogate model-based*.

To analyze the differences in the robustness of the explanation methods to our privacy attacks, we investigate the explanation utility in terms of *faithfulness* and *sparsity*. We find that the gradient-based methods are the most susceptible to graph reconstruction attacks even though the corresponding explanations are least faithful to the model. In other words, have low utility. This is an important finding as the corresponding explanations could release a large amount of private information without offering any utility to the user. The perturbation-based approach ZORRO and its variant show the highest explanation utility as well as a high success rate of the attack, pointing to the expected trade-off between privacy and explanation utility.

We perform our study over three types of datasets with varying properties. For instance, our first dataset has a large number of binary features but the feature space is very sparse. Our second dataset has fewer but denser features. Our final dataset has a very small number of features. We find that the information leakage varies with explanation techniques as well as the feature size of the dataset. The dataset with the smallest feature size (8) is the most difficult to attack. All baseline attacks which rely on knowledge of only features and labels perform no better than a random guess. In such a case, explanation-based attacks provide an improvement, though not very huge, in inferring private graph structure information.

Finally, we develop a perturbation-based defense for releasing feature-based explanations. Our defense employs a randomized response mechanism to perturb the individual explanation bits. We show that our defense reduces the attack to a random guess with a small drop in the explanation utility. Our anonymized code is available at xxxxxxxxx.

2 PRELIMINARIES

2.1 Graph Neural Networks

Graph neural networks (GNNs) [15, 19, 34] are a special family of deep learning model designed to perform inference on graph-structured data. The variants of GNNs, such as graph convolutional network (GCN), compute the representation of a node by utilizing its feature representation and that of the neighboring nodes. That is, GNNs compute node representations by recursive aggregation and transformation of feature representations of its neighbors.

Let $G = (V, E)$ denote a graph where V is the node-set, and E represents the edges or links among the nodes. Furthermore, let $x_i^{(\ell)}$ be the feature representation of node i at layer ℓ, $\mathcal{N}_i$ denote the set of its 1-hop neighbors and θ is a learnable weight matrix. Formally, the ℓ-th layer of a graph convolutional operation can be described as

$$z_i^{(\ell)} = \text{AGGREGATION}^{(\ell)} \left(\left\{ x_i^{(\ell-1)}, \left\{ x_j^{(\ell-1)} \mid j \in \mathcal{N}_i \right\} \right\} \right) \quad (1)$$

$$x_i^{(\ell)} = \text{TRANSFORMATION}^{(\ell)} \left(z_i^{(\ell)} \right) \quad (2)$$

Then a softmax layer is applied to the node representations at the last layer (say L) for the final prediction of the node classes C

$$y \leftarrow \text{argmax}(\text{softmax}(z_i^{(L)} \theta)), \quad (3)$$

GNNs have been shown to possess increased vulnerability to privacy attacks due to encoding of additional graph structure in the model [6, 23]. We further investigate the privacy risks of releasing post-hoc explanations for GNN models.

2.2 Explaining Graph Neural Networks

GNNs are deep learning models which are inherently black-box or non-interpretable. This black-box behavior becomes more critical for applying them in sensitive domains like medical, crime, and finance. Consequently, recent works have proposed post-hoc explainability techniques to explain the decisions of an already trained model. In this work, we are concerned with the task of node classification where the model is trained to predict node labels in a graph. For such a task, an instance is a single node. An explanation for a node usually consists of a subset of the most important features as well as a subset of its neighboring nodes/edges responsible for the model's prediction. Depending on the explanation method the importance is usually quantified either as a continuous score (also referred to as a soft mask) or a binary score (also called a hard mask). In this work, we consider three popular classes of explanation methods: *gradient-based* [2, 25, 28], *perturbation-based* [13, 42], and *surrogate* [17] methods.

Gradient-based methods. These approaches usually employ the gradients of the target model's prediction with respect to input features as importance scores for the node features. We use two gradient-based methods in our study, namely **GRAD**[31] and Grad-Input (**GRAD-I**) [32]. For a given graph G and the trained GNN model $f(X, G, \theta)$ (where θ is the set of parameters and X is the features matrix), GRAD generates an explanation $\mathcal{E}_X$ by assigning continuous valued importance scores to the features. For node i and x_i is its features vector, The score is calculated by $\frac{\partial f}{\partial x_i}$. GRAD-I

transforms GRAD explanation by an element-wise multiplication with the input features ($x_i \odot \frac{\partial f}{\partial x_i}$).

Perturbation-based methods. Perturbation-based methods obtain soft or hard masks over the features/nodes/edges as explanations by monitoring the change in prediction with respect to different input perturbations. We use two methods from this class: ZORRO [13] and **GNNEXP** [42]. ZORRO learns discrete masks over input nodes and node features as explanations using a greedy algorithm. It optimizes a fidelity-based objective that measures how the new predictions match the original predictions of the model by fixing the selected nodes/features and replacing the others with random noise values. This returns hard mask for the explanations. GNNEXP learns soft masks over edges and features by minimizing the cross-entropy loss between the predictions of the original graph and the predictions of the newly obtained (masked) graph. We also utilize the ZORRO variant that provides soft explanation masks called **ZORRO-S**. ZORRO-S relaxes the argmax in ZORRO's objective with a softmax, such that the masked are retrievable with standard gradient-based optimization. Together with the regularization terms of GNNEXP, ZORRO-S learns sparse soft masks.

Surrogate methods Surrogate methods fit a simple and interpretable model to a sampled local dataset corresponding to the query node. For example, the sampled dataset can be generated from the neighbors of the given query node. The explanations from the surrogate model are then used to explain the original predictions. We use GraphLime[17] which we denote as **GLIME** from this class. As an interpretable model, it uses the global feature selection method HSIC-Lasso [41]. The sampled dataset consists of the node and its neighborhood. The set of the most important features returned by HSIC-Lasso is used as an explanation for the GNN model.

Remark: Please note that in this work, we assume that only feature-based explanations are released to the user. In the presence of node/edge explanations, an adversary can trivially reconstruct large parts of the neighborhood as the returned nodes/edges are part of the node's original neighborhood.

2.2.1 Measuring explanation quality. We measure the quality of the explanation by its ability to approximate the model's behavior which is referred to as **_faithfulness_**. As the groundtruth for explanations is not available, we use the *RDT-Fidelity* proposed by [13] to measure faithfulness. The corresponding fidelity score measure how the original and new predictions match by fixing the selected nodes/features and replacing the others with random noise values. Formally, the *RDT-Fidelity* of explanation $\mathcal{E}_X$ corresponding to explanation mask $M(\mathcal{E}_X)$ with respect to the GNN f and the noise distribution $\mathcal{N}$ is given by

$$\mathcal{F}(\mathcal{E}_X) = \mathbb{E}_{Y_{\mathcal{E}_X}|Z \sim \mathcal{N}}\left[\mathbb{1}_{f(X)=f(Y_{\mathcal{E}_X})}\right], \quad (4)$$

where the perturbed input is given by

$$\tilde{I}_{\mathcal{E}_X} = X \odot M(\mathcal{E}_X) + Z \odot (\mathbb{1} - M(\mathcal{E}_X)), Z \sim \mathcal{N}, \quad (5)$$

where $\odot$ denotes an element-wise multiplication, and $\mathbb{1}$ is a matrix of ones with the corresponding size.

Sparsity. We further note that by definition, the complete input is faithful to the model. Therefore, we further measure the sparsity of

the explanation. A meaningful explanation should be sparse and should only contain a small subset of features most predictive of the model decision. We use the entropy-based sparsity definition from [13] as it is applicable for both soft and hard explanation masks. Let p be the normalized distribution of explanation (feature) masks. Then the sparsity of an explanation is given by the entropy $H(p)$ over the mask distribution,

$$H(p) = -\sum_{f \in M} p(f) \log p(f).$$

Note that the entropy here is bounded by $\log(|M|)$, where M corresponds to the size of feature set. The lower the entropy, the sparser the explanation. We will utilize these two metrics used in measuring explanation quality to argue about the differences in the attack performance.

3 THREAT MODEL AND ATTACK METHODOLOGY

3.1 Motivation

We consider the setting in which the features and graph (adjacency matrix) are held by different data holders in practice. Specifically, the central trusted server has access of the graph structure and trains a GNN model using the node features and labels provided by the user. The user can further query the trained GNN model by providing node features. As the features are already known to the user, revealing *feature explanations* for the prediction might be considered a safe way to increase the trust of the user in the model. We investigate such a scenario and uncover the increased privacy risks of releasing *feature explanations even if the original node features/labels are already known to the adversary.*

As an example of such a scenario, consider a GNN model trained on a SWIFT (Society for Worldwide Interbank Financial Telecommunication) [33] financial network spanning several financial institutions to detect malicious users. For the sake of transparency, each financial institution in the network has the capability of querying a GNN model trained with all the SWIFT data to easily identify malicious transactions by inputting the message/account number (*node features*) and the model returns a *label* indicating whether the transaction is malicious or not (node classification task). In addition, the central server returns a feature explanation for the model's prediction. Assuming that there exists an attacker (insider of a financial institution with malicious intentions) and has some messages (node features). She is interested in knowing if there was a transaction between two customer accounts (target nodes). Such leakage will put the different customers at risk of estimating their financial worth or publicizing their relationships.

3.2 Threat Model

We consider two scenarios: **first**, in which the adversary has access to node features and or labels and obtains additional access to feature explanations, and **second**, in which the adversary has access only to feature explanations. The goal of the adversary is to infer the private connections of the graph used to train the GNN model.

Corresponding to the above settings, we categorize our attacks as *explanation augmentation* (when explanations are augmented with other information to launch the attack) and *explanation-only*

attacks. Through our five proposed attacks, we investigate the privacy leakage from the explanations generated by six different explanation methods.

Our simple similarity-based explanation-only attack already allows us to quantify the additional information that the feature-based explanation encodes about the graph structure. Our explanation augmentation attacks are based on the graph structure learning paradigm which allows us to effectively integrate additional known information in the learning of the private graph structure. Besides, our explanation augmentation attacks also result in a successfully trained GNN model without the knowledge of the true graph structure offering an additional advantage to the adversary.

3.3 Attack Methodologies

Here, we provide a detailed description of our two types of attacks. We commence with the *explanation-only attack* in which we utilize only the provided explanation to launch the attack followed by the *explanation augmentation* attacks in which more information such as node labels or/and features are exploited in addition to the explanation. The taxonomy of our attacks based on the attacker's knowledge is presented in Table 1.

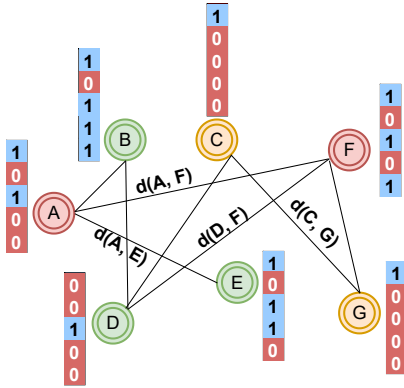**Explanation-only Attack.** This is an unsupervised attack in which

Figure 2: EXPLAINSIM attack. Each node is assigned a feature explanation vector where blue (1) and red (0) indicate whether the feature is part of the explanation or not. The attacker then assigns an edge by computing the pairwise similarity represented as $d(node_i, node_j)$ between nodes' explanation vectors. We show the representation of ZORRO's explanation for easier visualization.

the attacker only has access to the explanations and does not have access to the features or the labels. The attacker measures the distance between each pair of explanation vectors and assigns an edge between them if their distance is small. The intuition is that the model might assign similar labels to connected nodes which lead to similar explanations. We experimented with various distance metrics but cosine similarity performs best across all dataset. We refer to this similarity-based attack as **EXPLAINSIM**. This attack is illustrated in Figure 2.

Explanation Augmentation Attack. Towards explanation aug-

mentation attacks, we leverage the graph structure learning paradigm of [10]. In particular, we employ two *generator modules* for generating graph edges corresponding to features and explanations, respectively. The generators are trained using feature/explanation reconstruction-based losses as well as the node classification loss. We commence by describing the common architecture of the attack model followed by its concrete usage in the four attack variations in Section 3.3.1.

Generators. The two generators take the node features/explanations as input and output two adjacency matrices. We employ the full parameterization (FP) approach to model the generators similar to those used in [10, 12]. In other words, each element of the reconstructed adjacency matrix $\tilde{A}$ is treated as a separate learnable parameter. The adjacency matrix is parameterized using Bernoulli distribution. Let the output adjacency matrix function be given as $\tilde{A} = G_{FP}(X; \theta_G) = \theta_G$ where $\theta_G \in \mathbb{R}^{n \times n}$ and $G_{FP}(\cdot; \cdot)$ denotes the generator function. To obtain a symmetric adjacency matrix with all positive elements, we perform the following transformation

$$A = D^{-\frac{1}{2}} \left(\frac{P_{[0,1]}(\tilde{A}) + P_{[0,1]}(\tilde{A})^T}{2} \right) D^{-\frac{1}{2}},$$

where P is a non-negative function defined by

$$P_{[0,1]}[x] = \begin{cases} 0 & x < 0, \\ 1 & x > 1, \\ x & otherwise. \end{cases}$$

The final adjacency matrix is computed by adding the matrices corresponding to two generators. Any element greater than one is trimmed to 1. The final graph is then generated by sampling from the learnt Bernoulli distributions.

Training with self-supervision and node classification loss. The parameters of the generator/adjacency matrix are trained using a supervised loss as well as feature reconstruction losses. For supervised training, the graph sampled from the generator is fed to a graph convolution network (GCN) to predict node labels. Note that the GCN used here is not the target model. We use the cross-entropy loss (CE) between the predicted labels ($\tilde{Y}$) and the ground-truth labels (Y), $\mathcal{L}_C = CE(Y, \tilde{Y})$.

For self-supervision, we employ *denoising graph autoencoders* which aims at reconstructing node features and explanations given as input noisy features/explanations and the learnt graph structure. We represent our self-supervised task models as M_{DAE} and $M_{DAE_{\mathcal{E}_X}}$ for the node features and explanations respectively. They take $\tilde{X}$ ($\tilde{\mathcal{E}}_X$), the noisy node features (explanations) as input and produces a denoised version of the features (explanations) with the same dimension. The noise is added to random indices, represented by idx, of the node feature (explanation). Let X_{idx} ($\mathcal{E}_{X_{idx}}$) be the true values of the indices, we minimize the following objectives for the node features and explanations respectively:

$$\mathcal{L}_{DAE} = \mathcal{L}(X_{idx}, M_{DAE}(\tilde{X}, A_X; \theta_{M_{DAE}})_{idx}) \quad (6)$$

$$\mathcal{L}_{DAE_{\mathcal{E}_X}} = \mathcal{L}(\mathcal{E}_{X_{idx}}, M_{DAE_{\mathcal{E}_X}}(\tilde{\mathcal{E}}_X, A_{\mathcal{E}_X}; \theta_{M_{DAE_{\mathcal{E}_X}}})_{idx}) \quad (7)$$

where A_X and $A_{\mathcal{E}_X}$ are the generated adjacency matrices corresponding to the node features and explanations respectively and $\mathcal{L}$
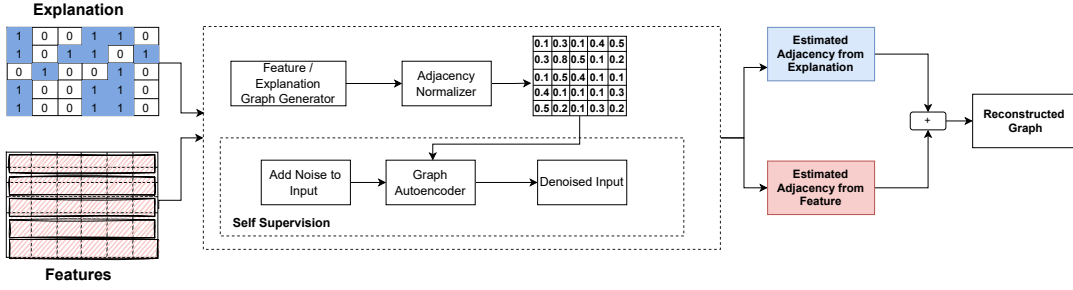
Figure 3: Overview of GSEF. The generator takes node features and explanations as input and outputs an adjacency matrix which may be non-normalized by the adjacency normalizer. The normalized adjacency matrix is used in predicting both the class labels and reconstructing the node features (explanations) by the denoising autoencoders. The final reconstructed adjacency is the one that minimizes the reconstruction error on each of the node features and explanations, and the loss of the class label prediction.

is either the binary cross entropy loss or the mean squared error depending on the dataset.

The final training loss for private graph extraction is then given by

$$\mathcal{L} = \mathcal{L}_{DAE} + \mathcal{L}_{DAE_{\mathcal{E}_X}} + \mathcal{L}_C.$$

Choice of noise and $\mathcal{L}$. For binary datasets and hard masking explanation (Zorro), we use the binary cross entropy loss as $\mathcal{L}$. We randomly flip $r = 20$ percent of the indices whose values are 1 to 0. For datasets and explanations with continuous values, we add independent Gaussian noise to $r = 20$ percent of the indices to each of the features or explanations. The loss in this case is the mean-squared error.

We refer to the above attack framework as **G**raph **S**tealing with **E**xplanations and **F**eatures (**GSEF**) and the schematic diagram is given in Figure 3. Besides, we have three attack variations which employs a single generator module as described below.

Table 1: Attack taxonomy based on attacker's knowledge of node features (X), labels (Y) and feature explanations ($\mathcal{E}_X$).

Attack	X	Y	$\mathcal{E}_X$
ExplainSim	✗	✗	✓
GSEF	✓	✓	✓
GSEF-concat	✓	✓	✓
GSEF-mult	✓	✓	✓
GSE	✗	✓	✓

3.3.1 Attack variations. Besides GSEF, we have three attack variations that employ explanations (i) **GSEF-concat** in which we concatenate the node features and explanations and feed the concatenated input to a single generator module (ii) **GSEF-mult** in which we perform element-wise multiplication between the features and the explanations and feed them into a single graph generator module. This is equivalent to assigning importance to the node features which emphasize the essential characteristics of the nodes. Similar to GSEF-concat, we reconstruct the adjacency matrix using one generator of Figure 3 and (iii) **GSE** in which the attacker

only has access to the explanations and labels. Here, we also employ only one generator with explanations as input.

4 EXPERIMENTS

In this section, we present the experimental results to show the effectiveness of explanation-based attacks. Specifically, our experiments are designed to answer the following research questions:

RQ 1. *How does the knowledge of feature explanations influence the reconstruction of private graph structure?*

RQ 2. *What are the differences between explanation methods with respect to privacy leakage?*

RQ 3. *What is the additional advantage of the adversary (for example in terms of the utility of the inferred information on a downstream task) on explanation augmentation attacks?*

RQ 4. *How much does the lack of knowledge about groundtruth node labels affect attack performance?*

4.1 Experimental Settings

4.1.1 Attack baselines without explanations.

FeatureSim. In this unsupervised attack, the attacker computes the pairwise similarity between pairs of the actual features to reconstruct the graph. Specifically, an edge exists between two nodes if the distance between their feature representation is low. We use the cosine similarity as a measure of similarity because it performs better than other distance metrics.

Lsa *[16].* Lsa (Link stealing attack) is a black-box attack that assumes that the attacker has access to a dataset drawn from a similar distribution as that of the target data (shadow dataset). Additionally, LSA knows the architecture of the target model and can train a corresponding shadow model that replicates the behavior of the target model. The goal of the attack is to infer sensitive links between nodes of interest. We compare our results with that of their proposed attack-2 where an attacker has access to the node features and labels. They trained a separate MLP model (reference model) using the available target attributes and their corresponding labels to obtain posteriors. Then, LSA computes the pairwise distance

between posteriors obtained from the target model and that of the reference model for the nodes of interest. We use cosine similarity as the distance metric.

GraphMI *[44].* GraphMI is a white-box attack in which the attacker has access to the parameters of the target model, node features, all the node labels, and other auxiliary information like edge density. The goal of an attacker is to reconstruct the sensitive links or connections between nodes. The attack model uses the cross entropy loss between the true labels and the output posterior distribution of the target model along with feature smoothness and adjacency sparsity constraints to train a fully parameterized adjacency matrix. The graph is then reconstructed using the graph autoencoder module in which the encoder is replaced by learnt parameters of the target model and the decoder is a logistic function.

Slaps *[10].* Since our attack model is built on top of the graph structure learning framework of SLAPS, we performed an experiment using the vanilla SLAPS. Given node features and labels, the goal of SLAPS is to reconstruct the graph that works best for the node classification task.

4.1.2 Target GNN model. We employ a 2-layer graph convolution network (GCN) [19] as our target GNN model. We use a learning rate of 0.001 and trained for 200 epochs.

4.1.3 Evaluation metrics. Following the existing works [16, 44], we use the area under the receiver operating cost curve (AUC) and average precision (AP) to evaluate our attack. For all experiments (including baselines), we randomly sample equal number pairs of connected and unconnected nodes from the original graph and the predicted graph for evaluation. We measure both the AUC and the AP on these randomly selected node pairs. All our experiments were conducted for 10 different instantiations using PyTorch Geometric library [11] on 11GB GeForce GTX 1080 Ti GPU, and we report the mean values across all runs.

4.1.4 Datasets. We use three commonly used datasets chosen based on varying graph properties such as their feature dimensions and structural properties. The task on all datasets is node classification. We present the data statistics in Table 2.

Cora*.* The Cora dataset [29] is a citation dataset where each research article is a node and there exists an edge between two articles if one article cites the other. Each node has a label that shows the article category. The features of each node are represented by a 0/1-valued word vector which indicates the presence or absence of the word from the abstract of the article.

CoraML*.* In CoraML dataset [29], each node is research article and its abstract is available as raw text. In contrast to the above dataset, the raw text of the abstract is transformed into a dense feature representation. We preprocess the text by removing stop words, web-page links and special characters. Then, we generate Word2Vec [21] embedding of each word. Finally, we generate the feature vector by taking the average over the embedding of all words in the abstract.

Bitcoin*.* The Bitcoin-Alpha dataset[20] is a signed network of trading accounts. Each account is represented as a node and there is a weighted edge between any two accounts which represents

Table 2: Dataset statistics. $|V|$ **and** $|E|$ **denotes the number of nodes and edges respectively,** C, X_d, **and deg denotes the number of classes, size of feature dimension and the average degree of the corresponding graph dataset.**

	Cora	CoraML	Bitcoin
$|V|$	2708	2995	3783
$|E|$	5429	4113	14124
X_d	1433	300	8
C	7	7	2
deg	3.9	2.75	7.5

the trust between accounts. The maximum weight value is +10, indicating total trust and the lowest is -10 which indicates total distrust. Each node is assigned a label which indicates whether the account is trustworthy or not-trustworthy. The features vector of each node is based on the rating by other users such as the average positive or negative rating. We follow the procedures in [35] for generating the feature vectors.

4.2 Result Analysis

*4.2.1 **Analysing the information leakage by explanations (RQ 1)**.* The detailed results of different attacks are provided in Table 3. Our results show that the explanation-only (ExplainSim) and explanation augmentation (GSEF) attacks for all explanation methods other than GLime and GNNExp outperform all baseline methods and by far reveal the most information about the private graph structure. We attribute the superior performance of GSEF to the multi-task learning paradigm that aims to reconstruct both the features and explanations. Our results also supports our assumption that a graph structure that is good for predicting the node labels is also good for predicting the node features and explanations.

In the following we provide an indepth result analysis.

○ **Baseline methods.** Among the baseline methods GraphMI is the best performing attack. This is not very surprising as GraphMI has white-box access to the target GNN in addition to access to node features and labels. On the contrary, the FeatureSim attack which only uses node features shows competitive performance. This highlights the fact that the features alone are very informative of the graph structure of the studied datasets (except Bitcoin in which all baseline attacks almost fail with an AUC score of close to 0.5).

○ **Comparison of the privacy leakage via explanations with that of features.** Figure 4 compares the performance of two similarity-based attacks using explanations (ExplainSim) and features (FeatureSim) respectively. Note that these attacks do not use any other information except explanations and features respectively. Hence, allowing us to compare their information content. We note that except for GLime and GNNExp, ExplainSim outperforms FeatureSim for all datasets except Bitcoin. Moreover, for Bitcoin, both of these attacks fail (with AUC close to 0.5) except for gradient-based explanation methods.

○ **Explanation augmentation with features and labels.** Next, we compare the explanation augmentation attack GSEF with

Table 3: Attack performance and baselines. The best performing attack(s) on each explanation method is(are) highlighted in bold, and the second best attack(s) is(are) underlined.

Exp	Method	Cora		CoraML		Bitcoin	
		AUC	AP	AUC	AP	AUC	AP
Baseline	FeatureSim	0.796	0.822	0.736	0.776	0.536	0.476
	Lsa [16]	0.794	0.829	0.728	0.759	0.530	0.500
	GraphMI [44]	0.859	0.834	0.815	0.810	0.583	0.515
	Slaps [10]	0.716	0.757	0.682	0.738	0.590	0.557
Grad	GSEF-concat	0.694	0.733	0.685	0.749	0.447	0.476
	GSEF-mult	0.692	0.749	0.683	0.762	0.266	0.381
	GSEF	<u>0.947</u>	<u>0.955</u>	**0.902**	<u>0.832</u>	**0.700**	**0.715**
	GSE	0.870	0.893	0.689	0.761	0.254	0.376
	ExplainSim	**0.983**	**0.980**	0.900	**0.904**	<u>0.694</u>	<u>0.656</u>
Grad-I	GSEF-concat	0.700	0.755	0.703	0.753	0.522	0.526
	GSEF-mult	0.665	0.702	0.710	0.743	0.228	0.363
	GSEF	<u>0.914</u>	<u>0.917</u>	<u>0.802</u>	**0.842**	**0.710**	**0.725**
	GSE	0.872	0.900	0.725	0.790	0.256	0.377
	ExplainSim	**0.983**	**0.978**	**0.908**	**0.911**	0.690	0.651
Zorro	GSEF-concat	0.823	0.860	0.735	0.786	<u>0.575</u>	0.529
	GSEF-mult	0.723	0.756	0.681	0.697	0.399	0.449
	GSEF	**0.884**	**0.880**	0.776	0.820	0.537	<u>0.527</u>
	GSE	0.779	0.810	0.722	0.777	**0.596**	**0.561**
	ExplainSim	<u>0.871</u>	<u>0.873</u>	**0.806**	**0.829**	0.427	0.485
Zorro-S	GSEF-concat	0.881	0.913	0.751	0.804	**0.602**	**0.586**
	GSEF-mult	0.752	0.784	0.710	0.727	0.536	0.524
	GSEF	**0.921**	**0.918**	**0.797**	**0.801**	<u>0.595</u>	<u>0.572</u>
	GSE	0.891	0.916	0.774	0.818	0.560	0.561
	ExplainSim	<u>0.912</u>	<u>0.932</u>	0.732	0.804	0.480	0.489
GLime	GSEF-concat	<u>0.634</u>	<u>0.685</u>	<u>0.627</u>	<u>0.664</u>	<u>0.536</u>	<u>0.538</u>
	GSEF-mult	0.517	0.529	0.563	0.570	0.238	0.362
	GSEF	**0.769**	**0.800**	**0.681**	**0.740**	**0.548**	**0.542**
	GSE	0.559	0.588	0.503	0.565	0.262	0.371
	ExplainSim	0.513	0.535	0.522	0.515	0.502	0.498
GNNExp	GSEF-concat	0.600	0.639	0.649	0.677	0.418	0.459
	GSEF-mult	<u>0.703</u>	<u>0.750</u>	<u>0.661</u>	<u>0.720</u>	0.391	0.451
	GSEF	**0.790**	**0.808**	**0.700**	**0.732**	**0.605**	**0.573**
	GSE	0.514	0.540	0.461	0.494	0.322	0.406
	ExplainSim	0.517	0.513	0.498	0.499	<u>0.539</u>	<u>0.523</u>

SUMMARY OF ATTACK COMPARISONS

- The amount of information contained in the explanation alone for the graph structure can be quantified using the ExplainSim attack.

- Explanation only (ExplainSim) and explanation augmentation (GSEF) attacks for all explanation methods other than GLime and GNNExp outperform all baseline methods.

- Among the baseline approaches, the white-box access-based attack, GraphMI, performs the best followed by FeatureSim.

- The relatively good performance of FeatureSim in datasets points to a high correlation of node features with node connections in all datasets other than Bitcoin.

- The information leakage for Bitcoin is limited by small feature size.

- For GLime and GNNExp, we observe that the explanation contains little information about the graph structure. The reason behind this is further revealed in the fidelity-sparsity analysis of the obtained explanations.

the vanilla graph structure learning approach Slaps which only uses node features and labels. GSEF outperforms Slaps (Figure 6) which points to the added utility of using explanations to reconstruct the graph structure.

- **Explanation augmentation attack variants.** In Figure 5, we compare the performance of GSE which have no access to the true features with Slaps which utilize the true features. We observe that on most datasets, the attack performance significantly outperform Slaps. This emphasizes that explanations encode the feature information albeit the importance. Comparing the performance of GSEF-concat and GSEF-mult with GSEF on all

dataset shows that independently extracting the adjacency matrix from the features and explanations respectively and then combining the two adjacency matrix is better than combining the features and explanations at input stage.

- **Attack on Bitcoin.** We observe that all baseline attacks fail for Bitcoin. We attribute this to the very small feature size (=8) and the small number of labels (=2). The attacks are not able to exploit the little information exposed by a small set of features/labels. Explanation-based attacks especially in the case of gradient-based explanations are more successful than baselines but less than for other datasets. The reason can be again attributed to the small feature size. Even though the explanations provide additional

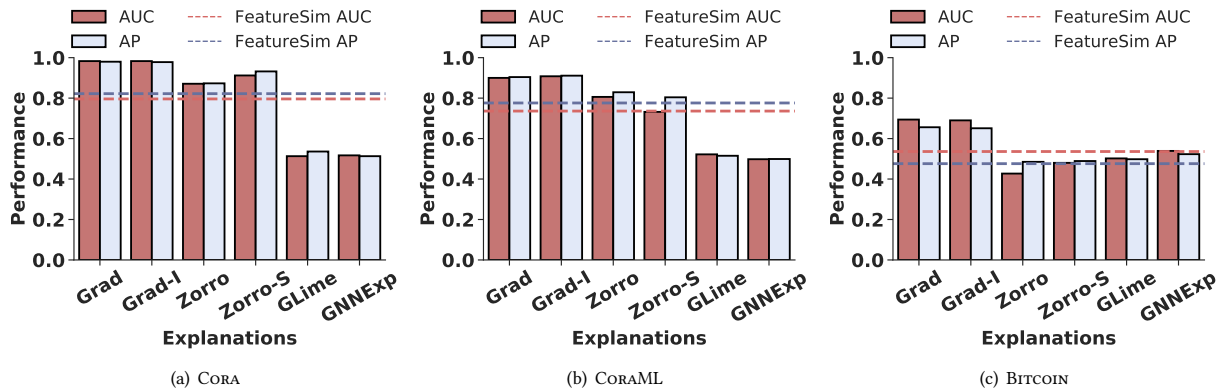(a) Cora   (b) CoraML   (c) Bitcoin

Figure 4: Performance of explanation-only attack (ExplainSim) on the different datasets. The adopted baseline is FeatureSim which performs the pairwise similarities using the true node features.

information, the number of revealed bits in the explanation is equal to the feature dimension which is quite small as compared to other datasets.

4.2.2 Differences in privacy leakage (RQ 2).

All explanation method leaks significant information via the reconstruction attacks except for GLime and GNNExp. We observe that for GLime and GNNExp, the explanation-based attacks do not perform better than the baselines which do not utilize any explanation. Moreover, gradient-based methods are most vulnerable to privacy leakage. To understand the reason behind these observations, we investigate the explanation quality. We measure the goodness of the explanation by its ability to approximate the model's behavior which is also referred to as *faithfulness*. As the groundtruth for explanations is not available, we use the RDT-Fidelity proposed by [13] to measure faithfulness. The results are shown in Table 4. We further note that by definition, the complete input is faithful to the model. Therefore, in addition, we measure the sparsity of the explanation. A meaningful explanation should be sparse and should only contain a small subset of features most predictive of the model decision. We use the entropy-based sparsity definition from [13] as it is applicable for both soft and hard explanation masks. The results are shown in Table 4. We analyse the tradeoffs of privacy and explanation goodness in the following.

- **First**, we observe that GNNExp has the lowest sparsity (the higher the entropy, the more uniform is the explanation mask distribution i.e. higher the explanation density). In other words, almost all features are marked equally important. Hence, it is not surprising that it shows a high fidelity. This is the main reason why ExplainSim fails because there is no distinguishing power contained in the explanations.
- **Second**, we observe that gradient-based explanations (Grad and Grad-I) contain the most information about the graph structure though they have low fidelity, i.e., they do not reflect the model's decision process. It appears that these two methods provide the most similar explanations for connected nodes. This is really the worst case when the explanations

are not useful but leak maximum private information about the graph structure.
- **Third**, GLime has the highest sparsity and lowest fidelity. GLime runs the HSIC-Lasso feature selection method over a local dataset created using the node and its neighborhood. HSIC-Lasso is known to output a very small set of most predictive features [4] when used for global feature selection. But for the current setting of instance-wise feature selection, i.e., finding the most predictive features for decision over an instance/node, GLime's explanation turns out to be too short which is neither faithful to the model nor contains any predictive information about the neighborhood.
- **Finally**, the explanations of Zorro and Zorro-S show the highest fidelity and intermediate sparsity pointing to their high quality. The GSEF attack also obtains high AUC scores for two datasets pointing to the expected increased privacy risk with an increase in explanation utility.

Table 4: RDT-Fidelity and sparsity (entropy) of different explanation methods. For fidelity, the higher the better. For sparsity, the lower the better

Exp	Cora		CoraML		Bitcoin	
	Fidelity	Sparsity	Fidelity	Sparsity	Fidelity	Sparsity
Grad	0.23	3.99	0.22	5.24	0.83	0.64
Grad-I	0.19	3.99	0.20	5.30	0.82	0.64
Zorro	0.89	1.83	0.96	3.33	0.99	0.37
Zorro-S	0.98	2.49	0.84	2.75	0.95	0.96
GLime	0.19	0.88	0.20	0.98	0.82	0.13
GNNExp	0.74	7.27	0.55	5.70	0.90	2.05

4.2.3 Adversary's advantage in terms of trained GNN model (RQ 3).

Here, we formalize a quantitative advantage measure that captures the privacy risk posed by the different attacks. The attacker is at an advantage if she can train a well-performing model (on a downstream task) using the reconstructed graph. As
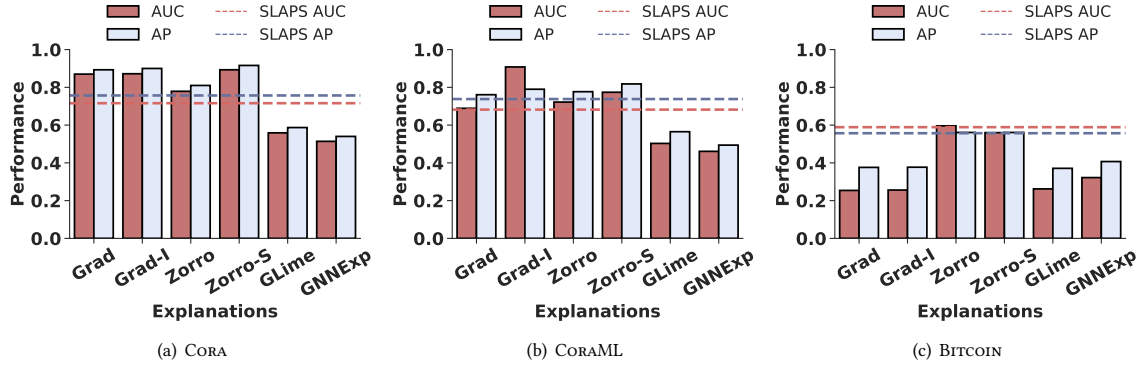
Figure 5: Average AUC and AP of GSE attack on the different datasets. The adopted baseline is SLAPS which use the true node features.
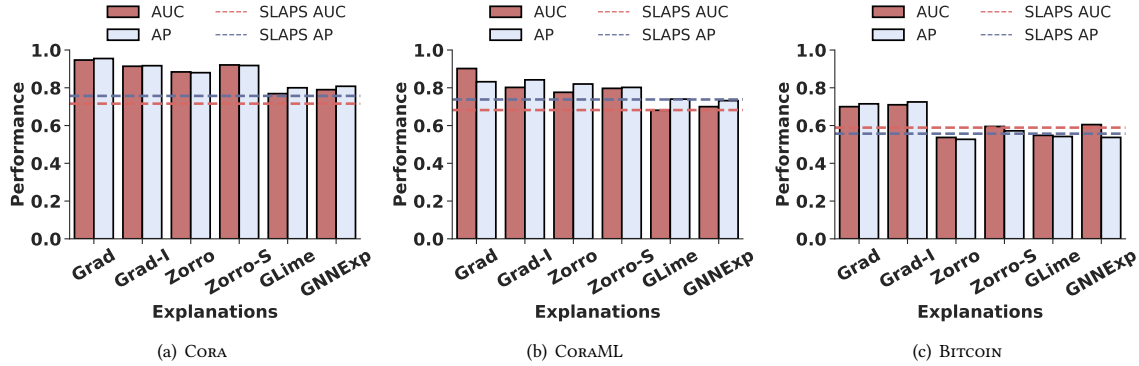

Figure 6: Average AUC and AP of GSEF attack on the different datasets. The adopted baseline is SLAPS which use the true node features.

the attack models based on graph structure learning implicit trains a GNN on the reconstructed graph, we quantify the attacker's advantage by the performance on the downstream task of node classification.

HYPOTHESIS 1. *If the explanations and the reconstructed graph can perform better on a downstream task with high confidence, then the reconstructed adjacency is a valid representation of the graph structure. Hence, the attacker has an advantage quantified by Equation 8.*

We define the attacker's advantage as

$$Advantage = \mathcal{R}(f(\mathcal{E}_X; Adj_{rec}; \theta_W), y), \qquad (8)$$

where f is a 2-layer GCN model parameterized by θ_W, $\mathcal{E}_X$ is the explanation matrix, Adj_{rec} is the reconstructed graph by the attacker, y is the groundtruth label and $\mathcal{R}$ is a advantage measure that compares the predictions on f and the groundtruth label. We use accuracy as the choice of $\mathcal{R}$.

We compare the results to that of SLAPS that use the actual features and groundtruth label for learning the graph structure, and the original performance (denoted by *Max*) in which the model is trained with true features, labels, and graph structure. We analyze

the attacker's advantage corresponding to four attacks GSEF-CONCAT, GSEF-MULT, GSEF, and GSE. The intuition is that if the attacker's advantage is not better than SLAPS, then the best advantage an attacker can have is similar to having the actual feature and performing a graph structure learning. Also, if the attacker's advantage is greater or equal to *Max*, then the attacker has an equivalent advantage as she would have by possessing the actual feature and graph. An example use case of the attacker's advantage is shown in Figure 7. Specifically, if a model trained with say Jane's full data (true features and graph) and another trained only with her explanations (no graph or true features), both models will make the same prediction about Jane. The detailed results for attacker's advantage are plotted in Figure 8. We observe that on CORA, the attacker obtains highest advantage for GRAD, GRAD-I, ZORRO and ZORRO-S explanations. On CORAML, the highest advantage is obtained with GRAD and ZORRO-S explanations. Usually, the attacker's advantage is positively correlated with the success rate of corresponding attacks for both CORA and CORAML. On BITCOIN, the attacker's advantage for all explanation methods is usually high. This is surprising as the success rate for attacks is relatively lower than for
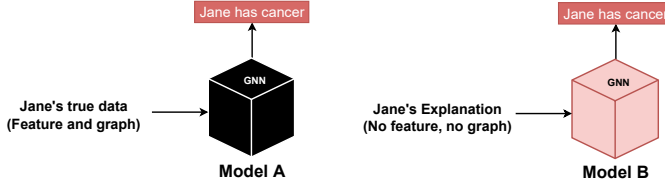
Figure 7: An example use case of the attacker's advantage

other datasets. This might imply that the reconstructed graph has the same semantics as the true graph, if not the exact structure.

4.2.4 ***Lack of groundtruth labels (RQ 4)***. We relax the assumption that the attacker has access to groundtruth labels. Instead, she has black-box access to the target model. This is made possible with the popularity of machine learning as a service (MLaaS) where a user can input a query and get the predictions as output. Therefore, the "groundtruth" label is the one obtained from the target model. As representative explanations, we show the performance on GRAD and ZORRO on all datasets.

As shown in Table 5, on the GRAD explanation, we observe a 3% gain in attack performance on GSEF-CONCAT when the attacker has access to the target model on the CORA and CORAML dataset. On the BITCOIN dataset, we observe a decrease of 2% in AUC and an increase of 6% in AP. The corresponding performance on ZORRO follows the same with no significant change in the performance on CORAML and a 2% decrease in performance on CORA.

The performance of GSEF-MULT and GSEF attack decreases across all datasets with BITCOIN having the worst performance reduction of up to 29% on GRAD. However, on ZORRO, there is a 2% gain on CORA, 4% decrease on CORAML, and upto 36% decrease on BITCOIN. We observe performance drop on GSEF and GSE on all datasets across both explanations except for GSE on GRAD, which has up to 5% gain in performance on the CORA and CORAML datasets. It is important to note that the BITCOIN dataset has the least performance across all attacks even when the true groundtruth label is used. Therefore, the large disparity in performance when the label is generated from the trained black-box model is not surprising.

Summary. For CORA and CORAML dataset, on GSEF-CONCAT and GSE attacks, having black-box access to the target model performs better than the attacker having access to the groundtruth label. For GSEF-MULT and GSEF attacks, it is better to have access to groundtruth label to achieve the best attack success rate. For BITCOIN dataset, the groundtruth labels performs better than blackbox access to target model on all attacks.

5 DEFENSE

Explanation perturbation. To limit the information leakage by the explanation, we perturb each explanation bit using a randomized response mechanism [18, 37]. Specifically for 0/1 feature (explanation) mask as in ZORRO, we flip each bit of the explanation with probability that depends on the privacy budget ϵ as follows

$$Pr(\mathcal{E}'_{x_i} = 1) = \begin{cases} \frac{e^\epsilon}{e^\epsilon+1}, & \text{if } \mathcal{E}_{x_i} = 1, \\ \frac{1}{e^\epsilon+1}, & \text{if } \mathcal{E}_{x_i} = 0, \end{cases} \quad (9)$$

where $\mathcal{E}_{x_i}$ and $\mathcal{E}'_{x_i}$ are true and perturbed i^{th} bit of explanation $\mathcal{E}_x$ respectively. Note that our defense mechanism satisfies $d\epsilon$-local differential privacy.

LEMMA 1. *For an explanation with d dimensions, the explanation perturbation defense mechanism in Equation 9 satisfies $d\epsilon$-local differential privacy.*

PROOF. Note that for an explanation corresponding to two graph datasets D and D' differing in a single edge, the ratio of probabilities of obtaining a certain explanation can be bounded as follows.

$$\frac{Pr[\mathcal{E}_X(D) = S]}{Pr[\mathcal{E}_X(D') = S]} = \prod_{i=1}^d \frac{Pr[\mathcal{E}_{x_i}(D) = S_i]}{Pr[\mathcal{E}_{x_i}(D') = S_i]} \leq \prod_{i=1}^d \frac{\frac{e^\epsilon}{e^\epsilon+1}}{\frac{1}{e^\epsilon+1}} = e^{d\epsilon}.$$

□

Defense evaluation. We evaluate our defense mechanism on EXPLAINSIM attack as it best quantifies the information leakage due to the explanations alone. All other attacks assume the availability of other information such as features and labels. We use two datasets: CORA and CORAML. As evaluation metrics, we use the AUC score and AP to compute the attack success rate after the defense. Besides, we measure the utility of the perturbed explanation in terms of fidelity, sparsity and the percentage of 1 bits that is retained from the original explanation (intersection).

Defense results. As shown in Figure 9, the explanation perturbation based on randomized response mechanism clearly defends against the attack. For instance, at a very high privacy level $\epsilon = 0.0001$ which gives $d\epsilon = 0.14$, the attack performance drastically dropped to 0.56 in AUC and 0.59 in AP which is about 36% decrease over the non-private released explanation. As expected, the attack performance decreases significantly with increase in the amount of noise (ϵ decreases). In Table 6, we analyse the change in explanation utility due to our perturbation mechanism. We observe that on CORA, with lowest privacy loss level, there is a drop of 5.61% in the fidelity when the attack is already reduced to a random guess. The entropy of the mask distribution increases, in other words the explanation sparsity decreases. For ZORRO, this implies that more bits are set to 1 than in the true explanation mask. Even though this decreases explanation utility to some extent, we point out that 74.68% of true explanation is still retained. Moreover, the sparsity is still lower than achieved by GNNEXP explanations even without any perturbations.

While quantitatively, the change in explanation sparsity seems to be acceptable, more application dependent qualitative studies would be required to evaluate the change in utility of explanations. Nevertheless, we provide a promising first defense for future development and possible improvements. We obtain similar results for CORAML which are provided in Appendix A.

Defense variant for soft explanation masks. Note that Equation 9 is only applicable to explanations that returns binary values. For explanations with continuous values, we can adapt the defense as follows. Keeping the original value ($\mathcal{E}_{\mathbf{x}i}$) when the flipped coin lands heads but when it lands tail, replace ($\mathcal{E}_{\mathbf{x}i}$) with $\mathcal{E}'_{x_i}$ where $\mathcal{E}'_{x_i}$ is a random number drawn from normal distribution ($\mathcal{E}'_{\mathbf{x}_i} \sim \mathcal{N}(0, 1)$).
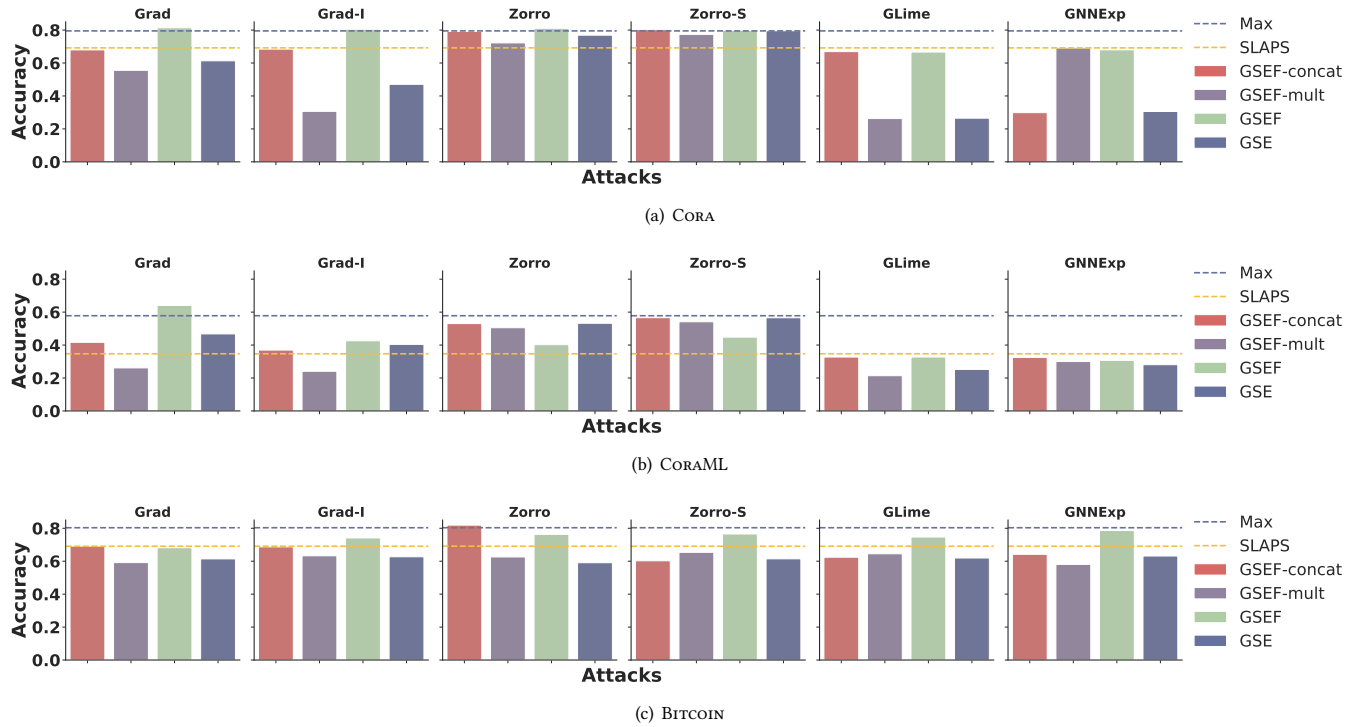
(a) CORA

(b) CORAML

(c) BITCOIN

Figure 8: Accuracy of reconstructed graph on a downstream node classification task by all models on different datasets. Blue line is the original accuracy using the true features and edges while the yellow line is the SLAPS accuracy.

Table 5: Performance comparison of relaxing the availability of groundtruth labels (Y) assumption. Here, the attacker has a black-box access to the target model ($\mathcal{M}$). We perform the experiment on the GRAD and ZORRO explanation method on all datasets. Δ is the percentage difference. Negative value implies that the groudtruth labels are preferred over black-box access.

Dataset	Method	Y_{GRAD}		$\mathcal{M}_{GRAD}$		Δ_{GRAD}		Y_{ZORRO}		$\mathcal{M}_{ZORRO}$		Δ_{ZORRO}	
		AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP	AUC	AP
CORA	GSEF-CONCAT	0.694	0.733	0.717	0.744	3.3	1.5	0.823	0.860	0.779	0.810	-2.8	-5.7
	GSEF-MULT	0.692	0.749	0.671	0.705	-3.1	-5.8	0.723	0.756	0.740	0.772	2.4	2.0
	GSEF	0.947	0.955	0.926	0.935	-2.2	-2.1	0.884	0.880	0.871	0.881	-1.5	0.1
	GSE	0.870	0.893	0.891	0.923	2.4	3.3	0.779	0.810	0.814	0.849	4.5	4.9
CORAML	GSEF-CONCAT	0.685	0.749	0.707	0.780	3.2	4.1	0.735	0.786	0.738	0.792	0.4	0.8
	GSEF-MULT	0.683	0.762	0.666	0.723	-2.5	-5.1	0.681	0.697	0.653	0.692	-4.1	-0.7
	GSEF	0.902	0.832	0.808	0.852	-10.4	2.4	0.776	0.820	0.751	0.796	-3.2	-2.9
	GSE	0.689	0.761	0.725	0.788	5.3	3.6	0.722	0.777	0.713	0.759	-1.2	-2.3
BITCOIN	GSEF-CONCAT	0.447	0.476	0.435	0.504	-2.7	6.0	0.575	0.529	0.523	0.517	-9.0	-2.3
	GSEF-MULT	0.266	0.381	0.208	0.365	-21.7	-4.1	0.399	0.449	0.255	0.369	-36.1	-17.8
	GSEF	0.700	0.715	0.497	0.540	-29.1	-24.5	0.537	0.527	0.312	0.412	-41.8	-21.7
	GSE	0.254	0.376	0.200	0.352	-21.3	-6.4	0.596	0.561	0.491	0.503	-17.6	-10.3

6 RELATED WORKS

Private graph extraction attacks. Given a black-box access to a GNN model that is trained on a target dataset and the adversary's background knowledge, He et al. [16] proposed link stealing attacks to infer whether there is a link between a given pair of nodes in the target dataset. Their attacks are specific to the adversary's background knowledge which range from simply exploiting the node feature similarities to a shadow model-based attack. Wu et al. [39] proposed an edge re-identification attack for vertically partitioned graph learning. Their attack setting is different from ours and is

Table 6: Fidelity, sparsity and percentage of 1 bits in the true explanation that is retained in the perturbed explanation (intersection) after defense for different ϵ on the Cora dataset for Zorro explanation. ∞ implies no privacy.

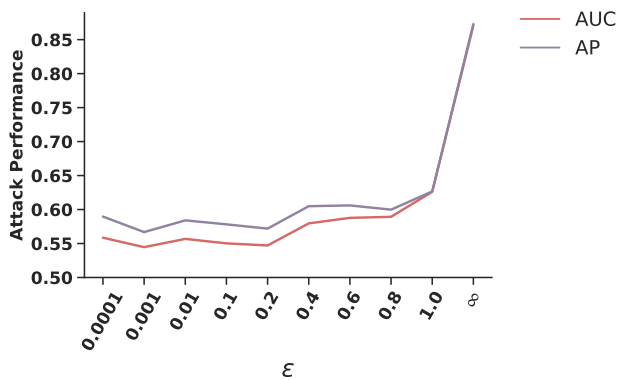| ϵ	Fidelity	Sparsity	Intersection
0.0001	0.84	5.91	74.68
0.001	0.84	5.91	74.70
0.01	0.84	5.89	75.03
0.1	0.84	5.80	75.10
0.2	0.83	5.71	75.60
0.4	0.82	5.49	76.45
0.6	0.81	5.25	77.16
0.8	0.81	5.00	78.66
1	0.81	4.73	80.10
∞	0.89	1.83	100

Figure 9: Privacy budget and corresponding attack performance of ExplainSim for Zorro explanation on the Cora dataset. ∞ implies that no perturbation is performed.

applicable in scenarios where the high-dimensional features and high-order adjacency information are usually heterogeneous and held by different data holders. GraphMI [44] aims at reconstructing the adjacency matrix of the target graph given whitebox access to the trained model, node features and labels.

Other inference attacks and defenses on GNNs. Several other attacks such as membership inference [7, 23] and model extraction attacks [38] have been proposed to quantify privacy leakage in GNNs. In membership inference attack, the goal of the attacker is to infer whether a node was part of the data used in training the GNN model via black-box access to the trained model. In model extraction attack, the attacker aims to steal the trained model's parameter and hyperparameters to duplicate or mimic the functionality of the target model via the predictions returned from querying the model[38]. Recently, several defenses against these attacks have been proposed which are mainly based on differential

privacy. Olatunji et al. [22] proposed a method for releasing GNN models by combining knowledge-distillation framework with two noise mechanisms, random subsampling, and noisy labeling. Their centralized setting approach trains a student model using public graph and private labels obtained from a teacher model trained exclusively for each query node (personalized teacher models). Their method, by design, defends against membership inference attacks and model extraction attack since only the student model (which has limited and perturbed information) is released. Sajadmanesh and Gatica-Perez [26] proposed a locally differentially private GNN model by considering a distributed setting where nodes and labels are private, and the graph structure is known to the central server. Their approach perturbs both the node features and labels to ensure differential privacy guarantee. However, all attacks and defenses are not applicable to explanations.

Membership Inference attack and explanations. On euclidean data such as images, Shokri et al. [30] analyzed the privacy risks of feature-based model explanations using membership inference attacks which quantifies the extent at which model predictions and their explanations leak information about the presence of a datapoint in the training set of a model. We emphasize that the goal of Shokri et al. [30] differ from ours in that we focus on reconstructing the entire graph structure from feature-based explanations. Also their investigations are limited to non-graph data and the corresponding target and explanation models.

Adversarial attacks and GNNs. Another line of research focuses on the vulnerability of GNNs to adversarial attacks [3, 36, 40, 43, 45, 46]. The goal of the attacker is to fool the GNN model into making a wrong prediction by manipulating node features or the structural information of nodes. A recent work [8] used explanation method such as GNNExplainer as a method for detecting adversarial perturbation on graphs. Hence, acting as a tool for inspecting adversarial attacks on GNN models. They further proposed an adversarial attack framework (GEAttack) that exploits the vulnerabilities of explanation methods and the GNN model. This allows the attacker to simultaneously fool the GNN model and misguide the inspection from the explanation method. Our work differs significantly from this work in that first, we aim to reconstruct the graph from the explanations and secondly, to quantify the privacy leakage of explanations on GNN models.

7 CONCLUSION

We initiate the first investigation on the privacy risks of releasing post-hoc explanations of graph neural networks. Concretely, we quantify the information leakage of explanations via our proposed *five* graph reconstruction attacks. The goal of the attacker is to reconstruct the private graph structure information used to train a GNN model. Our results show that even when the explanations alone is available without any additional auxiliary information, the attacker can reconstruct the graph structure with an AUC score of more than 90%. Our explanation-based attacks outperform all baseline methods pointing to the additional privacy risk of releasing explanations. We propose a perturbation-based defense mechanism which reduces the attack to a random guess. The defense leads to a slight decrease in fidelity. At the lowest privacy loss, the perturbed

explanation still contains around 75% of the true explanation. While quantitatively, the change in explanation sparsity seems to be acceptable, more application dependent qualitative studies would be required to evaluate the change in utility of explanations.

We emphasize that we strongly believe in transparency of graph machine learning and acknowledge the need of explaining trained models. At the same time, our work points out the associated privacy risks which cannot be ignored. We believe that our work would encourage future work on finding solutions to balance the complex trade-off between privacy and transparency.

REFERENCES

[1] David Ahmedt-Aristizabal, Mohammad Ali Armin, Simon Denman, Clinton Fookes, and Lars Petersson. 2021. Graph-Based Deep Learning for Medical Diagnosis and Analysis: Past, Present and Future. *Sensors (Basel, Switzerland)* 21, 14 (July 2021). https://doi.org/10.3390/s21144758

[2] Federico Baldassarre and Hossein Azizpour. 2019. Explainability techniques for graph convolutional networks. *arXiv preprint arXiv:1905.13686* (2019).

[3] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial attack on graph structured data. *arXiv preprint arXiv:1806.02371* (2018).

[4] Ngan Thi Dong and Megha Khosla. 2020. Revisiting Feature Selection with Data Complexity. In *2020 IEEE 20th International Conference on Bioinformatics and Bioengineering (BIBE)*. 211–216. https://doi.org/10.1109/BIBE50027.2020.00042

[5] Thi Ngan Dong, Stefanie Mucke, and Megha Khosla. 2022. MuCoMiD: A Multitask graph Convolutional Learning Framework for miRNA-Disease Association Prediction. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* (2022).

[6] Mengnan Du, Ninghao Liu, Qingquan Song, and Xia Hu. 2018. Towards explanation of dnn-based prediction with guided feature inversion. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1358–1367.

[7] Vasisht Duddu, Antoine Boutet, and Virat Shejwalkar. 2020. Quantifying Privacy Leakage in Graph Embedding. *arXiv preprint arXiv:2010.00906* (2020).

[8] Wenqi Fan, Wei Jin, Xiaorui Liu, Han Xu, Xianfeng Tang, Suhang Wang, Qing Li, Jiliang Tang, Jianping Wang, and Charu Aggarwal. 2021. Jointly Attacking Graph Neural Network and its Explanations. *arXiv preprint arXiv:2108.03388* (2021).

[9] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *The World Wide Web Conference*. 417–426.

[10] Bahare Fatemi, Layla El Asri, and Seyed Mehran Kazemi. 2021. SLAPS: Self-Supervision Improves Structure Learning for Graph Neural Networks. *Advances in Neural Information Processing Systems* 34 (2021).

[11] Matthias Fey and Jan Eric Lenssen. 2019. Fast graph representation learning with PyTorch Geometric. *arXiv preprint arXiv:1903.02428* (2019).

[12] Luca Franceschi, Mathias Niepert, Massimiliano Pontil, and Xiao He. 2019. Learning discrete structures for graph neural networks. In *International conference on machine learning*. PMLR, 1972–1982.

[13] Thorben Funke, Megha Khosla, and Avishek Anand. 2021. Zorro: Valid, Sparse, and Stable Explanations in Graph Neural Networks. *arXiv preprint arXiv:2105.08621* (2021).

[14] Bryce Goodman and Seth Flaxman. 2017. European Union regulations on algorithmic decision-making and a "right to explanation". *AI magazine* 38, 3 (2017), 50–57.

[15] William L. Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *NIPS*.

[16] Xinlei He, Jinyuan Jia, Michael Backes, Neil Zhenqiang Gong, and Yang Zhang. 2021. Stealing links from graph neural networks. In *30th USENIX Security Symposium (USENIX Security 21)*. 2669–2686.

[17] Qiang Huang, Makoto Yamada, Yuan Tian, Dinesh Singh, Dawei Yin, and Yi Chang. 2020. Graphlime: Local interpretable model explanations for graph neural networks. *arXiv preprint arXiv:2001.06216* (2020).

[18] Peter Kairouz, Keith Bonawitz, and Daniel Ramage. 2016. Discrete distribution estimation under local privacy. In *International Conference on Machine Learning*. PMLR, 2436–2444.

[19] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*.

[20] Srijan Kumar, Francesca Spezzano, VS Subrahmanian, and Christos Faloutsos. 2016. Edge weight prediction in weighted signed networks. In *2016 IEEE 16th International Conference on Data Mining (ICDM)*. IEEE, 221–230.

[21] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).

[22] Iyiola E Olatunji, Thorben Funke, and Megha Khosla. 2021. Releasing Graph Neural Networks with Differential Privacy Guarantees. *arXiv preprint arXiv:2109.08907* (2021).

[23] Iyiola E Olatunji, Wolfgang Nejdl, and Megha Khosla. 2021. Membership Inference Attack on Graph Neural Networks. In *IEEE International Conference on Trust, Privacy and Security in Intelligent Systems and Applications, TPS-ISA 2021*.

[24] European Parliament and Council of the European Union. 2016. General Data Protection Regulation. *Official Journal of the European Union* (2016). https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=CELEX:32016R0679

[25] Phillip E Pope, Soheil Kolouri, Mohammad Rostami, Charles E Martin, and Heiko Hoffmann. 2019. Explainability methods for graph convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 10772–10781.

[26] Sina Sajadmanesh and Daniel Gatica-Perez. 2020. Locally Private Graph Neural Networks. *arXiv preprint arXiv:2006.05535* (2020).

[27] Andrew Selbst and Julia Powles. 2018. "Meaningful Information" and the Right to Explanation. In *Conference on Fairness, Accountability and Transparency*. PMLR, 48–48.

[28] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*. 618–626.

[29] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.

[30] Reza Shokri, Martin Strobel, and Yair Zick. 2021. On the privacy risks of model explanations. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 231–241.

[31] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv:1312.6034* (2013).

[32] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *PMLR*.

[33] SWIFT. 2022. Swift Fin Traffic & Figures: Swift - the global provider of secure financial messaging services. https://www.swift.com/about-us/discover-swift/fin-traffic-figures

[34] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. *International Conference on Learning Representations* (2018).

[35] Minh Vu and My T Thai. 2020. Pgm-explainer: Probabilistic graphical model explanations for graph neural networks. *Advances in neural information processing systems* 33 (2020), 12225–12235.

[36] Binghui Wang and Neil Zhenqiang Gong. 2019. Attacking graph-based classification via manipulating the graph structure. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 2023–2040.

[37] Yue Wang, Xintao Wu, and Donghui Hu. 2016. Using Randomized Response for Differential Privacy Preserving Data Collection.. In *EDBT/ICDT Workshops*, Vol. 1558. 0090–6778.

[38] Bang Wu, Xiangwen Yang, Shirui Pan, and Xingliang Yuan. 2020. Model Extraction Attacks on Graph Neural Networks: Taxonomy and Realization. *arXiv preprint arXiv:2010.12751* (2020).

[39] Fan Wu, Yunhui Long, Ce Zhang, and Bo Li. 2021. LinkTeller: Recovering Private Edges from Graph Neural Networks via Influence Analysis. *arXiv preprint arXiv:2108.06504* (2021).

[40] Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. 2019. Adversarial examples on graph data: Deep insights into attack and defense. *arXiv preprint arXiv:1903.01610* (2019).

[41] Makoto Yamada, Wittawat Jitkrittum, Leonid Sigal, Eric P Xing, and Masashi Sugiyama. 2014. High-dimensional feature selection by feature-wise kernelized lasso. *Neural computation* 26, 1 (2014), 185–207.

[42] Rex Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. GNN Explainer: A tool for post-hoc explanation of graph neural networks. *arXiv:1903.03894* (2019).

[43] Zaixi Zhang, Jinyuan Jia, Binghui Wang, and Neil Zhenqiang Gong. 2020. Backdoor attacks to graph neural networks. *arXiv preprint arXiv:2006.11165* (2020).

[44] Zaixi Zhang, Qi Liu, Zhenya Huang, Hao Wang, Chengqiang Lu, Chuanren Liu, and Enhong Chen. 2021. GraphMI: Extracting Private Graph Data from Graph Neural Networks. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Zhi-Hua Zhou (Ed.). 3749–3755.

[45] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2847–2856.

[46] Daniel Zügner and Stephan Günnemann. 2019. Adversarial attacks on graph neural networks via meta learning. *arXiv preprint arXiv:1902.08412* (2019).

APPENDIX

A   RESULTS FOR DEFENSE ON CORAML

The attack performance for different values of ϵ is plotted in Figure 10. For the lowest privacy budget, we observe that the attack is reduced to a random guess (with AUC score close to 0.55). The variation in explanation utility and intersection with true explanation is shown in Table 7. Here also, the perturbed explanation is able to retain around 75% of the 1 bits of the true explanation. While there is a drop in fidelity and the explanation becomes denser, we note that the perturbed explanation still shows higher fidelity and sparsity than other explanation methods (c.f. Table 4).
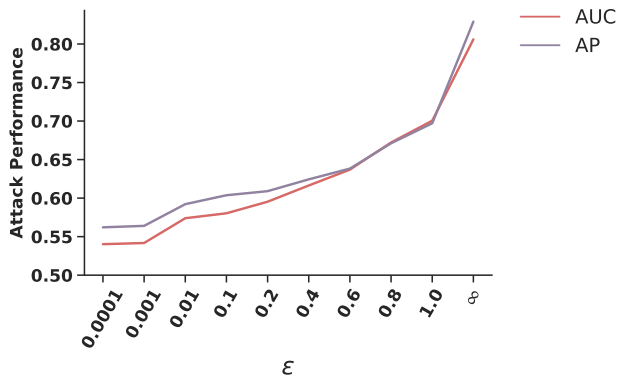
Table 7: Fidelity, sparsity and percentage of 1 bits in the true explanation that is retained in the perturbed explanation (intersection) after defense for different ϵ on the CoraML dataset for Zorro explanation. ∞ implies no privacy.

ϵ	Fidelity	Sparsity	Intersection
0.0001	0.86	4.53	74.96
0.001	0.86	4.53	74.98
0.01	0.86	4.53	75.12
0.1	0.87	4.46	75.08
0.2	0.87	4.39	75.20
0.4	0.87	4.24	75.70
0.6	0.87	4.08	77.07
0.8	0.88	3.95	78.70
1	0.89	3.81	80.75
∞	0.96	3.33	100

Figure 10: Privacy budget and corresponding attack performance of ExplainSim for Zorro explanation on the CoraML dataset. ∞ implies no privacy.