

Adversarial Attacks and Defenses on Graphs: A Review, A Tool and Empirical Studies

Wei Jin[†], Yaxin Li[†], Han Xu[†], Yiqi Wang[†], Shuiwang Ji[‡], Charu Aggarwal[§], Jiliang Tang[†]

[†]Data Science and Engineering Lab, Michigan State University

[‡]Texas A&M University

[§]IBM T. J. Watson Research Center

{jinwei2,liyaxin1,xuhan1,wangy206,tangjili}@msu.edu, sjj@tamu.edu,
charu@us.ibm.com

ABSTRACT

Deep neural networks (DNNs) have achieved significant performance in various tasks. However, recent studies have shown that DNNs can be easily fooled by small perturbation on the input, called adversarial attacks. As the extensions of DNNs to graphs, Graph Neural Networks (GNNs) have been demonstrated to inherit this vulnerability. Adversary can mislead GNNs to give wrong predictions by modifying the graph structure such as manipulating a few edges. This vulnerability has arisen tremendous concerns for adapting GNNs in safety-critical applications and has attracted increasing research attention in recent years. Thus, it is necessary and timely to provide a comprehensive overview of existing graph adversarial attacks and the countermeasures. In this survey, we categorize existing attacks and defenses, and review the corresponding state-of-the-art methods. Furthermore, we have developed a repository with representative algorithms¹. The repository enables us to conduct empirical studies to deepen our understandings on attacks and defenses on graphs.

1. INTRODUCTION

Graphs can be used as the representation of a large number of systems across various areas such as social science (social networks), natural science (physical systems, and protein-protein interaction networks), and knowledge graphs [2; 26]. With their prevalence, it is of great significance to learn effective representations for graphs that can facilitate various downstream tasks such as node classification, graph classification, link prediction and recommendation [2; 26]. Graph Neural Networks (GNNs), which generalize traditional deep neural networks (DNNs) to graphs, pave one effective way to learn representations for graphs [26; 73; 3; 84]. The power of GNNs relies on their capability of capturing the graph structure simultaneously with the node features. Instead of only considering the instances (nodes with their features) independently, GNNs also leverage the relationships between them. Specifically, GNNs follow a message-passing scheme [24], where the nodes aggregate and transform the information from their neighbors in each layer. By stacking multiple GNN layers, information can be propagated further

through the graph structure and we can embed nodes into low-dimensional representations. The obtained node representations can then be fed into any differentiable prediction layer so that the whole model can be trained in an end-to-end fashion. Due to their strong representation learning capability, GNNs have gained practical significance in various applications ranging from data mining [34; 63], natural language processing [48], and computer vision [37] to health-care and biology [45].

As new generalizations of traditional DNNs to graphs, GNNs inherit both advantages and disadvantages of traditional DNNs. Similar to traditional DNNs, GNNs are also powerful in learning representations of graphs and have permeated numerous areas of science and technology. Traditional DNNs are easily fooled by adversarial attacks [25; 75; 40; 39]. In other words, the adversary can insert slight perturbation during either the training or test phases, and the DNN models will totally fail. It is evident that GNNs also inherit this drawback [89; 17; 91]. The attacker can generate graph adversarial perturbations by manipulating the graph structure or node features to fool the GNN models. As illustrated in Figure 1, originally node 7 was classified by the GNN model as a green node; after node 7 creates a new connection with node 3 and modifies its own features, the GNN model misclassifies it as a blue node. Such vulnerability of GNNs has arisen tremendous concerns on applying them in safety-critical applications such as financial system and risk management. For example, in a credit scoring system, fraudsters can fake connections with several high-credit customers to evade the fraud detection models; and spammers can easily create fake followers to increase the chance of fake news being recommended and spread. Therefore, there is an urgent need to investigate graph adversarial attacks and their countermeasures.

Pushing this research has a great potential to facilitate the successful adoption of GNNs in a broader range of fields, which encourages increasing attention on graph adversarial attacks and defenses in recent years. Thus, it is necessary and timely to provide a comprehensive and systematic overview on existing algorithms. Meanwhile, it is of great importance to deepen our understandings on graph adversarial attacks via empirical studies. These understandings can not only provide knowledge about the behaviors of attacks but also offer insights for us to design defense strategies. These motivate this survey with the following key purposes:

¹<https://github.com/DSE-MSU/DeepRobust>

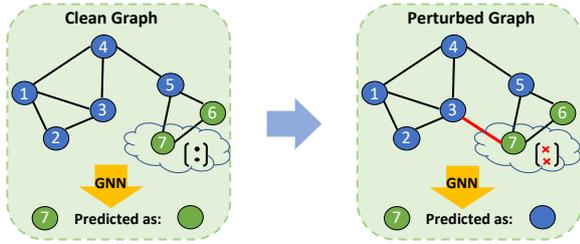


Figure 1: An example of adversarial attack on graph data. The goal of the GNN is to predict the color of the nodes. Here node 7 is the target node. Attacker aims to change the prediction of GNN on node 7 by modifying the edges and features.

- We categorize existing attack methods from various perspectives in Section 3 and review representative algorithms in Section 4.
- We classify existing countermeasures according to their defense strategies and give a review on representative algorithms for each category in Section 5.
- We perform empirical studies based on the repository we developed that provide comprehensive understandings on graph attacks and defenses in Section 6.
- We discuss some promising future research directions in Section 7.

2. PRELIMINARIES AND DEFINITIONS

Before presenting the review and empirical studies, we first introduce concepts, notations and definitions in this section.

2.1 Learning on Graph Data

In this survey, we use $G = (V, E)$ to denote the structure of a graph where $V = \{v_1, \dots, v_N\}$ is the set of N nodes and $E = \{e_1, \dots, e_K\}$ is the edge set. We use matrix $\mathbf{A} \in \{0, 1\}^{N \times N}$ to denote the adjacency matrix of G , where each entry $\mathbf{A}_{ij} = 1$ means nodes v_i and v_j are connected in G . Furthermore, we use $\mathbf{X} \in \mathbb{R}^{N \times D}$ to denote the node attribute matrix where D is the dimension of the node feature vectors. Thus, graph data can be denoted as $G = (\mathbf{A}, \mathbf{X})$. There are a lot of learning tasks on graphs and in this work, we focus on the classification problems on graphs. Furthermore, we use f_θ with parameters θ to denote the learning models in this survey.

Node-Level Classification. For node-level classification, each node in the graph G belongs to a class in the label set Y . The graph model aims to learn a neural network, based on labeled nodes (training nodes), denoted as V_L , to predict the class of unlabeled nodes (test nodes). The training objective function can be formulated as:

$$\min_{\theta} \mathcal{L}_{\text{train}}(f_\theta(G)) = \sum_{v_i \in V_L} \ell(f_\theta(\mathbf{X}, \mathbf{A})_i, y_i), \quad (1)$$

where $f_\theta(\mathbf{X}, \mathbf{A})_i$ and y_i are the predicted and the true label of node v_i and $\ell(\cdot, \cdot)$ is a loss function such as cross entropy.

Graph-Level Classification. For graph-level classification, each individual graph has a class in the label set Y . We use \mathcal{G} to denote a set of graphs, and G_L is the labeled set

(training set) of \mathcal{G} . The goal of graph-level classification is to learn a mapping function $f_\theta : \mathcal{G} \rightarrow Y$ to predict the labels of unlabeled graphs. Similar to node-level classification, the objective function can be formulated as

$$\min_{\theta} \mathcal{L}_{\text{train}}(\mathcal{G}) = \sum_{G_i \in G_L} \ell(f_\theta(G_i), y_i), \quad (2)$$

where G_i is the labeled graph with ground truth y_i and $f_\theta(G_i)$ is the prediction of the graph G_i .

2.2 A General Form of Graph Adversarial Attack

Based on the objectives in Section 2.1, we can define a general form of the objective for adversarial attacks, which aims to misguide the victim model by minimizing some attack loss. Thus, the problem of node-level graph adversarial attacks can be stated as:

Given $G = (\mathbf{A}, \mathbf{X})$ and victim nodes subset $V_t \subseteq V$. Let y_u denote the class for node u (predicted or using ground truth). The goal of the attacker is to find a perturbed graph $\hat{G} = (\hat{\mathbf{A}}, \hat{\mathbf{X}})$ that minimize the following attack objective \mathcal{L}_{atk} ,

$$\begin{aligned} \min \mathcal{L}_{\text{atk}}(f_\theta(\hat{G})) &= \sum_{u \in V_t} \ell_{\text{atk}}(f_{\theta^*}(\hat{G})_u, y_u) \\ \text{s.t.}, \quad \theta^* &= \arg \min_{\theta} \mathcal{L}_{\text{train}}(f_\theta(G')), \end{aligned} \quad (3)$$

where ℓ_{atk} is a loss function for attacking and one option is to choose $\ell_{\text{atk}} = -\ell$ and G' can either be G or \hat{G} . Note that \hat{G} is chosen from a constrained domain $\Phi(G)$. Given a fixed perturbation budget Δ , a typical $\Phi(G)$ can be implemented as,

$$\|\hat{\mathbf{A}} - \mathbf{A}\|_0 + \|\hat{\mathbf{X}} - \mathbf{X}\|_0 \leq \Delta. \quad (4)$$

We omit the definition of graph-level adversarial attacks since (1) the graph-level adversarial attacks can be defined similarly and (2) the majority of the adversarial attacks and defenses focus on node-level. Though adversarial attacks have been extensively studied in the image domain, we still need dedicated efforts for graphs due to unique challenges – (1) graph structure is discrete; (2) the nodes in the graph are not independent; and (3) it is difficult to measure whether the perturbation on the graph is imperceptible or not.

2.3 Notations

With the aforementioned definitions, we list all the notations which will be used in the following sections in Table 1.

3. TAXONOMY OF GRAPH ADVERSARIAL ATTACKS

In this section, we briefly introduce the main taxonomy of adversarial attacks on graph structured data. Attack algorithms can be categorized into different types based on different goals, resources, knowledge and capacity of attackers. We try to give a clear overview on the main components of graph adversarial attacks.

3.1 Attacker's Capacity

The adversarial attacks can happen at two phases, i.e., the model training and model testing. It depends on the attacker's capacity to insert adversarial perturbation:

Table 1: Commonly used notations

Notations	Description	Notations	Description
G	Graph	u	Target node
\hat{G}	Perturbed graph	y_u	Label of node u
V	The set of nodes	f_θ	Neural network model
V_L	The set of labeled nodes	\mathcal{L}	Loss function
E	The set of edges	$l(\cdot, \cdot)$	Pair-wise loss function
\mathbf{A}	Adjacency matrix	$\ \cdot\ _0$	ℓ_0 norm
$\hat{\mathbf{A}}$	Perturbed adjacency matrix	Δ	Perturbation budget
\mathbf{X}	Node attribute matrix	\mathbf{Z}	Predicted probability
$\hat{\mathbf{X}}$	Perturbed node attribute matrix	\mathbf{h}_u	Hidden representation of node u
D	Dimension of node features	e_{ij}	Edge between node v_i and v_j

- **Evasion Attack:** Attacking happens after the GNN model is trained or in the test phase. The model is fixed, and the attacker cannot change the model parameter or structure. The attacker performs evasion attack when $G' = G$ in Eq. (3).
- **Poisoning Attack:** Attacking happens before the GNN model is trained. The attacker can add “poisons” into the model training data, letting trained model have malfunctions. It is the case when $G' = \hat{G}$ in Eq. (3).

3.2 Perturbation Type

The attacker can insert adversarial perturbations from different aspects. The perturbations can be categorized as modifying node features, adding/deleting edges, and adding fake nodes. Attackers should also keep the perturbation unnoticeable, otherwise it would be easily detected.

- **Modifying Feature:** Attackers can slightly change the node features while maintaining the graph structure.
- **Adding or Deleting Edges:** Attackers can add or delete edges for the existing nodes under certain budget of total actions.
- **Injecting Nodes:** Attackers can insert fake nodes to the graph, and link it with some benign nodes in the graph.

3.3 Attacker’s Goal

According to the goals of attacks, we can divide the attacks for node-level classification into the following two categories:

- **Targeted Attack:** There is a small set of test nodes. The attacker aims to let the trained model misclassify these test samples. It is the case when $V_t \subset V$ in Eq. (3). We can further divide targeted attacks into (1) direct attack where the attacker directly modifies the features or edges of the target nodes and (2) influencer attack where the attacker can only manipulate other nodes to influence the targets.
- **Untargeted Attack:** The attacker aims to insert poisons to let the trained model have bad overall performance on all test data. It is the case when $V_t = V$ in Eq. (3).

Not that for graph-level classification, there also exist targeted and untargeted attacks. Targeted attack aims to induce the model to give a specific label to a given graph

sample, while untargeted attack only wants the model to perform incorrectly.

3.4 Attacker’s Knowledge

Attacker’s knowledge means how much information an attacker knows about the model that he aims to attack. Usually, there are three settings:

- **White-box Attack:** All information about the model parameters, training input (e.g, adjacency matrix and attribute matrix) and the labels are given to the attacker.
- **Gray-box Attack:** The attacker only has limited knowledge about the victim model. For example, the attacker cannot access the model parameters but can access the training labels. Then it can utilize the training data to train surrogate models to estimate the information from victim model.
- **Black-box Attack:** The attacker does not have access to the model’s parameters or training labels. It can access the adjacency matrix and attribute matrix, and do black-box query for output scores or labels.

3.5 Victim Models

In this part we are going to summarize the victim models that have been proven to be susceptible to adversarial examples.

Graph Neural Networks. Graph neural networks are powerful tools in learning representation of graphs [56; 73]. One of the most successful GNN variants is Graph Convolutional Networks (GCN) [34]. GCN learns the representation for each node by keeping aggregating and transforming the information from its neighbor nodes. Though GNNs can achieve high performance in various tasks, studies have demonstrated that GNNs including GCN are vulnerable to adversarial attacks [89; 56]. Besides, it is evident from recent works [89; 87; 44] that other graph neural networks including column network (CLN) [52], graph attention network (GAT) [63] and JKNet [77] have the same issue.

Other Graph Learning Algorithms. In addition to graph neural networks, adversary may attack some other important algorithms for graphs such as network embeddings including LINE [59] and Deepwalk [51], graph-based semi-supervised learning (G-SSL) [88], and knowledge graph embedding [8; 42].

4. GRAPH ADVERSARIAL ATTACKS

Table 2: Categorization of representative attack methods

Attack Methods	Attack Knowledge	Targeted or Non-targeted	Evasion or Poisoning	Perturbation Type	Application	Victim Model
PGD, Min-max [76]	White-box	Untargeted	Both	Add/Delete edges	Node Classification	GNN
IG-FGSM [72] IG-JSMA [72]	White-box	Both	Evasion	Add/Delete edges Modify features	Node Classification	GNN
Wang et al. [64]	White-box Gray-box	Targeted	Poisoning	Add/Delete edges	Node Classification	GNN
Nettack [89]	Gray-box	Targeted	Both	Add/Delete edges Modify features	Node Classification	GNN
Metattack [91]	Gray-box	Untargeted	Poisoning	Add/Delete edges	Node Classification	GNN
NIPA [57]	Gray-box	Untargeted	Poisoning	Inject nodes	Node Classification	GNN
RL-S2V [17]	Black-box	Targeted	Evasion	Add/Delete edges	Graph Classification Node Classification	GNN
ReWatt [46]	Black-box	Untargeted	Evasion	Add/Delete edges	Graph Classification	GNN
Liu et al. [43]	White-box Gray-box	Untargeted	Poisoning	Flip label Modify features	Classification Regression	G-SSL
FGA [13]	White-box	Targeted	Both	Add/Delete edges	Node Classification Community Detection	Network Embedding
GF-Attack [9]	Black-box	Targeted	Evasion	Add/Delete edges	Node Classification	Network Embedding
Bojchevski et al. [5]	Black-box	Both	Poisoning	Add/Delete edges	Node Classification Community Detection	Network Embedding
Zhang et al. [81]	White-box	Targeted	Poisoning	Add/Delete facts	Plausibility Prediction	Knowledge Graph Embedding
CD-Attack [38]	Black-box	Targeted	Poisoning	Add/Delete edges	Community Detection	Community Detection Algorithm

In this section, we review representative algorithms for graph adversarial attacks. Following the categorizations in the previous section, we first divide these algorithms into white-box, gray-box and black-box and then for algorithms in each category, we further group them into targeted and untargeted attacks. Note that without specific mention, the following attack methods focus on node-level classification. An overall categorization of representative attack methods is shown in Table 2. In addition, some open source implementations of representative algorithms are listed in Table 9.

4.1 White-box Attacks

In white-box attack setting, the adversary has access to any information about the victim model such as model parameters, training data, labels, and predictions. Although in most of the real world cases we do not have the access to such information, we can still assess the vulnerability of the victim models under the worst situation. Typically, white-box attacks use the gradient information from the victim model to guide the generation of attacks [13; 76; 72; 10].

4.1.1 Targeted Attack

Targeted attack aims to mislead the victim model to make wrong predictions on some target samples. A lot of studies follow the white-box targeted attack setting with a wide range of real-world applications. FGA [13] extracts the link gradient information from GCN, and then greedily selects the pair of nodes with maximum absolute gradient to modify the graph iteratively. Genetic algorithm based Q-Attack is proposed to attack a number of community detection algorithms [10]. Iterative gradient attack (IGA) based on

the gradient information in the trained graph auto-encoder, which is introduced to attack link prediction [11]. Furthermore, the vulnerability of knowledge graph embedding is investigated in [81] and the plausibility of arbitrary facts in knowledge graph can be effectively manipulated by the attacker. Recommender systems based on GNNs are also vulnerable to adversarial attacks, which is shown in [86]. In addition, there are great efforts on attacking node-level classification. Traditional attacks in the image domain always use models’ gradients to find adversarial examples. However, due to the discrete property of graph data, directly calculating gradients of models could fail. To solve this issue, the work [72] suggests to use integrated gradient [58] to better search for adversarial edges and feature perturbations. During the attacking process, the attacker iteratively chooses the edge or feature which has the strongest effect to the adversarial objective. By this way, it can cause the victim model to misclassify target nodes with a higher successful rate. The work [79] assumes there is a set of “bad actor” nodes in a graph. When they flip the edges with any target nodes in a graph, it will cause the GNN model to have a wrong prediction on the target node. These “bad actor” nodes are critical to the safety of GNN models. For example, Wikipedia has hoax articles which have few and random connections to real articles. Manipulating the connections of these hoax articles will cause the system to make wrong prediction of the categories of real articles.

4.1.2 Untargeted Attack

Currently there are not many studies on untargeted white-box attack, and topology attack [76] is one representative algorithm. It first constructs a binary symmetric perturba-

tion matrix $\mathbf{S} \in \{0, 1\}^n$ where $\mathbf{S}_{ij} = 1$ indicates to flip the edge between i and j and $\mathbf{S}_{ij} = 0$ means no modification on \mathbf{A}_{ij} . Thus, the goal of the attacker is to find \mathbf{S} that minimizes the predefined attack loss given a finite budget of edge perturbations Δ , i.e., $\|\mathbf{S}\|_0 \leq \Delta$. It considers two different attack scenarios: attacking pre-trained GNN with fixed parameters θ and attacking a re-trainable GNN f_θ . For attacking a fixed f_θ , the problem can be formulated as,

$$\min_{\mathbf{S} \in \{0,1\}^n} \mathcal{L}_{\text{atk}}(f_\theta(\mathbf{S}, \mathbf{A}, \mathbf{X})) \quad \text{s.t. } \|\mathbf{S}\|_0 \leq \Delta. \quad (5)$$

It utilizes the Projected Gradient Descent (PGD) algorithm in [47] to search the optimal \mathbf{S} . Note that the work [47] is also one popular attack algorithm in the image domain. For the re-trainable GNNs, parameter θ will be retrained after adversarial manipulation. Thus the attack problem is formulated as a min-max form where the inner maximization is to update θ by maximizing the attack loss and can be solved by gradient ascent while the outer minimization can be solved by PGD. Another work on untargeted white-box attack aims to attack multi-network mining by perturbing the networks (graphs) [85]. Specifically, in each iteration it measures the influence of network elements (edges and node attributes) to the mining results and attack one network element with the highest influence value.

4.2 Gray-box Attacks

White-box attacks assume that attackers can calculate gradient through model parameters, which is not always practical in real-world scenarios. Gray-box attacks are proposed to generate attacks with limited knowledge on the victim model [89; 91; 57]. Usually they first train a surrogate model with the labeled training data to approximate the information of the victim model and then generate perturbations to attack the surrogate model. It is noted that these models need the access to the labels of training data, thus they are not black-box attacks that will be introduced in the following subsection.

4.2.1 Targeted Attack

The early work on targeted gray-box attacks is for graph clustering [15]. It demonstrates that injecting noise to a domain name system (DNS) query graph can degrade the performance of graph embedding models. Different from [15], the work [89] proposes an attack method called Nettack to generate structure and feature attacks, aiming at solving Eq. (3). Besides, they argue that only limiting the perturbation budgets cannot always make the perturbation “unnoticeable”. They suggest the perturbed graphs should also maintain important graph properties, including degree distribution and feature co-occurrence. Therefore, Nettack first selects possible perturbation candidates not violating degree distribution and feature co-occurrence of the original graph. Then it greedily chooses the perturbation that has the largest score to modify the graph, where the score is defined as,

$$\max_{i \neq y} \ln(\mathbf{Z}_{u,i}(G')) - \ln(\mathbf{Z}_{u,y}(G')), \quad (6)$$

where $\mathbf{Z}_{u,i}$ is the probability of node u to be the class i predicted by the surrogate model. Thus, the goal of the attacker is to maximize the difference in the log-probabilities of the target node u . By doing this repeatedly until reaching the perturbation budget Δ , it can get the final modified

graph. Furthermore, it suggests that such graph attack can also transfer from model to model, just as the attacks in the image domain [25]. The authors also conduct influencer attacks where they can only manipulate the nodes except the target. It turns out that influencer attacks lead to a lower decrease in performance compared with directly modifying target node given the same perturbation budget.

However, Nettack cannot handle large-scale graphs due to its high time complexity. The work [90] employs statistical analysis to investigate the patterns exhibited by Nettack perturbations. According to the observations on perturbation patterns, a scalable attack method Fastack is proposed. It first ranks the possible perturbations by their impact on the victim model and chooses the top perturbations to generate attacks. In [67], AFGSM is proposed to derive an approximate closed-form solution with a lower time cost for the attack model and sequentially inject fake nodes into the graph.

Another line of works focus on generating backdoor attacks for graphs [83; 74]. Backdoor attack aims to find some “trigger” (or hidden pattern) that would produce misclassified results when added to an input [66; 14]. To generate backdoor attacks for graphs, the attacker injects triggers (e.g. carefully-crafted subgraphs) in the training samples, and aims to fool GNNs into misclassifying those samples with triggers while maintaining the performance on other testing samples. Note that graph backdoor attacks can be applied in both node-level and graph-level classification [74].

4.2.2 Untargeted Attack

Although following the same way of training a surrogate model as Nettack, Metattack [91] is a kind of untargeted poisoning attack. It tackles the bi-level problem in Eq. (3) by using meta-gradient. Basically, it treats the graph structure matrix as a hyper-parameter and the gradient of the attacker loss with respect to it can be obtained by:

$$\nabla_G^{\text{meta}} = \nabla_G \mathcal{L}_{\text{atk}}(f_{\theta^*}(G)). \quad (7)$$

Note that $\nabla_G \mathcal{L}_{\text{atk}}(f_{\theta^*}(G))$ is actually a function with respect to both G and θ . If θ^* is obtained by some differential operations, we can compute ∇_G^{meta} as follows,

$$\nabla_G^{\text{meta}} = \nabla_f \mathcal{L}_{\text{atk}}(f_{\theta^*}(G)) \cdot [\nabla_G f_{\theta^*}(G) + \nabla_{\theta^*} f_{\theta^*}(G) \cdot \nabla_G \theta^*] \quad (8)$$

where θ^* is often obtained by gradient descent in fixed iterations T . At iteration $t+1$, the gradient of θ_{t+1} with respect to G can be formulated as,

$$\nabla_G \theta_{t+1} = \nabla_G \theta_t - \alpha \nabla_G \nabla_{\theta_t} \mathcal{L}_{\text{train}}(f_{\theta_t}(G)), \quad (9)$$

where α denotes learning rate of the gradient descent operation. By unrolling the training procedure from θ_T back to θ_0 , we can get $\nabla_G \theta_T$ and then ∇_G^{meta} . A greedy approach is applied to select the perturbation based on the meta gradient. Specifically, given the meta-gradient for a node pair (u, v) , it defines a score $S(u, v) = \nabla_{\mathbf{A}_{uv}}^{\text{meta}} \cdot (-2 \cdot \mathbf{A}_{uv} + 1)$ and greedily picks the perturbation which has the highest score but satisfies the unnoticeable constraints as in Nettack.

Instead of modifying the connectivity of existing nodes, a novel reinforcement learning method for node injection poisoning attacks (NIPA) [57] is proposed to inject fake nodes into graph data. Specifically, NIPA first injects singleton n nodes into the original graph. Then in each action a_t , the attacker first chooses an injected node to connect with

another node in the graph and then assigns a label to the injected node. By doing this sequentially, the final graph is statistically similar to the original graph but can degrade the overall model performance.

4.3 Black-box Attacks

Different from gray-box attacks, black-box attacks [17; 46; 57; 5; 9; 44] are more challenging since the attacker can only access the adjacency matrix, attribute matrix and output of the victim model. The access of parameters, labels and predicted probability is prohibited.

4.3.1 Targeted Attack

As mentioned earlier, training a surrogate model requires access to the labels of training data, which is not always practical. We hope to find a way that we only need to do black-box query on the victim model [17] or attack the victim in an unsupervised fashion [5; 9].

To do black-box query on the victim model, reinforcement learning is introduced. RL-S2V [17] is the first work to employ reinforcement learning technique to generate adversarial attacks on graph data under the black-box setting. They model the attack procedure as a Markov Decision Process (MDP) and the attacker is allowed to modify m edges to change the predicted label of the target node u . They study both node-level (targeted) and graph-level (untargeted) attacks. For node-level attack, they define the MDP as follows,

- **State.** The state s_t is represented by the tuple $(G^{(t)}, u)$ where $G^{(t)}$ is the modified graph at time step t .
- **Action.** A single action at time step t is denoted as a_t . For each action a_t , the attacker can choose to add or remove an edge from the graph. Furthermore, a hierarchical structure is applied to decompose the action space.
- **Reward.** Since the goal of the attacker is to change the classification result of the target node u , RL-S2V gives non-zero reward r to the attacker at the end of the MDP:

$$r(s_m, a_m) = \begin{cases} 1 & \text{if } f_\theta(G^{(m)}, u) \neq y, \\ -1 & \text{if } f_\theta(G^{(m)}, u) = y. \end{cases} \quad (10)$$

In the intermediate steps, the attacker receives no reward, i.e., $\forall t = 1, 2, \dots, m-1, r(s_t, a_t) = 0$.

- **Termination.** The process terminates when the attacker finishes modifying m edges.

Since they define the MDP of graph-level attack in the similar way, we omit the details. Further, the Q-learning algorithm [50] is adopted to solve the MDP and guide the attacker to modify the graph.

Instead of attacking node-level classification, the work [5] shows a way to attack the family of node embedding models in the black-box setting. Inspired by the observation that DeepWalk can be formulated in matrix factorization form [53], they maximize the unsupervised DeepWalk loss with matrix perturbation theory by performing Δ edge flips. It is further demonstrated that the perturbed structure is transferable to other models like GCN and Label Propagation. However, this method only considers the structure information. GF-Attack [9] is proposed to incorporate

the feature information into the attack model. Specifically, they formulate the connection between the graph embedding method and general graph signal process with graph filter and construct the attacker based on the graph filter and attribute matrix. GF-Attack can also be transferred to other network embedding models and achieves better performance than the method in [5].

4.3.2 Untargeted Attack

It is argued that the perturbation constraining only the number of modified edges may not be unnoticeable enough. A novel framework ReWatt [46] is proposed to solve this problem and perform untargeted graph-level attack. Employing a reinforcement learning framework, ReWatt adopts the rewiring operation instead of simply adding/deleting an edge in one single modification to make perturbation more unnoticeable. One rewiring operation involves three nodes v_1, v_2 and v_3 , where ReWatt removes the existing edge between v_1 and v_2 and connects v_1 and v_3 . ReWatt also constrains v_3 to be the 2-hop neighbor of v_1 to make perturbation smaller. Such rewiring operation does not change the number of nodes and edges in the graph and it is further proved that such rewiring operation affects algebraic connectivity and effective graph resistance, both of which are important graph properties based on graph Laplacian, in a smaller way than adding/deleting edges.

To perform graph adversarial attacks in a more realistic way, the work [44] proposes a more restricted black-box setting where the attackers only have access to a subset of nodes and can only attack a small number of them. Under this setting, the attacker is required to generate attacks in two steps: 1) selecting a small subset of nodes to attack under the limits of node access; 2) modifying node features or edges under the perturbation budget. By exploiting the structural inductive biases of the GNN models [77; 36] as the information source for attacking, it proposes a practical greedy method of adversarial attacks for node-level classification tasks and effectively degrades the performance of GNNs.

5. COUNTERMEASURES AGAINST GRAPH ADVERSARIAL ATTACKS

In previous sections, we have shown that graph neural networks can be easily fooled by unnoticeable perturbation on graph data. The vulnerability of graph neural networks poses great challenges to apply them in safety-critical applications. In order to defend the graph neural networks against these attacks, different countermeasure strategies have been proposed. The existing methods can be categorized into the following types: (1) adversarial training, (2) adversarial perturbation detection, (3) certifiable robustness, (4) graph purification, and (5) attention mechanism.

5.1 Adversarial Training

Adversarial training is a widely used countermeasure for adversarial attacks in image data [25]. The main idea of adversarial training is to inject adversarial examples into the training set such that the trained model can correctly classify the future adversarial examples. Similarly, we can also adopt this strategy to defend graph adversarial attacks as follows,

$$\min_{\theta} \max_{\substack{\phi_{\mathbf{A}} \in \mathcal{P}_{\mathbf{A}} \\ \phi_{\mathbf{X}} \in \mathcal{P}_{\mathbf{X}}}} \mathcal{L}_{\text{train}}(f_{\theta}(\mathbf{A} + \delta_{\mathbf{A}}, \mathbf{X} + \delta_{\mathbf{X}})), \quad (11)$$

where $\delta_{\mathbf{A}}$, $\delta_{\mathbf{X}}$ denote the perturbation on \mathbf{A} , \mathbf{X} , respectively; $\mathcal{P}_{\mathbf{A}}$ and $\mathcal{P}_{\mathbf{X}}$ stand for the domains of imperceptible perturbation. The min-max optimization problem in Eq (11) indicates that adversarial training involves two processes: (1) generating perturbations that maximize the prediction loss and (2) updating model parameters that minimize the prediction loss. By alternating the above two processes iteratively, we can train a robust model against to adversarial attacks. Since there are two inputs, i.e., adjacency matrix \mathbf{A} and attribute matrix \mathbf{X} , adversarial training can be done on them separately. To generate perturbations on the adjacency matrix, it is proposed to randomly drop edges during adversarial training [17]. Though such simple strategy cannot lead to very significant improvement in classification accuracy (1% increase), it shows some effectiveness with such cheap adversarial training. Furthermore, projected gradient descent is used to generate perturbations on the discrete input structure, instead of randomly dropping edges [76]. On the other hand, an adversarial training strategy with dynamic regularization is proposed to perturb the input features [22]. Specifically, it includes the divergence between the prediction of the target example and its connected examples into the objective of adversarial training, aiming to attack and reconstruct graph smoothness. Furthermore, batch virtual adversarial training [19] is proposed to promote the smoothness of GNNs and make GNNs more robust against adversarial perturbations. Several other variants of adversarial training on the input layer are introduced in [12; 18; 69; 20].

The aforementioned adversarial training strategies face two main shortcomings: (1) they generate perturbations on \mathbf{A} and \mathbf{X} separately; and (2) it is not easy to perturb the graph structure due to its discreteness. To overcome the shortcomings, instead of generating perturbation on the input, a latent adversarial training method injects perturbations on the first hidden layer [32]:

$$\min_{\theta} \max_{\delta \in \mathcal{P}} \mathcal{L}_{\text{train}} \left(f_{\theta}(G; \mathbf{H}^{(1)} + \delta) \right), \quad (12)$$

where $\mathbf{H}^{(1)}$ denotes the representation matrix of the first hidden layer and $\delta \in \mathcal{P}$ is some perturbation on \mathbf{H} . It is noted that the hidden representation is continuous and it incorporates the information from both graph structure and node attributes.

5.2 Adversarial Perturbation Detection

To resist graph adversarial attacks during the test phase, there is one main strategy called adversary detection. These detection models protect the GNN models by exploring the intrinsic difference between adversarial edges/nodes and the clean edges/nodes [78; 28]. The work [78] is the first work to propose detection approaches to find adversarial examples on graph data. It introduces four methods to distinguish adversarial edges or nodes from the clean ones including (1) link prediction (2) sub-graph link prediction (3) graph generation models and (4) outlier detection. These methods have shown some help to correctly detect adversarial perturbations. The work [28] introduces a method to randomly draw subsets of nodes, and relies on graph-aware criteria to judiciously filter out contaminated nodes and edges before employing a semi-supervised learning (SSL) module. The proposed model can be used to detect different anomaly generation models, as well as adversarial attacks.

5.3 Certifiable Robustness

Previously introduced adversarial training strategies are heuristic and only show experimental benefits. However, we still do not know whether there exist adversarial examples even when current attacks fail. Therefore, there are works [92; 6; 30; 7; 93; 61; 65] considering to seriously reason the safety of graph neural networks which try to certify the GNN’s robustness. As we know, GNN’s prediction on one node v_t always depends on its neighbor nodes. In [92], they ask the question: which nodes in a graph are safe under the risk of any admissible perturbations of its neighboring nodes’ attributes. To answer this question, for each node v and its corresponding label y_v , they try to exactly calculate an upper bound $U(v)$ of the maximized margin loss :

$$U(v) \geq \max_{G' \in \mathcal{G}} \left(\max_{i \neq y} \mathbf{Z}_{v,i}(G') - \mathbf{Z}_{v,y}(G') \right), \quad (13)$$

where \mathcal{G} denotes the set of all allowed graph perturbations (in [92] only attribute perturbation). This upper bound U is called the *Certificate* of node v . During the certification, for v , if $U(v) \leq 0$, any perturbation cannot result the model to give a larger score to a wrong class ($i \neq y_v$) than the class y_v , so there is no adversarial attack in \mathcal{G} that can change the model’s prediction. During the test phase, they calculate the certificate for all test nodes, thus they can know how many nodes in a graph is absolutely safe under attributes perturbation. Moreover, this certificate is trainable, directly minimizing the certificates will help more nodes become safe. However, the work [92] only considers the perturbations on node attributes. Analyzing certifiable robustness from a different perspective, in [6], it deals with the case when the attacker only manipulates the graph structure. It derives the robustness certificates (similar to Eq. (13)) as a linear function of personalized PageRank [29], which makes the optimization tractable. In [93], it also tries to certify robustness of graph neural networks under graph structural perturbations. It successfully solves the certification problem using a jointly constrained bilinear programming method. In [65], it borrows the idea from *randomized smoothing* [16] to achieve certifiable robustness for graphs under structural perturbation. Meanwhile, the sparsity of graph data during certification is considered in [7]. It improves the efficiency and accuracy of the certifications against attacks on both graph features and structures. Immunization methods are introduced to improve graphs’ certifiable robustness [61]. Besides, there are also works studying certifiable robustness on GNN’s other applications such as community detection [30].

5.4 Graph Purification

Both adversarial training or certifiable defense methods only target on resisting evasion attacks, which means that the attack happens during the test time. However, graph purification defense methods mainly focus on defending poisoning attacks. Since the poisoning attacks insert poisons into the training graph, purification methods aim to purify the poisoned graph and learn robust graph neural network models based on it. There are two approaches to realize graph purification: pre-processing [72; 21] and graph learning [33; 35].

5.4.1 Pre-processing

Pre-processing methods first purify the perturbed graph data and then train the GNN model on the purified graph. In

this way, the GNN model is trained on a clean graph. The work [72] proposes a purification method based on two empirical observations of the attack methods: (1) Attackers usually prefer adding edges over removing edges or modifying features and (2) Attackers tend to connect dissimilar nodes. As a result, they propose a defense method by eliminating the edges whose two end nodes have small Jaccard Similarity [54]. Because these two nodes are different and it is not likely they are connected in reality, the edge between them may be adversarial. The experimental results demonstrate the effectiveness and efficiency of the proposed defense method. However, this method can only work when the node features are available. In [21], it is observed that Nettack [89] generates the perturbations which mainly changes the small singular values of the graph adjacency matrix. Thus it proposes to purify the perturbed adjacency matrix by using truncated SVD to get its low-rank approximation. it further shows that only keeping the top 10 singular values of the adjacency matrix is able to defend Nettack and improve the performance of GNNs.

5.4.2 Graph Learning

Pre-processing might not be an ideal choice for purifying the graph, as it is independent of GNN’s training process and could mistakenly remove normal edges. An alternative purification strategy is graph learning, which targets at removing adversarial patterns and obtaining clean graph structure. Traditional graph learning methods [23; 31] do not directly deal with adversarial attacks. Hence, to make GNNs more robust, it is of great importance to leverage the characteristics of adversarial attacks to guide the graph learning process. In [33], it has demonstrated that adversarial attacks can essentially violate some important graph properties, i.e., low-rank, sparsity and feature smoothness. Then it explores these graph properties to design robust graph neural networks, Pro-GNN, which jointly learns clean graph structure and robust GNN parameters. Specifically, Pro-GNN alternatively updates graph structure by preserving the aforementioned properties and trains GNN parameters on the updated graph structure. In [80], variational graph autoencoder [35] is employed to reconstruct graph structure, which can also reduce the effects of adversarial perturbations.

5.5 Attention Mechanism

Different from the purification methods which try to exclude adversarial perturbations, attention-based defense methods aim to train a robust GNN model by penalizing model’s weights on adversarial edges or nodes [87; 60; 82]. Basically, these methods learn an attention mechanism to distinguish adversarial edges and nodes from the clean ones, and then make the adversarial perturbations contribute less to the aggregation process of the GNN training. The work [87] assumes that adversarial nodes may have high prediction uncertainty, since adversary tends to connect the node with nodes from other communities. In order to penalize the influence from these uncertain nodes, they propose a defense method named RGCN to model the l -th layer hidden representation $\mathbf{h}_i^{(l)}$ of nodes as Gaussian distribution with mean value $\boldsymbol{\mu}_i^{(l)}$ and variance $\boldsymbol{\sigma}_i^{(l)}$,

$$\mathbf{h}_i^{(l)} \sim N\left(\boldsymbol{\mu}_i^{(l)}, \text{diag}\left(\boldsymbol{\sigma}_i^{(l)}\right)\right), \quad (14)$$

where the uncertainty can be reflected in the variance $\boldsymbol{\sigma}_i^{(l)}$.

When aggregating the information from neighbor nodes, it applies an attention mechanism to penalize the nodes with high variance,

$$\boldsymbol{\alpha}_i^{(l)} = \exp\left(-\gamma\boldsymbol{\sigma}_i^{(l)}\right), \quad (15)$$

where $\boldsymbol{\alpha}_i^{(l)}$ is the attention score assigned to node i and γ is a hyper-parameter. Furthermore, it is verified that the attacked nodes do have higher variances than normal nodes and the proposed attention mechanism does help mitigate the impact brought by adversarial attacks. From a different perspective, GNNGuard is proposed to employ the theory of network homophily [49] to assign higher scores to edges connecting similar nodes while pruning edges between unrelated node [82].

In [60], it suggests that to improve the robustness of one target GNN model, it is beneficial to include the information from other clean graphs, which share the similar topological distributions and node attributes with the target graph. For example, Facebook and Twitter have social network graph data that share similar domains; Yelp and Foursquare have similar co-review graph data. Thus, it first generates adversarial edges E_P on the clean graphs, which serve as the supervision of known perturbation. With this supervision knowledge, it further designs the following loss function to reduce the attention scores of adversarial edges:

$$\mathcal{L}_{dist} = -\min\left(\eta, \mathbb{E}_{e_{ij} \in E \setminus E_P} \boldsymbol{\alpha}_{ij}^{(l)} - \mathbb{E}_{e_{ij} \in E_P} \boldsymbol{\alpha}_{ij}^{(l)}\right), \quad (16)$$

where \mathbb{E} denotes the expectation, $E \setminus E_P$ represents normal edges in the graph, $\boldsymbol{\alpha}_{ij}^{(l)}$ is the attention score assigned to edge e_{ij} and η is a hyper-parameter controlling the margin between the expectation of two distributions. It then adopts meta-optimization to train a model initialization and fine-tunes it on the target poisoned graph to get a robust GNN model.

6. A REPOSITORY FOR GRAPH ATTACKS AND DEFENSES

In this section, we give a brief introduction to the repository we have developed for adversarial attacks and defenses, DeepRobust² [41]. To facilitate the research on adversarial attacks and defenses, DeepRobust includes the majority of representative attack and defense algorithms for both graph data and image data. The repository can enable researchers to deepen our understandings on attacks and defenses via empirical studies. Specifically, for graph adversarial learning, DeepRobust provides 7 datasets: four citation graphs including Cora [55], Cora-ML [4], Citeseer [4] and Pubmed [55], one co-author graph ACM [68], one blog graph Polblogs [1] and one social graph BlogCatalog [27]. On top of that, DeepRobust covers the following attacks and defenses: (1) 5 targeted attack algorithms, i.e., FGA [13], Nettack [89], RL-S2V [17], integrated gradient attack [72] and random attack [89]; (2) 3 untargeted algorithms, i.e., Metattack [91], Topology attack [76] and DICE [70]; (3) one victim model, i.e. GCN [34]; and (4) 5 defense algorithms, i.e., adversarial training [25], GCN-Jaccard [72], GCN-SVD [21], RGCN [87] and Pro-GNN [33].

DeepRobust is an easy-to-use platform for researchers who are working on adversarial attacks and defenses. With Deep-

²<https://github.com/DSE-MSU/DeepRobust>

Robust, users can generate graph adversarial attacks by training an attack model through the attacking APIs or loading the pre-attacked graphs provided by the repository. Robust models can be trained on the perturbed graph with the defense APIs. Once we obtain the perturbed graph and the trained defense model, they can be fed into the evaluation API to compete against each other, and the model performance under the corresponding attack will be reported. In the next section, we use DeepRobust for empirical studies on graph adversarial attacks.

7. EMPIRICAL STUDIES

With the help of DeepRobust, we are able to conduct empirical studies on graph adversarial attacks and discover their important patterns. Next we first introduce the experimental settings and then present the empirical results and findings.

7.1 Experimental Setup

Different attack and defense methods have been designed under different settings. We perform the experiments with one of the most popular settings – the untargeted poisoning setting. Correspondingly we choose representative attack and defense methods that have been designed for this setting. Three representative attack methods are adopted to generate perturbations including DICE [70], Metattack [91] and Topology attack [76]. It is noted that DICE is a white-box attack which randomly connects nodes with different labels or drops edges between nodes sharing the same label. To evaluate the performance of different defense methods under adversarial attacks, we compare the robustness of the natural trained GCN [34] and four defense methods on those attacked graphs, i.e., GCN [34], GCN-Jaccard [72], GCN-SVD [21], RGCN [87] and GAT [62]. Following [91], we use three datasets: Cora, Citeseer and Polblogs. For each dataset, we randomly choose 10% of nodes for training, 10% of nodes for validation and the remaining 80% for test. We repeat each experiment for 5 times and report the average performance. On Cora and Citeseer datasets, the most destructive variant CE-min-max [76] is adopted to implement Topology attack. But CE-min-max cannot converge on Polblogs dataset, we adopt another variant called CE-PGD [76] on this dataset.

7.2 Analysis on Attacked Graph

One way to understand the behaviors of attacking methods is to compare the properties of the clean graph and the attacked graph. In this subsection, we perform this analysis from both global and local perspectives.

Global Measure. We have collected five global properties from both clean graphs and perturbed graphs generated by the three attacks on the three datasets. These properties include the number of added edges, the number of deleted edges, the number of edges, the rank of the adjacent matrix, and clustering coefficient. We only show the results of Metattack in Table 3. Results for Topology attack and DICE can be found in Appendix A.1. Note that we vary the perturbation rates from 0 to 25% with a step of 5% and 0% perturbation denotes the original clean graph. It can be observed from the table:

- Attackers favor adding edges over deleting edges.

- Attacks are likely to increase the rank of the adjacency matrix.
- Attacks are likely to reduce the connectivity of a graph. The clustering coefficients of a perturbed graph decrease with the increase of the perturbation rate.

Table 3: Properties of attacked graphs under Metattack. Note that r denotes perturbation rate and 0% perturbation indicates the original clean graph.

Dataset	r (%)	edge+	edge-	edges	ranks	clustering coefficients
Cora	0	0	0	5069	2192	0.2376
	5	226	27	5268	2263	0.2228
	10	408	98	5380	2278	0.2132
	15	604	156	5518	2300	0.2071
	20	788	245	5633	2305	0.1983
	25	981	287	5763	2321	0.1943
Citeseer	0	0	0	3668	1778	0.1711
	5	181	2	3847	1850	0.1616
	1	341	25	3985	1874	0.1565
	15	485	65	4089	1890	0.1523
	20	614	119	4164	1902	0.1483
	25	743	174	4236	1888	0.1467
Polblogs	0	0	0	16714	1060	0.3203
	5	732	103	17343	1133	0.2719
	10	1347	324	17737	1170	0.2825
	15	1915	592	18038	1193	0.2851
	20	2304	1038	17980	1193	0.2877
	25	2500	1678	17536	1197	0.2723

Local Measure. We have also studied two local properties including the feature similarity and label equality between two nodes connected by three kinds of edges: the newly added edges, the deleted edges and the normal edges which have not been changed by the attack methods. Since features are binary in our datasets, we use jaccard similarity as the measure for feature similarity. For label equality, we report the ratio if two nodes share the same label or have different labels. The feature similarity and label equality results are demonstrated in Figures 2 and 3, respectively. We show the results for Metattack with 5% perturbations. Results for Topology attack and DICE can be found in Appendix A.2. Note that we do not have feature similarity results on Polblogs since this dataset does not have node features. We can make the following observations from the figures.

- Attackers tend to connect nodes with dissimilar features and different labels. As shown in Figure 2 and Figure 3, most of the added edges connect nodes with very dissimilar features and different labels.
- Attackers tend to remove edges from nodes which share similar features and same label. As shown in Figure 2 and Figure 3, most of the deleted edges originally connect nodes sharing similar features and same labels.

7.3 Attack and Defense Performance

In this subsection, we study how the attack methods perform and whether the defense methods can help resist to attacks. Similarly, we vary the perturbation rates from 0 to 25% with a step of 5%. The results are demonstrated in Table 4. We

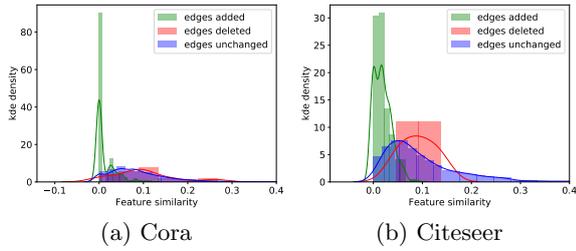


Figure 2: Node feature similarity for Metattack.

show the performance for Metattack. Results for Topology attack and DICE are shown in Appendix A.3. Note that we do not have the performance for Jaccard defense model in Polblogs since this mode requires node features and Polblogs does not provide node features. According to the results, we have the following observations:

- With the increase of the perturbations, the performance of GCN dramatically decreases. This result suggests that Metattack can lead to a significant reduce of accuracy on the GCN model.
- When the perturbations are small, we observe small performance reduction for defense methods which suggests their effectiveness. However, when the graphs are heavily poisoned, their performance also reduces significantly which indicates that efforts are needed to defend heavily poisoning attacks.

Table 4: Performance (accuracy) under Metattack

Dataset	r (%)	0	5	10	15	20	25
Cora	GCN	83.10	76.69	65.58	54.88	48.66	38.44
	Jaccard ¹	82.39	81.02	77.28	72.74	69.16	64.56
	SVD ²	77.97	75.67	70.51	64.34	55.89	45.92
	RGCN	84.81	81.32	72.12	60.25	49.75	37.76
	GAT	81.69	74.75	61.69	52.56	45.30	38.52
Citeseer	GCN	74.53	72.59	63.96	61.66	50.58	44.32
	Jaccard ¹	74.82	73.60	73.50	72.80	72.97	72.53
	SVD ²	70.32	71.30	67.58	63.86	56.91	45.28
	RGCN	74.41	72.68	71.15	69.38	67.93	67.24
	GAT	74.23	72.01	67.12	57.70	47.97	38.70
Polblogs	GCN	95.80	73.93	72.07	67.69	62.29	52.97
	SVD ²	94.99	82.64	71.27	66.09	61.37	52.82
	RGCN	95.60	72.01	67.12	57.70	47.97	38.70
	GAT	95.40	84.83	77.03	69.94	53.62	53.76

¹ Jaccard: GCN-Jaccard defense model.

² SVD: GCN-SVD defense model.

8. CONCLUSION & FUTURE DIRECTIONS

In this survey, we give a comprehensive overview of an emerging research field, adversarial attacks and defenses on graph data. We investigate the taxonomy of graph adversarial attacks, and review representative adversarial attacks and the corresponding countermeasures. Furthermore, we conduct empirical study to show how different defense methods behave under different attacks, as well as the changes in important graph properties by the attacks. Via this comprehensive study, we have gained deep understandings on

this area that enables us to discuss some promising research directions.

- **Imperceptible perturbation measure.** Different from image data, humans cannot easily tell whether a perturbation on graph is imperceptible or not. The ℓ_0 norm constraint on perturbation is definitely not enough. Currently only very few existing work study this problem, thus finding concise perturbation evaluation measure is of great urgency.
- **Different graph data.** Existing works mainly focus on static graphs with node attributes. Complex graphs such as graphs with edge attributes and dynamic graphs are not well-studied yet.
- **Existence and transferability of graph adversarial examples.** There are only a few works discussing about the existence and transferability of graph adversarial examples. Studying this topic is important for us to understand our graph learning algorithm, thus helping us build robust models.
- **Scalability.** The high complexity of attack methods has hindered their use in practical applications. However, there are only few works developing efficient attack methods in terms of time complexity. Furthermore, given that most of attack algorithms are gradient-based, how to reduce their memory complexity also remains a challenge.

Acknowledgments

This research is supported by the National Science Foundation (NSF) under grant number CNS1815636, IIS1928278, IIS1714741, IIS1845081, IIS1907704, and IIS1955285.

9. REFERENCES

- [1] L. A. Adamic and N. Glance. The political blogosphere and the 2004 us election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*, pages 36–43, 2005.
- [2] C. C. Aggarwal, H. Wang, et al. *Managing and mining graph data*, volume 40. Springer, 2010.
- [3] P. W. Battaglia, J. B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- [4] A. Bojchevski and S. Günnemann. Deep gaussian embedding of graphs: Unsupervised inductive learning via ranking. *arXiv preprint arXiv:1707.03815*, 2017.
- [5] A. Bojchevski and S. Günnemann. Adversarial attacks on node embeddings via graph poisoning. *arXiv preprint arXiv:1809.01093*, 2018.
- [6] A. Bojchevski and S. Günnemann. Certifiable robustness to graph perturbations. In *Advances in Neural Information Processing Systems*, pages 8317–8328, 2019.

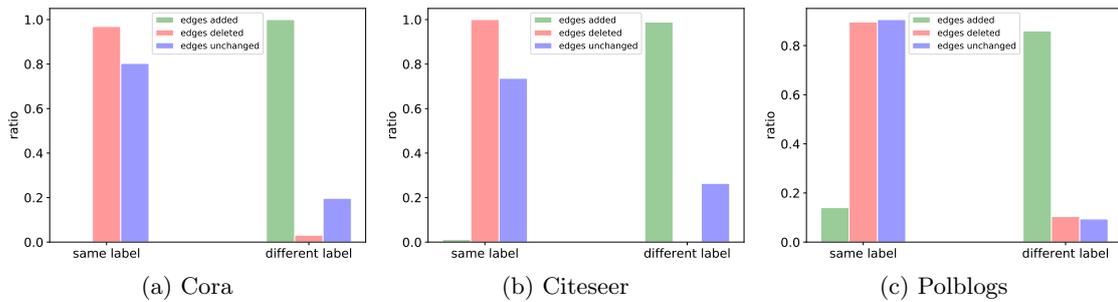


Figure 3: Label equality for Metattack.

- [7] A. Bojchevski, J. Klicpera, and S. Günnemann. Efficient robustness certificates for discrete data: Sparsity-aware randomized smoothing for graphs, images and more. *arXiv preprint arXiv:2008.12952*, 2020.
- [8] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko. Translating embeddings for modeling multi-relational data. In *Advances in neural information processing systems*, pages 2787–2795, 2013.
- [9] H. Chang, Y. Rong, T. Xu, W. Huang, H. Zhang, P. Cui, W. Zhu, and J. Huang. The general black-box attack method for graph neural networks. *arXiv preprint arXiv:1908.01297*, 2019.
- [10] J. Chen, L. Chen, Y. Chen, M. Zhao, S. Yu, Q. Xuan, and X. Yang. Ga-based q-attack on community detection. *IEEE Transactions on Computational Social Systems*, 6(3):491–503, 2019.
- [11] J. Chen, Z. Shi, Y. Wu, X. Xu, and H. Zheng. Link prediction adversarial attack. *arXiv preprint arXiv:1810.01110*, 2018.
- [12] J. Chen, Y. Wu, X. Lin, and Q. Xuan. Can adversarial network attack be defended? *arXiv preprint arXiv:1903.05994*, 2019.
- [13] J. Chen, Y. Wu, X. Xu, Y. Chen, H. Zheng, and Q. Xuan. Fast gradient attack on network embedding. *arXiv preprint arXiv:1809.02797*, 2018.
- [14] X. Chen, C. Liu, B. Li, K. Lu, and D. Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv preprint arXiv:1712.05526*, 2017.
- [15] Y. Chen, Y. Nadji, A. Kountouras, F. Monrose, R. Perdisci, M. Antonakakis, and N. Vasiloglou. Practical attacks against graph-based clustering. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1125–1142, 2017.
- [16] J. M. Cohen, E. Rosenfeld, and J. Z. Kolter. Certified adversarial robustness via randomized smoothing. *arXiv preprint arXiv:1902.02918*, 2019.
- [17] H. Dai, H. Li, T. Tian, X. Huang, L. Wang, J. Zhu, and L. Song. Adversarial attack on graph structured data. *arXiv preprint arXiv:1806.02371*, 2018.
- [18] Q. Dai, X. Shen, L. Zhang, Q. Li, and D. Wang. Adversarial training methods for network embedding. In *The World Wide Web Conference*, pages 329–339, 2019.
- [19] Z. Deng, Y. Dong, and J. Zhu. Batch virtual adversarial training for graph convolutional networks. *arXiv preprint arXiv:1902.09192*, 2019.
- [20] Y. Dou, G. Ma, P. S. Yu, and S. Xie. Robust spammer detection by nash reinforcement learning. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 924–933, 2020.
- [21] N. Entezari, S. A. Al-Sayouri, A. Darvishzadeh, and E. E. Papalexakis. 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*, pages 169–177, 2020.
- [22] F. Feng, X. He, J. Tang, and T.-S. Chua. Graph adversarial training: Dynamically regularizing based on graph structure. *IEEE Transactions on Knowledge and Data Engineering*, 2019.
- [23] L. Franceschi, M. Niepert, M. Pontil, and X. He. Learning discrete structures for graph neural networks. *arXiv preprint arXiv:1903.11960*, 2019.
- [24] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017.
- [25] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [26] W. L. Hamilton, R. Ying, and J. Leskovec. Representation learning on graphs: Methods and applications. *arXiv preprint arXiv:1709.05584*, 2017.
- [27] X. Huang, J. Li, and X. Hu. Label informed attributed network embedding. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 731–739, 2017.
- [28] V. N. Ioannidis, D. Berberidis, and G. B. Giannakis. Graphsac: Detecting anomalies in large-scale graphs. *arXiv preprint arXiv:1910.09589*, 2019.
- [29] G. Jeh and J. Widom. Scaling personalized web search. In *Proceedings of the 12th international conference on World Wide Web*, pages 271–279, 2003.

- [30] J. Jia, B. Wang, X. Cao, and N. Z. Gong. Certified robustness of community detection against adversarial structural perturbation via randomized smoothing. *arXiv preprint arXiv:2002.03421*, 2020.
- [31] B. Jiang, Z. Zhang, D. Lin, J. Tang, and B. Luo. Semi-supervised learning with graph learning-convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11313–11320, 2019.
- [32] H. Jin and X. Zhang. Latent adversarial training of graph convolution networks. In *ICML Workshop on Learning and Reasoning with GraphStructured Representations*, 2019.
- [33] W. Jin, Y. Ma, X. Liu, X. Tang, S. Wang, and J. Tang. Graph structure learning for robust graph neural networks. *arXiv preprint arXiv:2005.10203*, 2020.
- [34] T. N. Kipf and M. Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016.
- [35] T. N. Kipf and M. Welling. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308*, 2016.
- [36] J. Klicpera, A. Bojchevski, and S. Günnemann. Predict then propagate: Graph neural networks meet personalized pagerank. *arXiv preprint arXiv:1810.05997*, 2018.
- [37] L. Landrieu and M. Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4558–4567, 2018.
- [38] J. Li, H. Zhang, Z. Han, Y. Rong, H. Cheng, and J. Huang. Adversarial attack on community detection by hiding individuals, 2020.
- [39] Y. Li, S. Bai, C. Xie, Z. Liao, X. Shen, and A. Yuille. Regional homogeneity: Towards learning transferable universal adversarial perturbations against defenses. *arXiv preprint arXiv:1904.00979*, 2019.
- [40] Y. Li, S. Bai, Y. Zhou, C. Xie, Z. Zhang, and A. L. Yuille. Learning transferable adversarial examples via ghost networks. In *AAAI*, pages 11458–11465, 2020.
- [41] Y. Li, W. Jin, H. Xu, and J. Tang. Deeprobust: A pytorch library for adversarial attacks and defenses. *arXiv preprint arXiv:2005.06149*, 2020.
- [42] Y. Lin, Z. Liu, M. Sun, Y. Liu, and X. Zhu. Learning entity and relation embeddings for knowledge graph completion. In *Twenty-ninth AAAI conference on artificial intelligence*, 2015.
- [43] X. Liu, S. Si, X. Zhu, Y. Li, and C.-J. Hsieh. A unified framework for data poisoning attack to graph-based semi-supervised learning. *arXiv preprint arXiv:1910.14147*, 2019.
- [44] J. Ma, S. Ding, and Q. Mei. Black-box adversarial attacks on graph neural networks with limited node access. *arXiv preprint arXiv:2006.05057*, 2020.
- [45] T. Ma, C. Xiao, J. Zhou, and F. Wang. Drug similarity integration through attentive multi-view graph auto-encoders. *arXiv preprint arXiv:1804.10850*, 2018.
- [46] Y. Ma, S. Wang, L. Wu, and J. Tang. Attacking graph convolutional networks via rewiring. *arXiv preprint arXiv:1906.03750*, 2019.
- [47] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [48] D. Marcheggiani and I. Titov. Encoding sentences with graph convolutional networks for semantic role labeling. *arXiv preprint arXiv:1703.04826*, 2017.
- [49] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, 27(1):415–444, 2001.
- [50] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [51] B. Perozzi, R. Al-Rfou, and S. Skiena. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 701–710, 2014.
- [52] T. Pham, T. Tran, D. Phung, and S. Venkatesh. Column networks for collective classification. *arXiv preprint arXiv:1609.04508*, 2016.
- [53] J. Qiu, Y. Dong, H. Ma, J. Li, K. Wang, and J. Tang. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 459–467, 2018.
- [54] A. Said, E. W. De Luca, and S. Albayrak. How social relationships affect user similarities.
- [55] P. Sen, G. Namata, M. Bilgic, L. Getoor, B. Galligher, and T. Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- [56] L. Sun, J. Wang, P. S. Yu, and B. Li. Adversarial attack and defense on graph data: A survey. *arXiv preprint arXiv:1812.10528*, 2018.
- [57] Y. Sun, S. Wang, X. Tang, T.-Y. Hsieh, and V. Honavar. Node injection attacks on graphs via reinforcement learning. *arXiv preprint arXiv:1909.06543*, 2019.
- [58] M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3319–3328. JMLR. org, 2017.
- [59] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, and Q. Mei. Line: Large-scale information network embedding. In *Proceedings of the 24th international conference on world wide web*, pages 1067–1077, 2015.

- [60] X. Tang, Y. Li, Y. Sun, H. Yao, P. Mitra, and S. Wang. Transferring robustness for graph neural network against poisoning attacks. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, pages 600–608, 2020.
- [61] S. Tao, H. Shen, Q. Cao, L. Hou, and X. Cheng. Adversarial immunization for improving certifiable robustness on graphs. *arXiv preprint arXiv:2007.09647*, 2020.
- [62] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [63] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. 2018.
- [64] B. Wang and N. Z. Gong. Attacking graph-based classification via manipulating the graph structure. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 2023–2040, 2019.
- [65] B. Wang, J. Jia, X. Cao, and N. Z. Gong. Certified robustness of graph neural networks against adversarial structural perturbation. *arXiv preprint arXiv:2008.10715*, 2020.
- [66] B. Wang, Y. Yao, S. Shan, H. Li, B. Viswanath, H. Zheng, and B. Y. Zhao. Neural cleanse: Identifying and mitigating backdoor attacks in neural networks. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 707–723. IEEE, 2019.
- [67] J. Wang, M. Luo, F. Suya, J. Li, Z. Yang, and Q. Zheng. Scalable attack on graph data by injecting vicious nodes. *arXiv preprint arXiv:2004.13825*, 2020.
- [68] X. Wang, H. Ji, C. Shi, B. Wang, Y. Ye, P. Cui, and P. S. Yu. Heterogeneous graph attention network. In *The World Wide Web Conference*, pages 2022–2032, 2019.
- [69] X. Wang, X. Liu, and C.-J. Hsieh. Graphdefense: Towards robust graph convolutional networks, 2019.
- [70] M. Waniek, T. P. Michalak, M. J. Wooldridge, and T. Rahwan. Hiding individuals and communities in a social network. *Nature Human Behaviour*, 2(2):139–147, Jan 2018.
- [71] M. Waniek, T. P. Michalak, M. J. Wooldridge, and T. Rahwan. Hiding individuals and communities in a social network. *Nature Human Behaviour*, 2(2):139–147, 2018.
- [72] H. Wu, C. Wang, Y. Tyshetskiy, A. Docherty, K. Lu, and L. Zhu. Adversarial examples for graph data: deep insights into attack and defense. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 4816–4823. AAAI Press, 2019.
- [73] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *arXiv preprint arXiv:1901.00596*, 2019.
- [74] Z. Xi, R. Pang, S. Ji, and T. Wang. Graph backdoor. *arXiv preprint arXiv:2006.11890*, 2020.
- [75] H. Xu, Y. Ma, H. Liu, D. Deb, H. Liu, J. Tang, and A. Jain. Adversarial attacks and defenses in images, graphs and text: A review. *arXiv preprint arXiv:1909.08072*, 2019.
- [76] K. Xu, H. Chen, S. Liu, P.-Y. Chen, T.-W. Weng, M. Hong, and X. Lin. Topology attack and defense for graph neural networks: An optimization perspective. *arXiv preprint arXiv:1906.04214*, 2019.
- [77] K. Xu, C. Li, Y. Tian, T. Sonobe, K.-i. Kawarabayashi, and S. Jegelka. Representation learning on graphs with jumping knowledge networks. *arXiv preprint arXiv:1806.03536*, 2018.
- [78] X. Xu, Y. Yu, B. Li, L. Song, C. Liu, and C. Gunter. Characterizing malicious edges targeting on graph neural networks. 2018.
- [79] X. Zang, Y. Xie, J. Chen, and B. Yuan. Graph universal adversarial attacks: A few bad actors ruin graph learning models, 2020.
- [80] A. Zhang and J. Ma. Defensevae: Defending against adversarial attacks on graph data via a variational graph autoencoder. *arXiv preprint arXiv:2006.08900*, 2020.
- [81] H. Zhang, T. Zheng, J. Gao, C. Miao, L. Su, Y. Li, and K. Ren. Towards data poisoning attack against knowledge graph embedding. *ArXiv*, abs/1904.12052, 2019.
- [82] X. Zhang and M. Zitnik. Gnn-guard: Defending graph neural networks against adversarial attacks. *arXiv preprint arXiv:2006.08149*, 2020.
- [83] Z. Zhang, J. Jia, B. Wang, and N. Z. Gong. Backdoor attacks to graph neural networks. *arXiv preprint arXiv:2006.11165*, 2020.
- [84] J. Zhou, G. Cui, Z. Zhang, C. Yang, Z. Liu, and M. Sun. Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434*, 2018.
- [85] Q. Zhou, L. Li, N. Cao, L. Ying, and H. Tong. Admiring: Adversarial multi-network mining. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 1522–1527. IEEE, 2019.
- [86] Q. Zhou, Y. Ren, T. Xia, L. Yuan, and L. Chen. Data poisoning attacks on graph convolutional matrix completion. In *Algorithms and Architectures for Parallel Processing*, 2020.
- [87] D. Zhu, Z. Zhang, P. Cui, and W. Zhu. Robust graph convolutional networks against adversarial attacks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1399–1407, 2019.
- [88] X. Zhu and Z. Ghahramani. Learning from labeled and unlabeled data with label propagation. 2002.
- [89] D. Zügner, A. Akbarnejad, and S. Günnemann. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 2847–2856, 2018.

- [90] D. Zügner, O. Borchert, A. Akbarnejad, and S. Günnemann. Adversarial attacks on graph neural networks: Perturbations and their patterns. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 14(5):1–31, 2020.
- [91] D. Zügner and S. Günnemann. Adversarial attacks on graph neural networks via meta learning. *arXiv preprint arXiv:1902.08412*, 2019.
- [92] D. Zügner and S. Günnemann. Certifiable robustness and robust training for graph convolutional networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 246–256, 2019.
- [93] D. Zügner and S. Günnemann. Certifiable robustness of graph convolutional networks under structure perturbations. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1656–1665, 2020.

APPENDIX

A. ADDITIONAL RESULTS

A.1 Global Measures for Topology Attack and DICE

The global measures for Topology attack are shown in Table 5; the global measures for DICE are shown in Table 6.

Table 5: Properties of attacked graphs under Topology attack.

Dataset	$r(\%)$	edges+	edges-	edges	ranks	clustering coefficients
Cora	0	0	0	5069	2192	0.2376
	5	255	0	5324	2292	0.2308
	10	508	0	5577	2369	0.2185
	15	762	0	5831	2417	0.2029
	20	1015	0	6084	2442	0.1875
25	1269	0	6338	2456	0.1736	
Citeseer	0	0	0	3668	1778	0.1711
	5	185	0	3853	1914	0.1666
	10	368	0	4036	2003	0.1568
	15	552	0	4220	2058	0.1429
	20	735	0	4403	2077	0.1306
25	918	0	4586	2087	0.1188	
Polblogs	0	0	0	16714	1060	0.3203
	5	716	96	17334	1213	0.2659
	10	1532	128	18118	1220	0.2513
	15	2320	146	18887	1221	0.2408
	20	3149	155	19708	1221	0.2317
25	3958	163	20509	1221	0.2238	

A.2 Local Measures for Topology Attack and DICE

The node feature similarity and label equality for Topology attack are shown in Figure 4 and Figure 6, respectively; the node feature similarity and label equality for DICE are shown in Figure 5 and Figure 7, respectively.

Table 6: Properties of attacked graphs under DICE attack.

Dataset	$r(\%)$	edge+	edge-	edges	ranks	clustering coefficients
Cora	0	0	0	5069	2192	0.2376
	5	125	128	5066	2208	0.2165
	10	251	255	5065	2241	0.1963
	15	377	383	5063	2256	0.1768
	20	504	509	5063	2262	0.1588
25	625	642	5053	2271	0.1436	
Citeseer	0	0	0	3668	1798	0.1711
	5	91	92	3667	1829	0.1574
	10	183	183	3668	1843	0.1404
	15	276	274	3670	1858	0.1279
	20	368	365	3672	1872	0.1170
25	462	455	36755	1871	0.1068	
Polblogs	0	0	0	16714	1060	0.3203
	5	420	415	16719	1151	0.2742
	10	846	825	16736	1191	0.2341
	15	1273	1234	16752	1206	0.2077
	20	1690	1652	16752	1214	0.1862
25	2114	2064	16765	1216	0.1675	

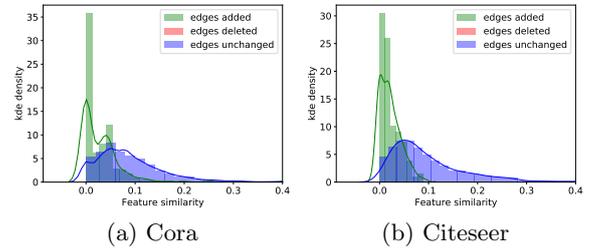


Figure 4: Node feature similarity for Topology attack.

A.3 Attack and Defense Performance for Topology Attack and DICE

The attack and defense performance for DICE is shown in Table 7; the performance for DICE is shown in Table 8.

B. OPEN SOURCE CODE

We list some open source implementations of representative algorithms in Table 9.

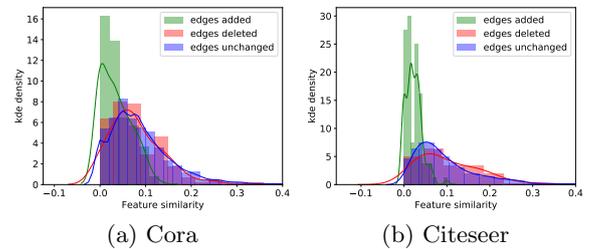


Figure 5: Node feature similarity for DICE attack.

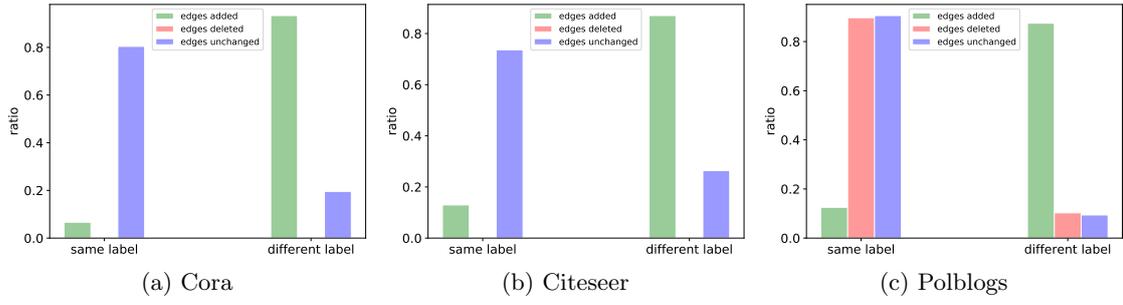


Figure 6: Label equality for Topology attack.

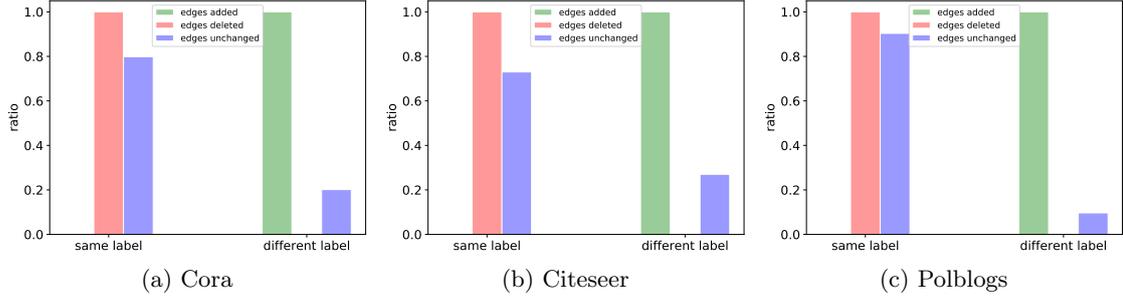


Figure 7: Label equality for DICE attack.

Table 7: Performance (accuracy) under Topology attack.

Dataset	r (%)	0	5	10	15	20	25
Cora	GCN	83.10	71.82	68.96	66.77	64.21	62.52
	Jaccard ¹	82.39	73.05	72.62	71.84	71.41	70.85
	SVD ²	77.97	78.17	75.92	73.69	72.03	70.11
	RGCN	84.81	72.68	71.15	69.38	67.92	67.23
	GAT	81.69	71.03	68.80	65.66	64.29	62.58
Citeseer	GCN	74.53	79.29	75.47	72.89	70.12	68.49
	Jaccard ¹	74.82	79.07	76.76	74.29	71.87	69.55
	SVD ²	70.32	78.17	75.92	73.69	72.03	70.11
	RGCN	74.41	78.13	75.93	73.93	72.32	70.60
	GAT	74.23	77.52	74.09	71.90	69.62	66.99
Polblogs	GCN	95.80	72.04	65.87	63.35	61.06	58.49
	SVD ²	94.99	71.90	65.42	63.01	60.74	58.26
	RGCN	95.60	71.27	65.30	62.76	60.25	57.89
	GAT	95.40	72.56	65.97	63.35	60.94	58.77

¹ Jaccard: GCN-Jaccard defense model.

² SVD: GCN-SVD defense model.

Table 8: Performance (accuracy) under DICE attack.

Dataset	r (%)	0	5	10	15	20	25
Cora	GCN	83.10	81.56	80.28	78.64	77.10	75.17
	Jaccard ¹	82.39	80.91	79.80	79.23	77.81	76.35
	SVD ²	77.97	75.24	73.33	70.92	69.47	67.29
	RGCN	84.81	83.53	82.56	80.70	79.30	77.89
	GAT	81.69	78.86	77.50	74.56	70.54	69.25
Citeseer	GCN	74.53	74.41	73.61	71.98	71.47	69.92
	Jaccard ¹	74.82	74.73	74.11	72.88	72.36	71.32
	SVD ²	70.32	70.68	69.89	68.59	67.09	66.65
	RGCN	74.41	74.46	73.93	72.37	71.61	70.85
	GAT	74.23	73.70	72.71	70.52	69.27	67.78
Polblogs	GCN	95.80	89.90	85.87	82.41	79.47	78.02
	SVD ²	94.99	90.88	87.79	85.09	83.64	81.25
	RGCN	95.60	89.86	86.11	82.25	78.81	77.18
	GAT	95.40	90.74	86.22	82.41	78.83	76.77

¹ Jaccard: GCN-Jaccard defense model.

² SVD: GCN-SVD defense model.

Table 9: A summary of open-source implementations

	Methods	Framework	Github Link
Attack	PGD, Min-max [76]	tensorflow	https://github.com/KaidiXu/GCN_ADV_Train
		pytorch	https://github.com/DSE-MSU/DeepRobust
	DICE [71]	python	https://github.com/DSE-MSU/DeepRobust
	FGA [13]	pytorch	https://github.com/DSE-MSU/DeepRobust
	IG-FGSM [72]	pytorch	https://github.com/DSE-MSU/DeepRobust
	Nettack [89]	tensorflow	https://github.com/danielzuegner/nettack
	Metattack [91]	tensorflow	https://github.com/danielzuegner/gnn-meta-attack
		pytorch	https://github.com/ChandlerBang/pytorch-gnn-meta-attack
	RL-S2V [17]	pytorch	https://github.com/Hanjun-Dai/graph_adversarial_attack
Bojchevski et al. [5]	tensorflow	https://github.com/abojchevski/node_embedding_attack	
GF-Attack [9]	tensorflow	https://github.com/SwiftieH/GFAttack	
Defense	RGCN [87]	tensorflow	https://github.com/thumanlab/nrlweb/blob/master/static/assets/download/RGCN.zip
		pytorch	https://github.com/DSE-MSU/DeepRobust
	GCN-Jaccard [72]	pytorch	https://github.com/DSE-MSU/DeepRobust
	GCN-SVD [21]	pytorch	https://github.com/DSE-MSU/DeepRobust
	Pro-GNN [33]	pytorch	https://github.com/ChandlerBang/Pro-GNN
	Adversarial Training [76]	tensorflow	https://github.com/KaidiXu/GCN_ADV_Train
		pytorch	https://github.com/DSE-MSU/DeepRobust
	PA-GNN [60]	tensorflow	https://github.com/tangxianfeng/PA-GNN
	Graph-Cert [6]	python	https://github.com/abojchevski/graph_cert
	DefenseVGAE [80]	pytorch	https://github.com/zhangao520/defense-vgae
Bojchevski et al. [7]	pytorch	https://github.com/abojchevski/sparse_smoothing	
Zügner et al. [93]	python	https://github.com/danielzuegner/robust-gcn-structure	