Adversarial Attack on Graph Neural Networks as An Influence Maximization Problem

Jiaqi Ma*
jiaqima@umich.edu
School of Information
University of Michigan
Ann Arbor, Michigan, USA

Junwei Deng*
junweid@umich.edu
School of Information
University of Michigan
Ann Arbor, Michigan, USA

Qiaozhu Mei
qmei@umich.edu
School of Information and
Department of EECS
University of Michigan
Ann Arbor, Michigan, USA

ABSTRACT

Graph neural networks (GNNs) have attracted increasing interests. With broad deployments of GNNs in real-world applications, there is an urgent need for understanding the robustness of GNNs under adversarial attacks, especially in realistic setups. In this work, we study the problem of attacking GNNs in a restricted and realistic setup, by perturbing the features of a small set of nodes, with no access to model parameters and model predictions. Our formal analysis draws a connection between this type of attacks and an influence maximization problem on the graph. This connection not only enhances our understanding on the problem of adversarial attack on GNNs, but also allows us to propose a group of effective and practical attack strategies. Our experiments verify that the proposed attack strategies significantly degrade the performance of three popular GNN models and outperform baseline adversarial attack strategies¹.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Security and privacy**;

KEYWORDS

graph neural networks; adversarial attack; influence maximization

ACM Reference Format:

Jiaqi Ma, Junwei Deng, and Qiaozhu Mei. 2022. Adversarial Attack on Graph Neural Networks as An Influence Maximization Problem. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM '22)*, February 21–25, 2022, Tempe, AZ, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3488560.3498497>

1 INTRODUCTION

There has been a surge of research interest recently in graph neural networks (GNNs) [28], a family of deep learning models on graphs, as they have achieved superior performance on various tasks such

as traffic forecasting [33], social network analysis [16], and recommender systems [7, 32]. Given the successful applications of GNNs in online Web services, there are increasing concerns regarding the robustness of GNNs under adversarial attacks, especially in realistic scenarios. In addition, the research about adversarial attacks on GNNs in turns helps us better understand the intrinsic properties of existing GNN models. Indeed, there have been a line of research investigating various adversarial attack scenarios for GNNs [1, 5, 19, 34, 35], and many of them have been shown to be, unfortunately, vulnerable in these scenarios. In particular, Ma et al. [19] examine an extremely restricted black-box attack scenario where the attacker has access to neither model parameters nor model predictions, yet they demonstrate that a greedy adversarial attack strategy can significantly degrade GNN performance due to the natural inductive biases of GNN binding to the graph structure. This scenario is motivated by real-world GNN applications on social networks, where attackers are only able to manipulate a limited number of user accounts, and they have no access to the GNN model parameters or predictions for the majority of users.

In this work, we study adversarial attacks on GNNs under the aforementioned restricted and realistic setup. Specifically, an attack in this scenario is decomposed into two steps: 1) select a small set of nodes to be perturbed; 2) alter the node features according to domain knowledge up to a per-node budget. The focus of the study lies on the node selection step, so as in Ma et al. [19]. The existing attack strategies, although empirically effective, are largely based on heuristics. We instead formulate the adversarial attack as an optimization problem to maximize the mis-classification rate over the selected set of nodes, and we carry out a formal analysis regarding this optimization problem. The proposed optimization problem is combinatorial and seems hard to solve in its original form. In addition, the mis-classification rate objective involves model parameters which are unknown in the restricted attack setup. We mitigate these difficulties by rewriting the problem and connecting it with *influence maximization* on a variant of *linear threshold model*² [14] related to the original graph structure. Inspired by this connection, we show that, under certain distributional assumptions about the GNN, the expected mis-classification rate is submodular with respect to the selected set of nodes to perturb. The expected mis-classification rate is independent of the model parameters and can be efficiently optimized by a greedy algorithm thanks to its submodularity. Therefore, by specifying concrete distributions, we are able to derive a group of practical attack strategies maximizing

¹Both authors contributed equally to this research.

²Code available at <https://github.com/TheaperDeng/GNN-Attack-InfMax>.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
WSDM '22, February 21–25, 2022, Tempe, AZ, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9132-0/22/02...\$15.00
<https://doi.org/10.1145/3488560.3498497>

²Strictly speaking, the diffusion model we connect with is slightly different from the standard linear threshold model, and should be called a *general threshold model* [8, 20]. But we view it as a (variant of) linear threshold model to avoid redundant notations.

the expected mis-classification rate. The connection with influence maximization also provides us nice interpretations regarding the problem of adversarial attack on GNNs.

To empirically verify the effectiveness of the theory, we implement two adversarial attack strategies and test them on three popular GNN models, Graph Convolutional Network (GCN) [15], Graph Attention Network (GAT) [26], and Jumping Knowledge Network (JKNet) [30] with common benchmark datasets. Both attack strategies significantly outperform baseline attack strategies in terms of decreasing model accuracy. Finally, we summarize the contributions of our study as follows.

- (1) We formulate the problem of adversarial attack on GNNs as an optimization problem to maximize the mis-classification rate.
- (2) We draw a novel connection between the problem of adversarial attacks on GNNs and that of influence maximization on a general threshold model. This connection helps us develop effective adversarial attack strategies under a restricted and realistic attack setup, and provides interpretations regarding the adversarial attack problem.
- (3) We implement two variants of the proposed attack strategies and empirically demonstrate their effectiveness on popular GNNs.

2 RELATED WORK

There has been increasing research interest in adversarial attacks on GNNs recently. Detailed expositions of existing literature are made available in a couple of survey papers [12, 23]. Given the heterogeneous nature of diverse graph structured data, there are numerous adversarial attack setups for GNN models. Following the taxonomy provided by Jin et al. [12], the adversarial attack setup can be categorized based on (but not limited to) the machine learning task, the goal of the attack, the phase of the attack, the form of the attack, and the model knowledge that attacker has access to. First, there are two common types of tasks, node-level classification [5, 6, 27, 34] and graph-level classification [5, 25]. The goal of the attack can be changing the predictions of a small and specific set of nodes (targeted attack) [5, 34] or degrading the overall GNN performance (untargeted attack) [24, 35]. The attack can happen at the model training phases (poisoning attack) [24, 35] or after training completes (evasion attack) [2, 5]. The form of the attack could be perturbing the node features [19, 34] or altering the graph topology [5, 24]. Finally, depending on the knowledge (e.g. model parameters, model predictions, features, and labels, etc.) the attacker has access to, the attacks can be roughly categorized into white-box attacks [29], grey-box attacks [24, 34], or black-box attacks [2, 5, 19]. However, it is worth noting that the borders of these three categories are blurry. In particular, the definition of a “black-box” attack varies much across literature.

Overall, the setup of interest in this paper can be categorized as *node-level, untargeted, evasional, and black-box* attacks by perturbing the node features. While each setup configuration might find its suitable application scenarios, we believe that the black-box setups are particularly important as they are associated with many realistic scenarios. Among the existing studies on node-level black-box attacks, most of them [1, 2, 5] still allow access to model predictions

or some internal representations such as node embeddings. In this paper, we follow the most strict black-box setup [19] to our knowledge, which prohibits any probing of the model³. Compared to Ma et al. [19], we develop attack strategies by directly analyzing the problem of maximizing mis-classification rate, rather than relying on heuristics.

3 PRELIMINARIES

In this section, we introduce necessary notations and backgrounds.

3.1 Notations

We start by introducing notations that will be used across this paper. Suppose we have an attributed graph $G = (V, E, X, y)$, where $V = \{1, 2, \dots, N\}$ is the set of N nodes, $E \subseteq V \times V$ is the set of edges, $X \in \mathbb{R}^{N \times D}$ is the node feature matrix with D -dimensional features, and $y \in \{1, 2, \dots, K\}^N$ is the node label vector with K classes.

We denote a random walk transition matrix on the graph as $M \in \mathbb{R}^{N \times N}$. For any $1 \leq i, j \leq N$,

$$M_{ij} = \begin{cases} \frac{1}{|\mathcal{N}_i|}, & \text{if } (i, j) \in E \text{ or } i = j, \\ 0, & \text{otherwise,} \end{cases}$$

where $\mathcal{N}_i = \{j \in V \mid (i, j) \in E\} \cup \{i\}$ is the set of neighbors of node i , including itself.

To ease the notation, for any matrix $A \in \mathbb{R}^{D_1 \times D_2}$ in this paper, we refer A_j to the transpose of the j -th row of the matrix, i.e., $A_j \in \mathbb{R}^{D_2}$.

We consider a GNN model $f : \mathbb{R}^{N \times D} \rightarrow \mathbb{R}^{N \times K}$ that maps from the node feature matrix X to the output logits of all nodes (denoted as $H \triangleq f(X) \in \mathbb{R}^{N \times K}$). We assume the GNN f has L layers, with the l -th layer ($0 < l < L$) at node i taking the form

$$H_i^{(l)} = \text{ReLU} \left(\sum_{j \in \mathcal{N}_i} \alpha_{ij} W^{(l)} H_j^{(l-1)} \right),$$

where $W^{(l)}$ is the learnable weight matrix, $\text{ReLU}(\cdot)$ is an element-wise ReLU activation function, and different GNNs have different normalization terms α_{ij} . We also define $H^{(0)} = X$ and $H = H^{(L)} = \sum_{j \in \mathcal{N}_i} \alpha_{ij} W^{(L)} H_j^{(L-1)}$. Later in Section 4, we carry out our analysis on a GCN model with $\alpha_{ij} = 1/|\mathcal{N}_i|$ [9].

3.2 The Adversarial Attack Setup

Next we briefly introduce the adversarial attack setup proposed by Ma et al. [19]. The goal of the attack is to perturb the node features of a few carefully selected nodes such that the model performance is maximally degraded. The attack is decomposed into two steps.

In the first step (*node selection step*), the attacker selects a set of nodes $S \subseteq V$ to be perturbed, under two constraints:

$$|S| \leq r \text{ and } |\mathcal{N}_i| \leq m, \forall i \in S,$$

for some $0 < r \ll N$ and $0 < m \ll \max_i |\mathcal{N}_i|$. These two constraints prevent the attacker from manipulating a lot of nodes or

³A similar setting in the computer vision literature [3, 17] is sometimes referred as a “no-box” attack to emphasize that neither model parameters nor model predictions can be accessed.

very important nodes as measured by the node degree, which makes the setup more realistic.

In the second step (*feature perturbation step*), the attacker is allowed to add a small constant perturbation $\epsilon \in \mathbb{R}^D$ to the feature X_i of each node $i \in S$ and obtain the perturbed feature X'_i , i.e.,

$$X'_i \triangleq X_i + \epsilon.$$

The perturbation vector ϵ is constructed based on the domain knowledge about the task but without access to the GNN model. For example, if the GNN model facilitates a recommender system for social media, an attacker may hack a handful of carefully selected users and manipulate their demographic features to get more other users exposed to certain political content the attacker desires. In practice, the perturbation vector ϵ can be tailored for different nodes being manipulated, given personalized knowledge about each node. But following Ma et al. [19], we consider the worst case where no personalization is available and hence ϵ remains constant for each node $i \in S$.

3.3 Influence Maximization on A Linear Threshold Model

Given an information/influence diffusion model on a social network, influence maximization is the problem of finding a small seed set of users such that they spread the maximum amount of influence over the network. In a linear threshold model [14], the influence among nodes is characterized by a weighted directed adjacency matrix $I \in \mathbb{R}^{N \times N}$ where $I_{ij} \geq 0$ for each $(i, j) \in E$ and $I_{ij} = 0$ for each $(i, j) \notin E$. Given a seed set of nodes being activated at initial state, the influence passes through the graph to activate other nodes. There is a threshold vector $\eta \in \mathbb{R}^N$ associated with the nodes, indicating the threshold of influence each node must have received from its active neighbors before it becomes activated. In particular, when the influence propagation comes to a stationary point, a node i outside the seed set will be activated if and only if

$$\sum_{j \in \mathcal{N}_i, j \text{ is activated}} I_{ij} \geq \eta_i. \quad (1)$$

4 ANALYSIS OF THE ADVERSARIAL ATTACK PROBLEM

In this section, we develop principled adversarial attack strategies under the setup stated in Section 3.2.

4.1 Node Selection for Mis-classification Rate Maximization

Suppose an attacker wants to attack a well-trained L -layer GCN model f . Following the two-step attack procedure, the attacker first selects a valid node set $S \in C_{r,m} \triangleq \{T \subseteq V \mid |T| \leq r, |\mathcal{N}_i| \leq m, \forall i \in T\}$ for some given constraints r and m . Then the constant perturbation ϵ is added to the feature of each node in S , which leads to a perturbed feature matrix $X(S, \epsilon)$. Since our primary interest is the design of the node selection step, we shall omit ϵ and just write the perturbed feature as $X(S)$ for simplicity. We denote the output logits of the model after perturbation as $H(S) = f(X(S))$. Clearly, $H(\emptyset)$ equals to the matrix of output logits without attack.

In an untargeted attack, the attacker wants the model to make as many mistakes as possible, which is best measured by the mis-classification rate. Therefore we formulate the problem of selecting the node set as an optimization problem maximizing the mis-classification rate over S , with the two constraints quantified by r, m :

$$\max_{S \in C_{r,m}} \sum_{j=1}^N \mathbf{1} \left[\max_{k=1, \dots, K} H_{jk}(S) \neq H_{jy_j}(S) \right], \quad (2)$$

where $\mathbf{1}[\cdot]$ is the indicator function. We drop normalizing constant $1/N$ in mis-classification rate.

At the first glance, the optimization problem (2) is a combinatorial optimization problem with a complicated objective function involving neural networks. In the following section, we demonstrate that, under a simplifying assumption, it can be connected to the influence maximization problem.

4.2 Connection to the Influence Maximization on A Linear Threshold Model

We first introduce a simplifying assumption of ReLU that has been widely used to ease the analysis of neural networks [4, 13], including GCN [30].

ASSUMPTION 1 (XU ET AL. [30]). *All the ReLU activations activate independently with the same probability, which implies that all paths in the computation graph of the GCN model are independently activated with the same probability of success ρ .*

Under Assumption 1, we are able to define $\bar{H}(S) \triangleq \mathbb{E}_{\text{path}}[H(S)]$ for any $S \subseteq V$, where $\mathbb{E}_{\text{path}}[H(S)]$ indicates the expectation of $H(S)$ over the random activations of ReLU functions in the model. Then we can rewrite problem (2) in a form that is similar to the influence maximization objective on a linear threshold model. The influence weight matrix is defined by the L -step random walk transition matrix $B \triangleq M^L$. And the threshold for each node is related to the original output logits $\bar{H}(\emptyset)$, the perturbation vector ϵ , and the product of the GCN weights $W \triangleq \rho \cdot \prod_{l=L}^1 W^{(l)} \in \mathbb{R}^{K \times D}$. Formally, we have the following Proposition 1, whose proof can be found in Appendix A.1.

PROPOSITION 1. *If we replace $H(\cdot)$ by $\bar{H}(\cdot)$ in problem (2), then we can rewrite the optimization problem as follows,*

$$\max_{S \in C_{r,m}} \sum_{j=1}^N \mathbf{1} \left[\sum_{i \in S} B_{ji} > \theta_j \right], \quad (3)$$

where, for $\hat{k}_j = \arg\max_{k=1, \dots, K} \bar{H}_{jk}(S)$,

$$\theta_j \triangleq \frac{\bar{H}_{jy_j}(\emptyset) - \bar{H}_{j\hat{k}_j}(\emptyset)}{(W_{\hat{k}_j} - W_{y_j})^T \epsilon}. \quad (4)$$

In particular, if $\hat{k}_j = y_j$, we define $\theta_j = \infty$.

Interpretations of the new objective (3). The new optimization objective (3) has nice interpretations. The L -step random walk transition matrix measures the pairwise influence from input nodes to target nodes in the GCN model and $\sum_{i \in S} B_{ji}$ can be viewed as measuring the influence of nodes in S on a target nodes j . In

Figure 1: An illustrative example of the linear threshold model on the derived directed bipartite graph. To simplify the visualization, the GNN is assumed to have 1 layer, and therefore the derived directed bipartite graph have links from its zero-th (itself) and first order neighbors in the original graph. For a GNN with k layers, the derived directed bipartite graph will have links from all its l -th order neighbors in the original graph, for any $0 \leq l \leq k$. Each target node i has its own threshold θ_i to be influenced (mis-classified). The edge weight depends on the random walk transition from the seed node to the target node.

each θ_j , the numerator $\tilde{H}_{jy_j}(\emptyset) - \tilde{H}_{jk_j}(\emptyset)$ can be viewed as the logit margin between the correct class and those wrong classes, which measures the robustness of the prediction on node j . The denominator $(W_{k_j} - W_{y_j})^T \epsilon$ measures how effective the perturbation is. In combination, θ_j measures how difficult it is to mis-classify the node j with perturbation ϵ . This new objective nicely separates the influence between nodes and the node-specific robustness.

Note the form of each term inside the summation over N in Eq. (3), $1 \left[\sum_{i \in S} B_{ji} > \theta_j \right]$, is very similar to that of Eq. (1). In fact, the objective (3) can be viewed as the influence maximization objective on a directed bipartite graph derived from the original graph, as shown in Figure 1. The derived bipartite graph has N nodes on both sides (assuming we call them the seed candidate side $\mathcal{S}$ and target node side $\mathcal{T}$), and there are edges pointing from side $\mathcal{S}$ to side $\mathcal{T}$ but not the converse way. The edge weight from the node i on the side $\mathcal{S}$ to the node j on the side $\mathcal{T}$ ($1 \leq i, j \leq N$) is defined as B_{ji} . Then it is easy to see that the problem (3) is equivalent to the influence maximization problem on the bipartite graph with the node-specific thresholds being $\theta_j, j = 1, \dots, N$.

Two difficulties for solving the problem (3). While we now have got better interpretations of the original mis-classification rate maximization problem in terms of influence maximization, we still face two major difficulties before we can develop an algorithm to solve the problem. The first difficulty is that we do not know the value of θ in a black-box attack setup as it involves the model parameters. The second difficulty is that, even if θ is given, influence maximization on the seemingly simple bipartite graph is still NP-hard, as we show in Lemma 1.

LEMMA 1. *The influence maximization problem on a directed bipartite graph with linear threshold model is NP-hard.*

PROOF. The proof follows similarly as the proof of Theorem 2.4 in Kempe et al. [14]. We prove by reducing the NP-complete Set Cover problem to the influence maximization problem on directed bipartite graph with a linear threshold model. The Set Cover problem is defined as following. Suppose we have a ground set

$U = \{u_1, u_2, \dots, u_n\}$ and a group of m subsets of U , $S_1, S_2, \dots, S_m$. The goal is to determine whether there exists r ($r < n$ and $r < m$) of the subsets whose union equals to U .

For any instance of the Set Cover problem, we can construct a bipartite graph with the first side having m nodes (each one corresponding to a given subset of U), and the second side having n nodes (each one corresponding to an element of U). There are only links going from the first side to the second side. There will be a link with constant influence score $\alpha > 0$ from a node on the first side to the second side if and only if the corresponding subset contains that element in U . Finally the node-specific thresholds of each node on the second side is set as $\alpha/2$. And the influence maximization problem asks to select r nodes on the graph to maximize the number of activated nodes. The Set Cover problem is then solved by deciding if the maximized number of activated nodes on the bipartite graph is greater than $n + r$. $\square$

4.3 Assumptions on the Thresholds

Next, we mitigate the aforementioned two difficulties by making distribution assumptions on the thresholds θ .

It is well-known that if the threshold θ_j of each node j is drawn uniformly at random from the interval $[0, 1]$, the expected objective of a general linear threshold model is submodular, which leads to an efficient greedy algorithm that solves the expected influence maximization problem with a performance guarantee [14]. In light of this fact regarding the general linear threshold model, we show (in Proposition 2) that a mild assumption on the distribution of θ will guarantee the expectation of the objective (3) to be submodular, thanks to the simple bipartite structure. The proof of Proposition 2 can be found in Appendix A.2.

PROPOSITION 2. *Suppose the individual thresholds are random variables drawn from some distributions, and the marginal cumulative distribution function of the threshold θ_j for node j is $F_j, j = 1, \dots, N$. If $F_1, \dots, F_N$ are individually concave in the domain $[0, +\infty)$, then*

the expectation of the objective (3),

$$h(S) \triangleq \mathbb{E}_{\theta_1, \dots, \theta_N} \sum_{j=1}^N \mathbf{1} \left[\sum_{i \in S} B_{ji} > \theta_j \right], \quad (5)$$

is submodular.

Note that here we do not need the thresholds θ to be independent from each other, and we only require the marginal probability density function of each θ_j to be non-increasing on the positive region.

Proposition 2 partially addresses the second difficulty. While we still do not have a solution to the original problem (3), we now know that for a wide range of distributions of θ , the expected mis-classification rate is submodular and can be approximated efficiently through a greedy algorithm.

For the first difficulty, we propose to explicitly specify a simple distribution for θ and optimize the expected mis-classification rate $h(S)$, which no longer involves any model parameters and gives us a black-box attack strategy. While this seems to radically deviate from the original optimization objective (3), in the following Section 5, we empirically show that we only need a crude characterization of the distribution of θ to obtain effective attack strategies.

The concrete black-box attack strategies. Below we derive two concrete black-box attack strategies by specifying the distribution of θ to be uniform distributions and normal distributions respectively.

COROLLARY 1. If $a, b > 0$ and $\theta_j \xrightarrow{i.i.d.} \text{uniform}(-b, a)$, then

$$h(S) = \frac{1}{a+b} \sum_{j=1}^N \left(\min \left(\sum_{i \in S} B_{ji}, a \right) + b \right), \quad (6)$$

and $h(S)$ is submodular.

COROLLARY 2. If $\sigma > 0$ and $\theta_j \xrightarrow{i.i.d.} \mathcal{N}(0, \sigma^2)$, then

$$h(S) = \frac{1}{2} \sum_{j=1}^N \left(1 + \text{erf} \left(\frac{\sum_{i \in S} B_{ji}}{\sigma \sqrt{2}} \right) \right), \quad (7)$$

where $\text{erf}(\cdot)$ is the Gauss error function. And $h(S)$ is submodular.

Corollary 1 and 2 follow directly from Proposition 2 given the cumulative distribution functions of the uniform distribution and the normal distribution as well as the fact that they are concave at the positive region. In particular, Eq. (6) belongs to a well-known submodular function family named the *saturated coverage function* [11, 18]. Under assumptions in Corollary 1, the adversarial attack problem reduces to the classic influence maximization problem under the linear threshold model where the thresholds follow uniform distributions.

We name the attack strategies obtained by greedily maximizing the objectives (6) and (7) as **InfMax-Unif** and **InfMax-Norm** respectively. Specifically, each strategy iteratively selects nodes into the set to be perturbed up to a given size. At each iteration, the node, combining with the existing set, that maximizes Eq. (6) or Eq. (7) will be selected.

4.4 Discussions on the Approximations

From problem (3) to our final attack strategies, we have made two major approximations to address the two difficulties that we raised at the end of Section 4.2.

The first approximation is we go from the original optimization problem to its expected version. Note that θ depends on both the model parameters and the data, which we do not have full access to. The first approximation treats them as random, and takes expectation over θ , which integrates out the randomness in data and the model training process. And the resulted expected objective function $h(S)$ is submodular under the conditions in Proposition 2. A natural question regarding this approximation is how does the mis-classification rate (3) concentrate around its expectation (5)? If θ are independent, the indicator variables in (5) are also independent, and it is easy to show the mis-classification rate is well-concentrated for a large graph size N through Hoeffding's inequality. However, the independence assumption is unrealistic in the case of GNN as the predictions of adjacent nodes should be correlated. Further note that θ can be written in terms of linear combinations of node features. With extra assumptions on the node features and the graph structure, one may be able to carry out finer analysis on the covariance of θ , and thus how well the mis-classification rate concentrates. We leave this analysis for future work.

The second approximation is that we further specify simple distributions of θ , which highly likely deviate much from the real distribution. On one hand, our superior empirical results shown in Section 5 suggest that these simple strategies are practical enough for some applications. On the other hand, this leaves room for further improvement in real-world scenarios if we have more knowledge regarding the distribution of θ . For example, if an attacker has a very limited number of API calls to access the model predictions, these calls are likely insufficient to train a reinforcement-learning-based attack strategies but they can be effectively used to better estimate the distribution of θ .

5 EXPERIMENTS

In this section, we empirically evaluate the performance of the proposed attack strategies (InfMax-Unif and InfMax-Norm) against several baseline attack strategies, closely following the experiment setup of Ma et al. [19]. We also visualize the distributions of θ and have a case study of the selected nodes to gain better qualitative understandings of the proposed methods.

5.1 Attack Strategies for Comparison

Implementation of InfMax-Unif and InfMax-Norm. For the proposed InfMax-Unif and InfMax-Norm, there are two hyper-parameters respectively to be specified. Recall $B = M^L$, the first hyper-parameter for both method is L . We set $L = 4$ following RWCS and GC-RWCS. We note that, for the attack strategies to be effective in practice, the hyper-parameter L does not have to be the same as the number of layers of the GNN being attacked, as we will show in the experiments. For InfMax-Unif, there are two additional distribution hyper-parameters a, b . However, b does not influence the selection of nodes so we only need to specify a . For InfMax-Norm, we need to specify the distribution parameter σ . We fix $a = 0.01$ and $\sigma = 0.01$ across all the experiment setups.

Table 1: Summary of the attack performance in terms of test accuracy (%), the lower the better attack. Bold denotes the best performing strategy in each setup. Underline indicates our strategy outperforms all the baseline strategies. Asterisk (*) means the difference between our strategy and the best baseline strategy is statistically significant by a pairwise t-test at significance level 0.05. The error bar ($\pm$) denotes the standard error of the mean by 40 independent trials. We test on two setups of the node degree threshold, m , by setting it equal to the lowest degree of the top 10% and 30% nodes respectively.

Method	Cora			Citeseer			Pubmed		
	JKNetMaxpool	GCN	GAT	JKNetMaxpool	GCN	GAT	JKNetMaxpool	GCN	GAT
None	85.9 $\pm$ 0.1	85.5 $\pm$ 0.2	87.7 $\pm$ 0.2	73.0 $\pm$ 0.2	75.0 $\pm$ 0.2	74.8 $\pm$ 0.2	85.7 $\pm$ 0.1	85.7 $\pm$ 0.1	85.2 $\pm$ 0.1
Threshold 10%									
Random	70.5 $\pm$ 1.2	77.6 $\pm$ 0.4	71.5 $\pm$ 1.0	59.0 $\pm$ 0.9	69.7 $\pm$ 0.3	70.8 $\pm$ 0.4	74.7 $\pm$ 0.9	79.1 $\pm$ 0.3	72.8 $\pm$ 1.0
Degree	63.9 $\pm$ 1.3	73.4 $\pm$ 0.4	65.9 $\pm$ 1.2	51.1 $\pm$ 0.9	63.9 $\pm$ 0.3	65.6 $\pm$ 0.7	61.2 $\pm$ 1.5	73.2 $\pm$ 0.6	62.2 $\pm$ 1.5
Pagerank	71.6 $\pm$ 0.9	75.7 $\pm$ 0.3	72.7 $\pm$ 0.8	61.3 $\pm$ 0.8	69.5 $\pm$ 0.3	70.4 $\pm$ 0.4	73.7 $\pm$ 0.8	78.7 $\pm$ 0.2	73.3 $\pm$ 0.6
Betweenness	65.2 $\pm$ 1.6	72.4 $\pm$ 0.5	64.3 $\pm$ 1.2	50.8 $\pm$ 1.0	65.6 $\pm$ 0.3	67.1 $\pm$ 0.6	66.9 $\pm$ 1.4	74.6 $\pm$ 0.6	62.9 $\pm$ 1.6
RWCS	75.0 $\pm$ 0.6	75.7 $\pm$ 0.3	69.6 $\pm$ 0.9	61.4 $\pm$ 0.8	69.5 $\pm$ 0.3	70.3 $\pm$ 0.4	73.0 $\pm$ 0.8	78.1 $\pm$ 0.4	71.5 $\pm$ 0.8
GC-RWCS	57.6 $\pm$ 2.0	70.0 $\pm$ 0.6	57.5 $\pm$ 1.5	45.0 $\pm$ 0.8	60.5 $\pm$ 0.4	62.4 $\pm$ 1.0	56.7 $\pm$ 2.0	70.4 $\pm$ 0.8	55.4 $\pm$ 2.2
InfMax-Unif	56.5 $\pm$ 2.1*	68.8 $\pm$ 0.7*	55.9 $\pm$ 1.6*	44.4 $\pm$ 0.8*	60.7 $\pm$ 0.5	62.4 $\pm$ 1.0	55.6 $\pm$ 2.0*	69.5 $\pm$ 0.9*	53.7 $\pm$ 2.3*
InfMax-Norm	56.9 $\pm$ 2.1*	69.1 $\pm$ 0.7*	57.1 $\pm$ 1.5*	45.0 $\pm$ 0.8	60.5 $\pm$ 0.5	62.3 $\pm$ 1.0	54.0 $\pm$ 2.0*	67.5 $\pm$ 1.0*	52.5 $\pm$ 2.4*
Threshold 30%									
Random	73.0 $\pm$ 0.9	78.9 $\pm$ 0.3	74.1 $\pm$ 0.9	61.2 $\pm$ 1.1	71.1 $\pm$ 0.4	71.9 $\pm$ 0.4	77.0 $\pm$ 0.6	80.0 $\pm$ 0.3	75.8 $\pm$ 0.8
Degree	68.8 $\pm$ 1.2	76.3 $\pm$ 0.4	69.2 $\pm$ 1.0	56.0 $\pm$ 1.0	68.4 $\pm$ 0.3	69.1 $\pm$ 0.5	70.9 $\pm$ 1.1	76.9 $\pm$ 0.5	69.8 $\pm$ 1.2
Pagerank	77.2 $\pm$ 0.5	80.5 $\pm$ 0.3	81.9 $\pm$ 0.4	68.2 $\pm$ 0.8	72.4 $\pm$ 0.3	73.0 $\pm$ 0.3	80.3 $\pm$ 0.3	81.5 $\pm$ 0.2	79.6 $\pm$ 0.3
Betweenness	68.9 $\pm$ 1.4	74.3 $\pm$ 0.5	66.5 $\pm$ 1.0	54.1 $\pm$ 1.1	67.9 $\pm$ 0.3	68.7 $\pm$ 0.6	75.5 $\pm$ 0.7	78.4 $\pm$ 0.4	72.5 $\pm$ 1.0
RWCS	77.1 $\pm$ 0.5	80.6 $\pm$ 0.3	81.9 $\pm$ 0.4	68.2 $\pm$ 0.8	72.4 $\pm$ 0.3	73.1 $\pm$ 0.3	78.5 $\pm$ 0.3	80.0 $\pm$ 0.2	77.8 $\pm$ 0.4
GC-RWCS	63.5 $\pm$ 1.8	72.1 $\pm$ 0.5	61.5 $\pm$ 1.3	48.2 $\pm$ 1.0	64.0 $\pm$ 0.4	65.1 $\pm$ 0.8	70.8 $\pm$ 1.0	75.9 $\pm$ 0.6	69.1 $\pm$ 1.3
InfMax-Unif	62.0 $\pm$ 2.0*	71.7 $\pm$ 0.6	59.1 $\pm$ 1.4*	47.2 $\pm$ 0.9*	62.8 $\pm$ 0.4*	64.6 $\pm$ 0.9*	66.3 $\pm$ 1.6*	75.2 $\pm$ 0.6*	64.1 $\pm$ 1.8*
InfMax-Norm	62.3 $\pm$ 2.0*	72.1 $\pm$ 0.6	59.5 $\pm$ 1.4*	47.3 $\pm$ 0.9*	62.8 $\pm$ 0.4*	64.6 $\pm$ 0.8*	66.8 $\pm$ 1.5*	74.5 $\pm$ 0.7*	63.9 $\pm$ 1.8*

Theoretically, the optimal choice of a or σ should depend on the perturbation vector ϵ as well as the dataset. However, we find the proposed InfMax-Unif and InfMax-Norm strategies are fairly robust with respect to the choice of a or σ .

Baseline strategies. We compare with five baseline strategies, **Degree**, **Betweenness**, **PageRank**, Random Walk Column Sum (**RWCS**), and Greedily-Corrected RWCS (**GC-RWCS**).

The first three strategies, as suggested by their names, correspond to three node centrality scores. These strategies select nodes with the highest node centrality scores subject to the constraint $C_{r,m}$. For each node i , the Degree centrality score is defined as $C_D(i) \triangleq \frac{|N_i|}{N}$; the Betweenness centrality score is defined as $C_B(i) \triangleq \sum_{j \neq i, k \neq i, j < k} \frac{g_{jk}(i)}{g_{jk}}$, where g_{jk} is the number of shortest paths connecting node j and k and $g_{jk}(i)$ is the number of shortest paths that node i is on; the PageRank centrality score is defined as the stationary scores achieved by iteratively updating $PR(i) = \frac{1-\alpha}{N} + \alpha \sum_{j \in N_i} \frac{PR(j)}{|N_j|}$ and we set the hyper-parameter $\alpha = 0.85$.

RWCS and GC-RWCS are two black-box attack strategies proposed by Ma et al. [19]. RWCS is derived by maximizing the cross-entropy classification loss but with certain approximations. In practice, RWCS has a simple form: selects nodes with highest importance scores defined as $I(i) = \sum_{j=1}^N [M^L]_{ji}$ (recall that M is the random walk transition matrix). We set the hyper-parameter $L = 4$ following Ma et al. [19]. GC-RWCS further applies a few heuristics on top of RWCS to achieve better mis-classification rate. Specifically, it dynamically updates the RWCS importance score based on a heuristic. It also removes a local neighborhood of the selected

node after selecting each node. In the experiment, we set the hyper-parameters of GC-RWCS $L = 4$, $l = 30$, and $k = 1$ as suggested in their original paper. Interestingly, RWCS can be viewed as a special case of InfMax-Unif if we set $a = \infty$ (or large enough). And GC-RWCS without removing the local neighborhood step can also be viewed a modified version of InfMax-Unif. GC-RWCS further applies a few heuristics on top of RWCS to achieve better mis-classification rate. Specifically, it iteratively selects nodes one by one up to r nodes, based on a dynamic importance score, i.e., $I_t(i) = \sum_{j=1}^N [Q_t]_{ji}$ for the t -th iteration. $Q_t \in \{0, 1\}^{N \times N}$ is a binary matrix that is dynamically updated over t . At the initial iteration, Q_1 is obtained by binarizing M^L , assigning 1 to the top l nonzero entries in each row of M^L and 0 to other entries. For $t > 1$, suppose the node i is selected at the $t - 1$ iteration, then Q_t is obtained from Q_{t-1} by setting to zero for all the rows where the elements of the i -th column is 1 in Q_{t-1} . GC-RWCS also applies another heuristic that, after each iteration, remove the k -hop neighbors of the selected node from the candidate set in the subsequent iterations. In the experiment, we set the hyper-parameters of GC-RWCS $L = 4$, $l = 30$, and $k = 1$ as suggested in their original paper. The iterative-selection process in GC-RWCS (without removing the k -hop neighbors) gives equivalent results as InfMax-Unif if we replace the matrix B in InfMax-Unif by Q_1 and set $a = 1$.

5.2 Attack Experiments on Benchmark Datasets

Experiment setup. We closely follow the experiment setup in Ma et al. [19]. We test attack strategies on 3 popular GNN models, GCN [15], GAT [26], and JK-Net [30], with 3 public benchmark

Figure 2: Each figure shows a histogram of θ_j for a fixed node j over 1000 independent trials of GCN on Cora. The 3 nodes are randomly selected from the union of the validation set and test set.

Figure 3: Visualization of the selected nodes by different strategies on a synthetic graph generated by a stochastic blockmodel with 4 communities (A, B, C, and D). The red nodes are the selected nodes.

datasets, Cora, Citeseer, and Pubmed [22, 31]. We set the number of layers of GCN and GAT as 2, and that of JK-Net as 7. We use the implementations of all models in Deep Graph Library⁴. We randomly split each dataset by 60%, 20% and 20% as the training, validation, and test sets and run 40 independent trials for each model and dataset combination. We apply the attack strategies following the two-step procedure stated in Section 3.2. For the node selection step, we limit the number of nodes to be attacked, r , as 1% of the graph size for each dataset. We test on two setups of the node degree threshold, m , by setting it equal to the lowest degree of the top 10% and 30% nodes respectively.

For the feature perturbation step, we follow a similar way as in Ma et al. [19] to construct the constant perturbation vector ϵ . But our construction is more strictly black-box compared to Ma et al. [19]. Specifically, we first train 20 GCN models on each dataset. Then for each trained GCN (indexed by $k = 1, \dots, 20$), we calculate a set of (signed) importance score $\bar{g}_j^{(k)} = \frac{1}{N} \sum_{i=1}^N \frac{\partial \mathcal{L}(H, y)}{\partial X_{i,j}}$ for each dimension of feature, $j = 1, \dots, D$, and $\mathcal{L}(\cdot, \cdot)$ is the classification loss. Intuitively, $\bar{g}_j^{(k)}$'s are the average gradients over all nodes and provides an indicator of how important a feature is to the classification task at a coarse population level. We then select the top 2% important features by the following procedure: 1) we first filter the features j 's to make sure that at least 80% of $\{\bar{g}_j^{(k)}\}_{k=1}^{20}$ have the same sign; 2) then we select the top $0.02D$ features with largest average scores, $\frac{1}{20} \sum_{k=1}^{20} |\bar{g}_j^{(k)}|$. Then we construct the perturbation vector $\epsilon \in \{-\lambda, 0, \lambda\}^D$ (λ is the perturbation strength and is set as 10 in our experiments) by setting the dimensions corresponding

to the (98%) non-important features as 0, and the dimensions corresponding to the (2%) important features as $+\lambda$ or $-\lambda$, with the sign determined by the majority of $\{\bar{g}_j^{(k)}\}_{k=1}^{20}$. The same constant perturbation vectors are then added to all models in all trials of the experiment. We highlight that the 20 GCN models used to construct the perturbation vector are independently trained with different training splits compared to the models (GCN, GAT, and JK-Net) we attack in Table 1. This makes our experiment setup completely black-box. Notably, when attacking GAT and JK-Net, the construction of perturbation is even unaware of the exact GNN architectures of the victim models.

Note that both the proposed strategies and the baseline strategies are designed for the node selection step, and the feature perturbation step (as detailed below) is the same for all methods.

Construction of the perturbation vector. In a real-world scenario, the perturbation vector should be designed according to domain knowledge about the task (e.g., *age* is an important feature that influences the prediction of *income*). For the benchmark datasets, however, we do not know the semantic meaning of the features. So we use proxy models to identify a handful of important features to be perturbed. In particular, for each dataset, we first independently train 20 GCN models as the proxy models of the task. Then we calculate the gradients of the classification loss with respect to the node features, averaged over all the nodes. We treat the top 2% of features with the largest average gradients as the most important features. And we construct the perturbation vector $\epsilon \in \{-\lambda, 0, \lambda\}^D$ (D is the feature dimension and $\lambda > 0$ is the perturbation strength) by setting the dimensions corresponding to the (98%) non-important features as 0, and the dimensions corresponding to the (2%) important features as $+\lambda$ or $-\lambda$, depending on the sign of the average gradients. The same constant perturbation vector ϵ is then used for

⁴Website: <https://www.dgl.ai>. Apache license: <https://github.com/dmlc/dgl/blob/master/LICENSE>.

attacking all victim models (GCN, GAT, and JK-Net) in all trials on each dataset.

We note that the GCNs used as proxy models are trained independently of the victim models to be attacked, using different data splits. The construction of ϵ is therefore completely agnostic of the model information of the victim models.

Experiment results. We provide the attack experiment results in Table 1. We show the model accuracy after applying each attack strategy in each dataset and model combination, the lower the better. We also include the model accuracy without attack (**None**) and with an attack under random node selection (**Random**) for reference.

As can be seen in Table 1, both the proposed attack strategies achieve better attack performance than all baselines on all but 3 setups, out of the 18 setups in total. And most of the differences are statistically significant. We highlight that, compared to the strongest baseline, GC-RWCS, our methods are simpler, and have fewer hyper-parameters and better interpretation. In addition, the neighbor-removal heuristic also contributes to the performance of GC-RWCS method, while our methods outperform GC-RWCS without such additional heuristics.

5.3 Visualizing the Distributions of θ

We also empirically investigate the distributions of θ to see how likely their PDFs are non-increasing on the positive domain. In particular, given the parameters of a well-trained GCN, we are able to approximately calculate θ with Eq. (4)⁵. We train a GCN on Cora and get one set of θ . We repeat this process with 1000 independent model initializations and get 1000 sets of θ . Then we can visualize a histogram over the 1000 values of θ_j for each node j . In Figure 2, we show the histograms of 3 randomly selected nodes. In most cases, the empirical probability density decreases when $\theta_j > 0$, which is the assumption required for the expected mis-classification rate to be submodular in Proposition 2.

5.4 Case Study of the Selected Nodes

To gain better qualitative understandings of the selected nodes returned by different strategies, we generate a small graph using stochastic blockmodel [10] with 4 communities, so that we can visualize the selected nodes. As can be seen in Figure 3, better performed strategies (GC-RWCS and ours) tend to spread the selected nodes more evenly over the communities. This aligns well with the common belief in the influence maximization literature [21] that a good solution of the selected nodes under linear threshold models tend to have nodes far apart from each other on the graph.

6 CONCLUSION

We present a formal analysis of a restricted and realistic attack setup on Graph Neural Networks (GNNs). By establishing a novel connection between the original attack problem and an influence maximization problem on a linear threshold model, we develop a group of efficient and effective black-box attack strategies with nice interpretations. Extensive empirical results demonstrate the

effectiveness of the proposed methods on multiple types of GNNs, which outperform state-of-the-art attack strategies on multiple types of GNNs. In future work, we plan to explore how to perturb the graph structure under this realistic attack setup, as well as how to perturb node features under extra constraints (e.g. binary or nonnegative).

While the effectiveness of the proposed attack strategies reflects a systematic weakness of existing popular GNN architectures, which could be mis-used for malicious purposes. On the other hand, we believe the interpretable connection between the adversarial attack and the well-known influence maximization problem may motivate more robust GNN designs in future work.

Acknowledgements

This work was in part supported by the National Science Foundation under grant number 1633370.

REFERENCES

- [1] Aleksandar Bojchevski and Stephan Günnemann. 2018. Adversarial attacks on node embeddings via graph poisoning. *arXiv preprint arXiv:1809.01093* (2018).
- [2] Heng Chang, Yu Rong, Tingyang Xu, Wenbing Huang, Honglei Zhang, Peng Cui, Wenwu Zhu, and Junzhou Huang. 2020. A Restricted Black-box Adversarial Framework Towards Attacking Graph Embedding Models. In *AAAI Conference on Artificial Intelligence*.
- [3] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. 2017. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*. 15–26.
- [4] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. 2015. The loss surfaces of multilayer networks. In *Artificial intelligence and statistics*. 192–204.
- [5] 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).
- [6] Negin Entezari, Saba A Al-Sayouri, Amirali Darvishzadeh, and Evangelos E Papalexakis. 2020. All You Need Is Low (Rank) Defending Against Adversarial Attacks on Graphs. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 169–177.
- [7] 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.
- [8] Mark Granovetter. 1978. Threshold models of collective behavior. *American journal of sociology* 83, 6 (1978), 1420–1443.
- [9] Will Hamilton, Zitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in neural information processing systems*. 1024–1034.
- [10] Paul W Holland, Kathryn Blackmond Laskey, and Samuel Leinhardt. 1983. Stochastic blockmodels: First steps. *Social networks* 5, 2 (1983), 109–137.
- [11] Rishabh Iyer and Jeffrey Bilmes. 2015. Submodular point processes with applications to machine learning. In *Artificial Intelligence and Statistics*. 388–397.
- [12] Wei Jin, Yaxin Li, Han Xu, Yiqi Wang, and Jiliang Tang. 2020. Adversarial Attacks and Defenses on Graphs: A Review and Empirical Study. *arXiv preprint arXiv:2003.00653* (2020).
- [13] Kenji Kawaguchi. 2016. Deep learning without poor local minima. In *Advances in neural information processing systems*. 586–594.
- [14] David Kempe, Jon Kleinberg, and Eva Tardos. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*. 137–146.
- [15] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [16] Cheng Li, Jiaqi Ma, Xiaoxiao Guo, and Qiaozhu Mei. 2017. Deepcas: An end-to-end predictor of information cascades. In *Proceedings of the 26th international conference on World Wide Web*. 577–586.
- [17] Qizhang Li, Yiwen Guo, and Hao Chen. 2020. Practical No-box Adversarial Attacks against DNNs. *Advances in Neural Information Processing Systems* 33 (2020).
- [18] Hui Lin and Jeff Bilmes. 2011. A class of submodular functions for document summarization. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. 510–520.
- [19] Jiaqi Ma, Shuangrui Ding, and Qiaozhu Mei. 2020. Towards More Practical Adversarial Attacks on Graph Neural Networks. *Advances in neural information*

⁵We can only do it approximately because we do not know ρ . For the visualization, we just set $\rho = 1$.

- processing systems* (2020).
- [20] Elchanan Mossel and Sebastien Roch. 2007. On the submodularity of influence in social networks. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*. 128–134.
 - [21] Grant Schoenebeck, Biaoshuai Tao, and Fang-Yi Yu. 2019. Think Globally, Act Locally: On the Optimal Seeding for Nonsubmodular Influence Maximization. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)* 145 (2019), 39.
 - [22] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.
 - [23] Lichao Sun, Yingdong Dou, Carl Yang, Ji Wang, Philip S Yu, and Bo Li. 2018. Adversarial attack and defense on graph data: A survey. *arXiv preprint arXiv:1812.10528* (2018).
 - [24] Yiwei Sun, Suhang Wang, Xianfeng Tang, Tsung-Yu Hsieh, and Vasant Honavar. 2019. Node injection attacks on graphs via reinforcement learning. *arXiv preprint arXiv:1909.06543* (2019).
 - [25] Haoteng Tang, Guixiang Ma, Yurong Chen, Lei Guo, Wei Wang, Bo Zeng, and Liang Zhan. 2020. Adversarial Attack on Hierarchical Graph Pooling Neural Networks. *arXiv preprint arXiv:2005.11560* (2020).
 - [26] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.
 - [27] Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. 2019. Adversarial examples for graph data: deep insights into attack and defense. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*. AAAI Press, 4816–4823.
 - [28] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* (2020).
 - [29] Kaidi Xu, Hongge Chen, Sijia Liu, Pin-Yu Chen, Tsui-Wei Weng, Mingyi Hong, and Xue Lin. 2019. Topology attack and defense for graph neural networks: An optimization perspective. *arXiv preprint arXiv:1906.04214* (2019).
 - [30] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. 2018. Representation learning on graphs with jumping knowledge networks. *arXiv preprint arXiv:1806.03536* (2018).
 - [31] Zhilin Yang, William W Cohen, and Ruslan Salakhutdinov. 2016. Revisiting semi-supervised learning with graph embeddings. *arXiv preprint arXiv:1603.08861* (2016).
 - [32] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 974–983.
 - [33] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2017. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *arXiv preprint arXiv:1709.04875* (2017).
 - [34] 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.
 - [35] Daniel Zügner and Stephan Günnemann. 2019. Adversarial Attacks on Graph Neural Networks via Meta Learning. In *International Conference on Learning Representations (ICLR)*.

A PROOFS

We first give a more precise and restated version (Assumption 2) of Assumption 1, and introduce Lemma 2 about GCN, which is proved by Xu et al. [30].

ASSUMPTION 2 (XU ET AL. [30] RESTATED.). Recall that a ReLU function can be written as

$$\text{ReLU}(x) = x \cdot \mathbf{1}[x > 0].$$

Suppose there are R ReLU functions in the GCN model and we index them with $i = 1, 2, \dots, R$. This assumption assumes that the i -th ReLU functions, for $i = 1, 2, \dots, R$, is replaced by the following function,

$$\text{ReLU}_i(x) = x \cdot z_i,$$

where $z_1, z_2, \dots, z_R \stackrel{i.i.d.}{\sim} \text{Bernoulli}(\gamma)$.

This assumption implies that all paths in the computation graph of an L -layer GCN model are independently activated with the same probability $\rho = \gamma^L$.

LEMMA 2 (XU ET AL. [30].) Given an L -layer GCN, under Assumption 1, for any node $i, j \in V$,

$$\mathbb{E}_{\text{path}} \left[\frac{\partial H_j}{\partial X_i} \right] = \rho [M^L]_{ji} \cdot \left(\prod_{l=L}^1 W^{(l)} \right), \quad (8)$$

where $M \in \mathbb{R}^{N \times N}$ is the random walk transition matrix, i.e., for any $1 \leq i, j \leq N$, $M_{ij} = 1/|N_i|$ if $(i, j) \in E$ or $i = j$, and $M_{ij} = 0$ otherwise.

A.1 Proof of Proposition 1

PROOF. Recall that $\bar{H}(S) = \mathbb{E}_{\text{path}}[H(S)] = \mathbb{E}_{\text{path}}[f(X(S))]$. We first show $\bar{H}(S)$ is a linear function of $X(S)$, which suffices to show that, for any $i \in V$ and $1 \leq l \leq L$, $\mathbb{E}_{\text{path}}[H_i^{(l)}(S)]$ is a linear function of $\mathbb{E}_{\text{path}}[H^{(l-1)}(S)]$. When $l = L$,

$$\mathbb{E}_{\text{path}}[H_i^{(L)}(S)] = \sum_{j \in N_i} \alpha_{ij} W^{(L)} \mathbb{E}_{\text{path}}[H_j^{(L-1)}(S)],$$

so the statement holds. When $1 \leq l < L$, under Assumption 1, suppose each ReLU activates independently with probability p .

$$\begin{aligned} \mathbb{E}_{\text{path}} H_i^{(l)} &= \mathbb{E}_{\text{path}} \left[\sigma \left(\sum_{j \in N_i} \alpha_{ij} W^{(l)} H_j^{(l-1)} \right) \right] \\ &= p \sum_{j \in N_i} \alpha_{ij} W^{(l)} \mathbb{E}_{\text{path}}[H_j^{(l-1)}(S)], \end{aligned}$$

so the statement also holds. Therefore $\bar{H}(S)$ is a linear function of $X(S)$. In particular, $\mathbb{E}_{\text{path}}[H] = \bar{H}(\emptyset)$ is a linear function of X .

We know that $X_i(S) = X_i + \epsilon$ for $i \in S$ and $X_i(S) = X_i$ for $i \notin S$. And by Lemma 2, we can rewrite $\bar{H}(S)$ in terms of $\bar{H}(\emptyset)$ and ϵ . For any $j \in V$,

$$\bar{H}_j(S) = \bar{H}_j(\emptyset) + \sum_{i \in S} \rho [M^L]_{ji} \cdot \left(\prod_{l=L}^1 W^{(l)} \right)^T \epsilon.$$

In Section 4.2, we have defined $B = M^L$ and $W = \rho \prod_{l=L}^1 W^{(l)}$, so

$$\bar{H}_j(S) = \bar{H}_j(\emptyset) + W^T \epsilon \sum_{i \in S} B_{ji}. \quad (9)$$

Now we look at the objective (2). If we replace $H(S)$ with $\bar{H}(S)$ in this objective and plug Eq. (9) into it, then for each $j \in V$, we have

$$\begin{aligned} & \mathbf{1} \left[\max_{k \in \{1, \dots, K\}} \bar{H}_{jk}(S) \neq \bar{H}_{jy_j}(S) \right] \\ &= \mathbf{1} \left[\bar{H}_{j\hat{k}_j}(S) > \bar{H}_{jy_j}(S) \right] \\ &= \mathbf{1} \left[\bar{H}_{j\hat{k}_j}(\emptyset) + W_{\hat{k}_j}^T \epsilon \cdot \sum_{i \in S} B_{ji} > \bar{H}_{jy_j}(\emptyset) + W_{y_j}^T \epsilon \cdot \sum_{i \in S} B_{ji} \right] \\ &= \mathbf{1} \left[\sum_{i \in S} B_{ji} > \frac{\bar{H}_{jy_j}(\emptyset) - \bar{H}_{j\hat{k}_j}(\emptyset)}{(W_{\hat{k}_j} - W_{y_j})^T \epsilon} \right] \\ &= \mathbf{1} \left[\sum_{i \in S} B_{ji} > \theta_j \right], \end{aligned}$$

where we have defined $\hat{k}_j = \text{argmax}_{k=1, \dots, K} \bar{H}_{jk}(S)$ and recall the definition of θ_j in Eq. (4).

Therefore we get the optimization problem (3)

$$\max_{S \in C_{r,m}} \sum_{j=1}^N \mathbf{1} \left[\sum_{i \in S} B_{ji} > \theta_j \right].$$

□

A.2 Proof of Proposition 2

PROOF. We first show that the expected mis-classification rate $h(S)$ can be written in terms of the marginal CDFs of θ .

$$\begin{aligned} h(S) &= \mathbb{E}_{\theta_1, \dots, \theta_N} \sum_{j=1}^N \mathbf{1} \left[\sum_{i \in S} B_{ji} > \theta_j \right] \\ &= \sum_{j=1}^N \mathbb{E}_{\theta_1, \dots, \theta_N} \mathbf{1} \left[\sum_{i \in S} B_{ji} > \theta_j \right] \\ &= \sum_{j=1}^N \mathbb{E}_{\theta_j} \mathbf{1} \left[\sum_{i \in S} B_{ji} > \theta_j \right] \\ &= \sum_{j=1}^N P_j \left(\sum_{i \in S} B_{ji} > \theta_j \right) \\ &= \sum_{j=1}^N F_j \left(\sum_{i \in S} B_{ji} \right), \end{aligned}$$

where P_j is the marginal probability of θ_j .

Since $B_{ji} \geq 0$, so $\sum_{i \in S} B_{ji}$ is a non-decreasing submodular function of S with a lower bound 0. Each CDF F_j is non-decreasing by definition, if it is also individually concave at the domain $[0, +\infty)$, we know $F_j \left(\sum_{i \in S} B_{ji} \right)$ is submodular w.r.t. S and hence $h(S)$ is submodular. □

Table 2: Experiment results under the setup by Ma et al. [19]. Summary of the attack performance in terms of test accuracy (%), the lower the better attack. Bold denotes the best performing strategy in each setup. Underline indicates our strategy outperforms all the baseline strategies. Asterisk (*) means the difference between our strategy and the best baseline strategy is statistically significant by a pairwise t-test at significance level 0.05. The error bar ($\pm$) denotes the standard error of the mean by 40 independent trials. We test on two setups of the node degree threshold, m , by setting it equal to the lowest degree of the top 10% and 30% nodes respectively.

Method	Cora			Citeseer			Pubmed		
	JKNet	GCN	GAT	JKNet	GCN	GAT	JKNet	GCN	GAT
None	85.9 $\pm$ 0.1	85.5 $\pm$ 0.2	87.7 $\pm$ 0.2	73.0 $\pm$ 0.2	75.0 $\pm$ 0.2	74.8 $\pm$ 0.2	85.7 $\pm$ 0.1	85.7 $\pm$ 0.1	85.2 $\pm$ 0.1
Threshold 10%									
Random	69.9 $\pm$ 1.1	81.7 $\pm$ 0.3	72.6 $\pm$ 0.6	61.5 $\pm$ 0.9	71.6 $\pm$ 0.2	70.2 $\pm$ 0.5	76.1 $\pm$ 0.6	82.0 $\pm$ 0.3	73.5 $\pm$ 0.3
Degree	63.0 $\pm$ 1.4	78.7 $\pm$ 0.4	66.6 $\pm$ 0.7	53.7 $\pm$ 0.9	68.2 $\pm$ 0.3	63.9 $\pm$ 0.5	63.5 $\pm$ 0.9	78.9 $\pm$ 0.5	65.8 $\pm$ 0.7
Pagerank	71.7 $\pm$ 0.9	80.1 $\pm$ 0.3	74.2 $\pm$ 0.5	62.3 $\pm$ 0.6	70.0 $\pm$ 0.3	69.7 $\pm$ 0.3	71.8 $\pm$ 0.8	80.2 $\pm$ 0.3	71.2 $\pm$ 0.3
Betweenness	63.6 $\pm$ 1.4	80.2 $\pm$ 0.4	64.9 $\pm$ 0.5	54.9 $\pm$ 1.0	70.0 $\pm$ 0.3	65.5 $\pm$ 0.5	67.0 $\pm$ 1.0	78.4 $\pm$ 0.5	62.6 $\pm$ 0.6
RWCS	71.8 $\pm$ 0.8	80.3 $\pm$ 0.4	70.8 $\pm$ 0.5	61.9 $\pm$ 0.6	69.9 $\pm$ 0.3	69.4 $\pm$ 0.3	70.8 $\pm$ 0.8	79.7 $\pm$ 0.3	68.9 $\pm$ 0.4
GC-RWCS	55.2 $\pm$ 1.5	78.3 $\pm$ 0.5	57.1 $\pm$ 0.6	47.5 $\pm$ 1.0	66.3 $\pm$ 0.5	58.5 $\pm$ 0.6	61.7 $\pm$ 1.1	77.4 $\pm$ 0.6	57.8 $\pm$ 0.8
InfMax-Unif	54.3 $\pm$ 1.5*	77.9 $\pm$ 0.5*	55.6 $\pm$ 0.6*	47.1 $\pm$ 1.0*	66.2 $\pm$ 0.5	58.4 $\pm$ 0.6	60.0 $\pm$ 1.2*	77.1 $\pm$ 0.7*	57.0 $\pm$ 0.9*
InfMax-Norm	54.6 $\pm$ 1.5*	78.1 $\pm$ 0.5	56.9 $\pm$ 0.6	47.1 $\pm$ 1.0*	65.6 $\pm$ 0.5*	58.1 $\pm$ 0.6*	58.8 $\pm$ 1.1*	76.2 $\pm$ 0.7*	55.9 $\pm$ 1.0*
Threshold 30%									
Random	71.5 $\pm$ 1.1	82.1 $\pm$ 0.3	74.1 $\pm$ 0.6	64.0 $\pm$ 0.8	72.4 $\pm$ 0.2	71.7 $\pm$ 0.3	78.0 $\pm$ 0.4	82.4 $\pm$ 0.3	76.0 $\pm$ 0.3
Degree	67.5 $\pm$ 1.2	81.0 $\pm$ 0.4	70.4 $\pm$ 0.6	58.4 $\pm$ 1.0	70.5 $\pm$ 0.3	67.7 $\pm$ 0.4	73.2 $\pm$ 0.8	81.1 $\pm$ 0.4	71.0 $\pm$ 0.4
Pagerank	79.4 $\pm$ 0.5	82.5 $\pm$ 0.3	82.3 $\pm$ 0.3	70.2 $\pm$ 0.3	72.7 $\pm$ 0.2	73.8 $\pm$ 0.2	79.9 $\pm$ 0.3	82.6 $\pm$ 0.2	79.0 $\pm$ 0.2
Betweenness	66.9 $\pm$ 1.3	81.4 $\pm$ 0.3	67.5 $\pm$ 0.5	57.7 $\pm$ 1.0	70.8 $\pm$ 0.3	67.8 $\pm$ 0.5	75.3 $\pm$ 0.5	80.9 $\pm$ 0.4	71.7 $\pm$ 0.4
RWCS	79.2 $\pm$ 0.5	82.5 $\pm$ 0.3	82.3 $\pm$ 0.3	69.9 $\pm$ 0.3	72.7 $\pm$ 0.2	73.7 $\pm$ 0.2	78.2 $\pm$ 0.3	81.7 $\pm$ 0.3	77.8 $\pm$ 0.2
GC-RWCS	61.9 $\pm$ 1.5	80.2 $\pm$ 0.4	63.2 $\pm$ 0.5	50.6 $\pm$ 1.1	67.8 $\pm$ 0.4	62.1 $\pm$ 0.6	71.1 $\pm$ 0.8	79.9 $\pm$ 0.5	68.8 $\pm$ 0.4
InfMax-Unif	58.2 $\pm$ 1.5*	79.9 $\pm$ 0.4	59.6 $\pm$ 0.5*	49.6 $\pm$ 1.0*	67.3 $\pm$ 0.5*	61.2 $\pm$ 0.6*	69.4 $\pm$ 1.0*	80.1 $\pm$ 0.5	65.4 $\pm$ 0.5*
InfMax-Norm	58.0 $\pm$ 1.5*	79.9 $\pm$ 0.5	60.0 $\pm$ 0.5*	49.5 $\pm$ 1.0*	67.6 $\pm$ 0.5	61.6 $\pm$ 0.6*	69.6 $\pm$ 1.0*	79.7 $\pm$ 0.5	65.2 $\pm$ 0.5*

Table 3: Experiments with 10% training data. All other settings are the same as in Table 2.

Method	Cora			Citeseer			Pubmed		
	JKNetMaxpool	GCN	GAT	JKNetMaxpool	GCN	GAT	JKNetMaxpool	GCN	GAT
None	76.8 $\pm$ 0.5	81.1 $\pm$ 0.4	83.3 $\pm$ 0.3	65.1 $\pm$ 0.7	72.5 $\pm$ 0.2	72.6 $\pm$ 0.2	84.1 $\pm$ 0.1	85.4 $\pm$ 0.1	84.7 $\pm$ 0.1
Threshold 10%									
Random	67.9 $\pm$ 0.9	76.9 $\pm$ 0.5	71.6 $\pm$ 0.6	50.6 $\pm$ 1.0	68.6 $\pm$ 0.2	68.6 $\pm$ 0.4	76.1 $\pm$ 0.5	79.7 $\pm$ 0.3	74.2 $\pm$ 0.6
Degree	62.9 $\pm$ 1.2	73.9 $\pm$ 0.5	66.0 $\pm$ 0.7	58.0 $\pm$ 0.7	64.4 $\pm$ 0.4	63.5 $\pm$ 0.7	64.6 $\pm$ 0.9	75.0 $\pm$ 0.6	68.8 $\pm$ 0.8
Pagerank	67.4 $\pm$ 0.8	75.2 $\pm$ 0.5	72.0 $\pm$ 0.5	52.8 $\pm$ 0.9	67.3 $\pm$ 0.2	67.5 $\pm$ 0.3	71.8 $\pm$ 0.6	78.2 $\pm$ 0.3	72.1 $\pm$ 0.6
Betweenness	63.6 $\pm$ 1.1	74.6 $\pm$ 0.5	63.9 $\pm$ 0.8	57.9 $\pm$ 0.9	66.3 $\pm$ 0.3	65.5 $\pm$ 0.6	67.6 $\pm$ 0.9	74.0 $\pm$ 0.7	65.1 $\pm$ 1.0
RWCS	67.4 $\pm$ 0.7	75.3 $\pm$ 0.5	69.0 $\pm$ 0.6	57.6 $\pm$ 0.7	67.2 $\pm$ 0.2	67.4 $\pm$ 0.3	70.7 $\pm$ 0.6	77.1 $\pm$ 0.4	70.0 $\pm$ 0.7
GC-RWCS	58.8 $\pm$ 1.4	73.1 $\pm$ 0.8	58.2 $\pm$ 1.0	45.5 $\pm$ 1.1	61.5 $\pm$ 0.5	60.6 $\pm$ 1.0	63.6 $\pm$ 1.1	71.8 $\pm$ 0.9	61.7 $\pm$ 1.2
InfMax-Unif	58.1 $\pm$ 1.4	72.6 $\pm$ 0.8	57.3 $\pm$ 1.1*	45.2 $\pm$ 1.1	61.6 $\pm$ 0.5	60.8 $\pm$ 1.0	62.1 $\pm$ 1.2*	71.3 $\pm$ 1.0	61.2 $\pm$ 1.2
InfMax-Norm	58.3 $\pm$ 1.5	72.8 $\pm$ 0.8	58.0 $\pm$ 1.1	45.1 $\pm$ 1.1	60.8 $\pm$ 0.5*	60.4 $\pm$ 1.0	60.6 $\pm$ 1.2*	70.1 $\pm$ 0.9*	60.6 $\pm$ 1.2*

B ADDITIONAL EXPERIMENTS

B.1 Results under the Experiment Setup by Ma et al. [19]

In a previous version of this paper, we conducted experiments under a setup exactly the same as that by Ma et al. [19]. The main difference between our current setup (as detailed in Section 5.2) and their setup is that we now use GCN models independent of the victim models to construct the constant perturbation vector, while they directly use the victim models to construct the perturbation vector. Although the use of the model information is extremely

limited in the setup by Ma et al. [19], it is not completely black-box. Nevertheless, in this section, we provide the experiment results under the setup by Ma et al. [19] for easier comparison with existing literature. As can be seen in Table 2, the overall trend is almost the same as Table 1.

B.2 10% Training Data

We also conduct experiments with 10% training data, with other settings following the experiments in Table 2. The results have similar trends as the experiments with 60% training data.