

Turning Strengths into Weaknesses: A Certified Robustness Inspired Attack Framework against Graph Neural Networks

Binghui Wang^{§*}, Meng Pang^{†*}, and Yun Dong[‡]

[§]Department of Computer Science, Illinois Institute of Technology

[†]School of Mathematics and Computer Sciences, Nanchang University

[‡]Department of Visual and Performing Arts, Education, and Sciences, Waubesa Community College

Email: [§]bwang70@iit.edu, [†]mengpang@ncu.edu.cn, [‡]ydong@waubesa.edu

Abstract

Graph neural networks (GNNs) have achieved state-of-the-art performance in many graph learning tasks. However, recent studies show that GNNs are vulnerable to both test-time evasion and training-time poisoning attacks that perturb the graph structure. While existing attack methods have shown promising attack performance, we would like to design an attack framework to further enhance the performance. In particular, our attack framework is inspired by certified robustness, which was originally used by defenders to defend against adversarial attacks. We are the first, from the attacker perspective, to leverage its properties to better attack GNNs. Specifically, we first derive nodes' certified perturbation sizes against graph evasion and poisoning attacks based on randomized smoothing, respectively. A larger certified perturbation size of a node indicates this node is theoretically more robust to graph perturbations. Such a property motivates us to focus more on nodes with smaller certified perturbation sizes, as they are easier to be attacked after graph perturbations. Accordingly, we design a certified robustness inspired attack loss, when incorporated into (any) existing attacks, produces our certified robustness inspired attack counterpart. We apply our framework to the existing attacks and results show it can significantly enhance the existing base attacks' performance.

1. Introduction

Learning with graphs, such as social networks, citation networks, chemical networks, has attracted significant attention recently. Among many methods, graph neural networks (GNNs) [14,33,38,41,44] have achieved state-of-the-art performance in graph related tasks such as node classification, graph classification, and link prediction. However, recent studies [8,19,20,23,30,34,36,37,39,40,50,51] show that GNNs are vulnerable to both test-time graph evasion

attacks and training-time graph poisoning attacks¹. Take GNNs for node classification as an instance, graph evasion attacks mean that, given a learnt GNN model and a (clean) graph, an attacker carefully perturbs the graph structure (i.e., inject new edges to or remove the existing edges from the graph) such that as many testing nodes as possible are misclassified by the GNN model. Whereas, graph poisoning attacks mean that, given a GNN algorithm and a graph, an attacker carefully perturbs the graph structure in the training phase, such that the learnt GNN model misclassifies as many testing nodes as possible in the testing phase. While existing methods have shown promising attack performance, we want to ask: Can we design a general *attack framework* that can further enhance both the existing graph evasion and poisoning attacks to GNNs? The answer is yes.

We design an attack framework inspired by certified robustness. Certified robustness was originally used by *defenders* to guarantee the robustness of classification models against evasion attacks. Generally speaking, a testing example (e.g., an image or a node) with a better certified robustness guarantee indicates this example is *theoretically* more robust to adversarial (e.g., pixel or graph) perturbations. While certified robustness is mainly derived for doing the good, *attackers*, on the other hand, can also leverage its property to do the bad. For instance, when an attacker knows the certified robustness of nodes in a graph, he can base on nodes' certified robustness to *reversely* reveal the vulnerable region of the graph and leverage this vulnerability to design better attacks. We are inspired by such property of certified robustness and design the first certified robustness inspired attacks to GNNs.

Our attack framework consists of three parts: i) Inspired by the state-of-the-art randomized smoothing based certified robustness against *evasion attacks* to image models [7,28] and GNN models [35], we first propose to generalize randomized smoothing and derive the node's cer-

¹We mainly consider the graph structure attack in the paper, as it is more effective than the feature attack. However, our attack framework can be easily extended to the feature attack.

*Corresponding authors

tified perturbation size against graph *poisoning attacks* to GNNs. Particularly, a larger certified perturbation size of a node indicates this node is *theoretically* more robust to adversarial graph perturbations. In other words, an attacker needs to perturb more edges during the training phase in order to make this node wrongly predicted by the learnt GNN model. This property inspires us to focus more on disrupting nodes with relatively smaller certified perturbation sizes under a given perturbation budget. ii) We design a certified robustness inspired attack loss. Specifically, we modify the classic node-wise loss by assigning each node a weight based on its certified perturbation size—A node with a larger/smaller certified perturbation size will be assigned a smaller/larger weight. In doing so, losses for nodes with smaller certified perturbation sizes will be enlarged, and most of the perturbation budget will be automatically allocated to perturb these nodes. Thus, more nodes will be misclassified with the given perturbation budget. iii) We design the certified robustness inspired attack framework to generate adversarial graph perturbations to GNNs, based on our certified robustness inspired attack loss. We emphasize that, as our new attack loss only modifies the existing attack loss with certified perturbation size defined node weights, any existing graph evasion or poisoning attack method can be used as the base attack in our framework.

We apply our certified robustness inspired attack framework to the state-of-the-art graph evasion and poisoning attacks [40, 51] to GNNs. Evaluation results on multiple benchmark datasets show our attack framework can substantially enhance the attack performance of the base attacks. Our contributions are as follows:

- We propose a certified robustness inspired attack framework to GNNs. Our framework can be plugged into any existing graph evasion and poisoning attacks.
- To our best knowledge, we are the first work to use certified robustness for an attack purpose.
- Evaluation results validate the effectiveness of our attack framework when applied to the existing attacks to GNNs.

2. Background and Preliminaries

2.1. Graph Neural Networks (GNNs)

Let $G = (\mathcal{V}, \mathcal{E})$ be a graph, where $u \in \mathcal{V}$ is a node, $(u, v) \in \mathcal{E}$ is an edge between u and v . Let $\mathbf{A} \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$ be the adjacency matrix. As \mathbf{A} contains all graph structure information, we will interchangeably use \mathbf{A} and G to indicate the graph in the paper. We mainly consider GNNs for node classification. Each node $u \in \mathcal{V}$ has a label y_u from a label set \mathcal{Y} . Let \mathcal{V}_{Tr} and \mathcal{V}_{Te} be the set of training nodes and testing nodes, respectively. Given a GNN algorithm \mathcal{A} , which takes the graph $G(\mathbf{A})$ and training nodes \mathcal{V}_{Tr} as an input and produces a node classifier f_θ parameterized by θ , i.e., $f_\theta = \mathcal{A}(\mathbf{A}, \mathcal{V}_{Tr})$. The node classifier f_θ inputs $G(\mathbf{A})$ and outputs labels for all nodes, i.e.,

$f_\theta : \mathbf{A} \rightarrow \mathcal{Y}^{|\mathcal{V}|}$. To learn f_θ , a common way is to minimize a loss function \mathcal{L} defined on the training nodes \mathcal{V}_{Tr} and the graph $G(\mathbf{A})$ as follows:

$$\min_{\theta} \mathcal{L}(f_\theta, \mathbf{A}, \mathcal{V}_{Tr}) = \sum_{u \in \mathcal{V}_{Tr}} \ell(f_\theta(\mathbf{A}; u), y_u), \quad (1)$$

where $f_\theta(\mathbf{A}; u)$ is the predicted label of a node u . After learning f_{θ^*} , a testing node $v \in \mathcal{V}_{Te}$ is then predicted a label as $\hat{y}_v = f_{\theta^*}(\mathbf{A}; v)$.

2.2. Adversarial Attacks to GNNs

We denote by $\delta \in \{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$ the adversarial *graph perturbation*, where $\delta_{s,t} = 1$ (or 0) means the attacker perturbs (or keeps) the edge status between a node pair (s, t) . Moreover, we denote $\mathbf{A} \oplus \delta$ as the perturbed graph, with \oplus the element-wise XOR operator. For instance, if there is an (or no) edge between (u, v) , i.e., $A_{uv} = 1$ (or $A_{uv} = 0$), perturbing this edge status (i.e. $\delta_{u,v} = 1$) means removing the edge (or injecting a new edge), i.e., $A_{u,v} \oplus \delta_{u,v} = 0$ (or $A_{u,v} \oplus \delta_{u,v} = 1$) to the graph. We assume an attacker has a perturbation budget Δ , i.e., $\|\delta\|_0 \leq \Delta$, meaning at most Δ number of edges can be perturbed by the attacker.

Graph evasion attacks to GNNs. In graph evasion attacks, given a learnt node classifier f_{θ^*} , an attacker carefully crafts a graph perturbation δ to the graph G such that f_{θ^*} predicts nodes' labels using the perturbed graph $\mathbf{A} \oplus \delta$ as the attacker desires. For instance, an attacker desires as many testing nodes as possible to be misclassified by f_{θ^*} (called *untargeted attack*) under the perturbation budget Δ . Formally, an attacker aims to maximize the following 0-1 (*attack*) loss:

$$\max_{\delta} \sum_{v \in \mathcal{V}_{Te}} \mathbf{1}[f_{\theta^*}(\mathbf{A} \oplus \delta; v) \neq y_v], \text{ s.t. } \|\delta\|_0 \leq \Delta, \quad (2)$$

where $\mathbf{1}[\cdot]$ is an indicator function, whose value is 1 if the condition satisfies and 0, otherwise.

The above problem is challenging to solve in that the indicator function is hard to be optimized. In practice, an attacker will solve an alternative optimize problem as below:

$$\max_{\delta} \sum_{v \in \mathcal{V}_{Te}} \ell(f_{\theta^*}(\mathbf{A} \oplus \delta; v), y_v), \text{ s.t. } \|\delta\|_0 \leq \Delta. \quad (3)$$

For instance, [40] design the state-of-the-art PGD evasion attack by solving Equation 3.

Graph poisoning attacks to GNNs. In graph poisoning attacks, an attacker specifies a GNN algorithm \mathcal{A} and carefully perturbs the graph G with a graph perturbation δ in the training phase, such that the learnt node classifier f_{θ^*} misclassifies as many testing nodes as possible on the perturbed graph $\mathbf{A} \oplus \delta$ in the testing phase. Formally, it solves the following bilevel optimization problem:

$$\begin{aligned} & \max_{\delta} \sum_{v \in \mathcal{V}_{Te}} \mathbf{1}[f_{\theta^*}(\mathbf{A} \oplus \delta; v) \neq y_v], \\ & \text{s.t. } \theta^* = \arg \min_{\theta} \sum_{u \in \mathcal{V}_{Tr}} \mathbf{1}[f_{\theta^*}(\mathbf{A} \oplus \delta; u) \neq y_u], \|\delta\|_0 \leq \Delta, \end{aligned} \quad (4)$$

where the inner optimization problem is learning the node classifier f_{θ^*} on the perturbed graph $\mathbf{A} \oplus \delta$ with training nodes \mathcal{V}_{Tr} , while the outer optimization problem is learning to generate the graph perturbation δ to maximally misclassify testing nodes \mathcal{V}_{Te} with the learnt node classifier f_{θ^*} .

In practice, the testing nodes \mathcal{V}_{Te} 's labels are unavailable during training, and thus we cannot directly optimize Equation 4. In addition, the indicator function in Equation 4 is hard to optimize. A common strategy to address this issue is by instead maximizing the loss on the *training nodes* \mathcal{V}_{Tr} [40, 51] and using an alternative continuous loss. Specifically, it solves the following alternative bilevel optimization problem:

$$\begin{aligned} \max_{\delta} \quad & \sum_{v \in \mathcal{V}_{Tr}} \ell(f_{\theta^*}(\mathbf{A} \oplus \delta; v), y_v), \\ \text{s.t. } \theta^* = \arg \min_{\theta} \quad & \sum_{u \in \mathcal{V}_{Tr}} \ell(f_{\theta}(\mathbf{A} \oplus \delta; u), y_u), \|\delta\|_0 \leq \Delta. \end{aligned} \quad (5)$$

This is based on the intuition that if a node classifier misclassifies a large number of training nodes, then it generalizes poorly and thus is also very likely to misclassify a large number of testing nodes.

2.3. Certified Robustness to Graph Evasion Attacks

We introduce certified robustness achieved via the state-of-the-art randomized smoothing [7, 15, 17]. Randomized smoothing was originally designed to build certified robust machine learning classifiers against evasion attacks. It is applicable to any classifier and scalable to large models, e.g., deep neural networks. Here, we introduce randomized smoothing that defends against graph evasion attacks to GNNs [35]. It consists of the following three steps.

Constructing a smoothed node classifier. Given a base node classifier f , a graph G , and a testing node u with label y_u , randomized smoothing builds a *smoothed node classifier* g via adding a random noise matrix ϵ to G . Formally,

$$g(\mathbf{A}; u) = \arg \max_{c \in \mathcal{Y}} \Pr(f(\mathbf{A} \oplus \epsilon; u) = c), \quad (6)$$

where $\Pr(f(\mathbf{A} \oplus \epsilon; u) = c)$ is the probability that the base node classifier f predicts label c on the noisy graph $\mathbf{A} \oplus \epsilon$ and $g(\mathbf{A}; u)$ is the predicted label for u by the smoothed node classifier g . ϵ has the following probability distribution in the binary space $\{0, 1\}^{|\mathcal{V}| \times |\mathcal{V}|}$:

$$\Pr(\epsilon_{s,t} = 0) = \beta, \Pr(\epsilon_{s,t} = 1) = 1 - \beta, \forall s, t \in \mathcal{V}. \quad (7)$$

Equation 7 means that for each pair of nodes (s, t) in the graph, we keep its edge status (i.e., $A_{s,t}$) with probability β and change its edge status with probability $1 - \beta$.

Deriving the certified robustness of graph evasion attacks to GNNs. Suppose $g(\mathbf{A}; u) = y_u$, meaning that the smoothed node classifier g correctly predicts u . Then, g

provably predicts the correct label for u once the graph perturbation δ is bounded. Formally [35]:

$$g(\mathbf{A} \oplus \delta; u) = y_u, \forall \|\delta\|_0 \leq K(p_{y_u}), \quad (8)$$

where $p_{y_u} \leq \Pr(f(\mathbf{A} \oplus \epsilon; u) = y_u)$ is a lower bound of the probability that f predicts the correct label y_u on the noisy graph $\mathbf{A} \oplus \epsilon$. $K(p_{y_u})$ is called node u 's *certified perturbation size*, indicating that g provably predicts the correct label when an attacker *arbitrarily* perturbs (at most) $K(p_{y_u})$ edge status in the graph G . In other words, if a node has a larger certified perturbation size, then it is certifiably more robust to adversarial graph perturbation.

Computing the certified perturbation size in practice.

Note that $K(p_{y_u})$ is (positively) related to p_{y_u} , which can be estimated via the Monte Carlo algorithm [7, 35]. Specifically, given a node classifier f , a graph $G(\mathbf{A})$, and a testing node u , we first sample N random noise matrices $\epsilon^1, \dots, \epsilon^N$ from the noise distribution defined in Equation 7 and add each noise matrix ϵ^j to the graph G to construct N noisy graphs $\mathbf{A} \oplus \epsilon^1, \dots, \mathbf{A} \oplus \epsilon^N$. Then, we use the node classifier f to predict u 's label on the N noisy graphs and compute the frequency of each label c , i.e., $N_c = \sum_{j=1}^N \mathbb{I}(f(\mathbf{A} \oplus \epsilon^j, u) = c)$ for $c \in \mathcal{Y}$. Then, we can estimate p_{y_u} as

$$p_{y_u} = B(\alpha; N_{y_u}, N - N_{y_u} + 1), \quad (9)$$

where $1 - \alpha$ is the confidence level and $B(\alpha; a, b)$ is the α -th quantile of Beta distribution with shape parameters a and b . With p_{y_u} , we can compute $K(p_{y_u})$, and details of computing $K(p_{y_u})$ can be seen in [35].

3. Certified Robustness to Graph Poisoning Attacks via Randomized Smoothing

Existing randomized smoothing mainly certifies robustness of *evasion attacks*. In this section, we generalize it and derive certified robustness of graph poisoning attacks. Our key idea is to extend randomized smoothing from the *classifier* perspective to a general *function* perspective. In particular, we will build a base function, a smoothed function, and then adapt randomized smoothing to certify robustness to poisoning attacks using the smoothed function. Such certified robustness guides us to design more effective graph poisoning attacks, as shown in Section 4.

Building a base function. Suppose we have a graph $G(\mathbf{A})$, training nodes \mathcal{V}_{Tr} , and a GNN algorithm \mathcal{A} that takes the graph and training nodes as an input and learns a node classifier f , i.e., $f = \mathcal{A}(\mathbf{A}, \mathcal{V}_{Tr})$. We use the learnt f to predict the label for a testing node v . Then, we can integrate the entire process of training the node classifier f and testing the node v as a function $\hat{f}(\mathbf{A}, \mathcal{V}_{Tr}; v)$. In other words, the function \hat{f} is the composition of learning the node classifier f and predicting a node v . We view \hat{f} as the base function.

Constructing a smoothed function. In graph poisoning attacks, an attacker aims to perturb the graph in the training phase. To apply randomized smoothing, we first add a random noise matrix ϵ to the graph, where each entry $\epsilon_{s,t}$ is drawn from a discrete distribution, e.g., defined in Equation 7. As we add random noise ϵ to the graph G , the output of the base function \tilde{f} is also random. Then, inspired by Equation 6, we define the smoothed function \tilde{g} as follows:

$$\tilde{g}(\mathbf{A}, \mathcal{V}_{Tr}; v) = \arg \max_{c \in \mathcal{Y}} \Pr(\tilde{f}(\mathbf{A} \oplus \epsilon, \mathcal{V}_{Tr}; v) = c), \quad (10)$$

where $\Pr(\tilde{f}(\mathbf{A} \oplus \epsilon, \mathcal{V}_{Tr}; v) = c)$ is the probability that v is predicted to be a label c by a GNN model trained on a noisy graph $\mathbf{A} \oplus \epsilon$ using training nodes \mathcal{V}_{Tr} . $\tilde{g}(\mathbf{A}, \mathcal{V}_{Tr}; v)$ is the predicted label for v by the smoothed function \tilde{g} .

Deriving the certified robustness of graph poisoning attacks to GNNs. An attacker adds an adversarial graph perturbation δ to the graph $G(\mathbf{A})$ to produce a perturbed graph $\mathbf{A} \oplus \delta$, where $\delta_{s,t}$ is the perturbation added to change the edge status of the node pair (s, t) in the graph G during training. Then, we can leverage the results in Equation 8 to derive the certified perturbation size against graph poisoning attacks. Specifically, we have:

$$\tilde{g}(\mathbf{A} \oplus \delta, \mathcal{V}_{Tr}; v) = y_v, \quad \forall \|\delta\|_0 \leq K(p_{y_v}), \quad (11)$$

where $p_{y_v} \leq \Pr(\tilde{f}(\mathbf{A} \oplus \epsilon, \mathcal{V}_{Tr}; v) = y_v)$ is a lower bound probability. Our result means the smoothed function \tilde{g} probably predicts the correct label for v when (at most) $K(p_{y_v})$ edge statuses in the graph are *arbitrarily* poisoned by an attacker *in the training phase*.

Computing the certified perturbation size in practice. Given a GNN algorithm \mathcal{A} , a graph $G(\mathbf{A})$, training nodes \mathcal{V}_{Tr} , a discrete noise distribution defined in Equation 7, and a node v , we first sample N random noise matrices $\epsilon^1, \dots, \epsilon^N$ from the discrete noise distribution and add each noise to the graph $G(\mathbf{A})$ to construct N noisy graphs $\mathbf{A} \oplus \epsilon^1, \dots, \mathbf{A} \oplus \epsilon^N$. Then, we train N node classifiers $\tilde{f}^1 = \mathcal{A}(\mathbf{A} \oplus \epsilon^1, \mathcal{V}_{Tr}), \dots, \tilde{f}^N = \mathcal{A}(\mathbf{A} \oplus \epsilon^N, \mathcal{V}_{Tr})$. We use each of the N node classifiers to predict v 's label and compute the frequency of each label c , i.e., $N_c = \sum_{j=1}^N \mathbb{I}(\tilde{f}^j(\mathbf{A} \oplus \epsilon^j, \mathcal{V}_{Tr}; v) = c)$ for $c \in \mathcal{Y}$. Finally, we estimate p_{y_v} using Equation 9 and use it to calculate the certified perturbation size, following [35]. Note the trained N node classifiers is re-used to predict node labels and compute certified perturbation size for different nodes.

4. Certified Robustness Inspired Attack Framework against GNNs

In this section, we will design our attack framework to GNNs inspired by certified robustness. Our attack framework can be seamlessly plugged into the existing graph evasion and poisoning attacks to design more effective attacks.

4.1. Motivation and Observation

Certified robustness, more specifically, certified perturbation size derived in Section 2.3 and Section 3, was used by *defenders* to defend GNN models against attacks. On the other hand, from the *attacker* perspective, he can leverage the properties of certified robustness to better attack GNN models. Specifically, certified perturbation size of a node characterizes the extent to which the GNN model *provably* and accurately predicts this node against the worst-case graph perturbation. An attacker can base on nodes' certified perturbation sizes to *reverse* reveal the vulnerable region of the graph and leverage this vulnerability to design better attacks. In particular, we have the following observation that reveals the *inverse* relationship between a node's certified perturbation size and the perturbation associated with this node when designing the attack.

Observation 1: *A node with a larger (smaller) certified perturbation size should be disrupted with a smaller (larger) number of perturbed edges.*

If a node has a larger (smaller) certified perturbation size, it means this node is more (less) robust to graph perturbations. To misclassify this node, an attacker should allocate a larger (smaller) number of perturbed edges. Thus, to design more effective attacks (i.e., misclassify more nodes) with a perturbation budget, an attacker should avoid disrupting nodes with larger certified perturbation sizes, but focus on nodes with smaller certified perturbation sizes.

Based on the above observation, our attack needs to solve three correlated problems: i) How to obtain the node's certified perturbation size for both graph evasion and poisoning attacks? ii) How to allocate the perturbation budget in order to disrupt the nodes with smaller certified perturbation sizes? iii) How to generate the adversarial graph perturbation for both evasion and poisoning attacks? To address i), we adopt the derived node's certified perturbation size against graph evasion attacks (Section 2.3) and graph poisoning attacks (Section 3). To address ii), we design a certified robustness inspired loss, by maximizing which an attacker will put more effort into disrupting nodes with smaller certified perturbation sizes. To address iii), we design a certified robustness inspired attack framework, where any existing graph evasion/poisoning attacks to GNNs can be adopted as the base attack in our framework.

4.2. Certified Robustness Inspired Loss Design

Suppose we have obtained nodes' certified perturbation sizes. To perform a more effective attack, a naive solution is that the attacker sorts all nodes' certified perturbation sizes in an ascending order, and then carefully perturbs the edges to misclassify the sorted nodes one-by-one until reaching the perturbation budget. However, this solution is both computationally intensive—as it needs to solve an optimization

problem for each node; and suboptimal—as all nodes and the associated edges collectively make predictions and perturbing an edge could affect predictions of many nodes.

We design a certified perturbation size inspired loss that assists to *automatically* seeking the “ideal” edges to be perturbed for both evasion attacks and poisoning attacks. Particularly, we notice that the loss function of evasion attacks in Equation 3 or poisoning attacks in Equation 5 is defined per node. Then, we propose to modify the loss function in Equation 3 or Equation 5 by assigning each node with a weight and multiplying each node loss with the corresponding weight, where the node weight has a strong connection with the node’s certified perturbation size. Formally, we design the certified perturbation size inspired loss as follows:

$$\mathcal{L}_{CR}(f_\theta, \mathbf{A}, \mathcal{V}_T) = \sum_{u \in \mathcal{V}_T} w(u) \cdot \ell(f_\theta(\mathbf{A}; u), y_u), \quad (12)$$

where $\mathcal{V}_T = \mathcal{V}_{T_e}$ for evasion attacks and $\mathcal{V}_T = \mathcal{V}_{T_r}$ for poisoning attacks; and $w(u)$ is the weight of the node u . Note that when setting all nodes with an equal weight, our certified perturbation size inspired loss reduces to the conventional loss in Equation 3 or Equation 5. Next, we show the *inverse* relationship between the node’s certified perturbation size and the assigned weight.

Observation 2: *A node with a larger (smaller) certified perturbation size is assigned a smaller (larger) weight.*

As shown in **Observation 1**, we should disrupt more nodes with smaller certified perturbation sizes, as these nodes are more vulnerable. In other words, we should put more weights on nodes with smaller certified perturbation sizes to enlarge these nodes’ losses—making these nodes easier to be misclassified with graph perturbations. In contrast, we should put smaller weights on nodes with larger certified perturbation sizes, in order to save the usage of the perturbation budget. Formally, we assign the node weight such that $w(u) \sim 1/K(p_{y_u})$. There are many ways to assign node weights satisfying the inverse relationship. In this paper, for instance, we propose to define node weights as

$$w(u) = \frac{1}{1 + \exp(a \cdot K(p_{y_u}))}, \quad (13)$$

where a is a tunable hyperparameter. We can observe that the node weight is exponentially decreased as the node’s certified perturbation size increases. Such a property ensures that the majority of perturbed edges are used for disrupting nodes with smaller certified perturbation sizes (See Figures 3) when performing the attack.

4.3. Certified Robustness Inspired Attack Design

Based on the derived certified perturbation size and our certified robustness inspired loss, we now propose to gener-

ate graph perturbations against GNNs with both graph evasion and poisoning attacks.

Certified robustness inspired graph evasion attacks to generate graph perturbations. We can choose any graph evasion attack to GNNs as the base evasion attack. In particular, given the attack loss from any existing evasion attack, we only need to modify the loss by multiplying it with our certification perturbation sizes defined node weights. For instance, we can use the PGD attack [40] as the base evasion attack. We replace its attack loss by our certified robustness inspired loss \mathcal{L}_{CR} in Equation 12. Then, we have our certified robustness inspired PGD (CR-PGD) evasion attack that iteratively generates graph perturbations as follows:

$$\delta = \text{Proj}_{\mathbb{B}}(\delta + \eta \cdot \nabla_{\delta} \mathcal{L}_{CR}(f_\theta, \mathbf{A} \oplus \delta, \mathcal{V}_{T_e})), \quad (14)$$

where η is the learning rate in PGD, $\mathbb{B} = \{\delta : \mathbf{1}^T \delta \leq \Delta, \delta \in [0, 1]^{|\mathcal{V}| \times |\mathcal{V}|}\}$ is the allowable perturbation set, and

$$\text{Proj}_{\mathbb{B}}(\mathbf{a}) = \begin{cases} \Pi_{[0,1]}(\mathbf{a} - \mu \mathbf{1}), & \text{if } \mathbf{1}^T \Pi_{[0,1]}(\mathbf{a} - \mu \mathbf{1}) = \Delta, \\ \Pi_{[0,1]}(\mathbf{a}), & \text{if } \mathbf{1}^T \Pi_{[0,1]}(\mathbf{a}) \leq \Delta, \end{cases} \quad (15)$$

where $\mu > 0$, $\Pi_{[0,1]}(x) = x$ if $x \in [0, 1]$, 0 if $x < 0$, and 1 if $x > 1$. The final graph perturbation is used to perform the evasion attack.

Certified robustness inspired graph poisoning attacks to generate graph perturbations. Likewise, we can choose any graph poisoning attack to GNNs as the base poisoning attack. Given the bilevel loss from any existing poisoning attack, we simply modify each loss by multiplying it with our certified perturbation sizes’ defined node weights. Specifically, we have

$$\begin{aligned} & \max_{\delta} \mathcal{L}_{CR}(f_{\theta^*}, \mathbf{A} \oplus \delta, \mathcal{V}_{T_r}), \quad (16) \\ & \text{s.t. } \theta^* = \arg \min_{\theta} \mathcal{L}_{CR}(f_\theta, \mathbf{A} \oplus \delta, \mathcal{V}_{T_r}), \|\delta\|_0 \leq \Delta, \quad (17) \end{aligned}$$

where $\mathcal{L}_{CR}(f_\theta, \mathbf{A} \oplus \delta, \mathcal{V}_{T_r}) = \sum_{v \in \mathcal{V}_{T_r}} w(v) \cdot \ell(f_\theta(\mathbf{A} \oplus \delta, v), y_v)$. Then, solving Equation 16 and Equation 17 produces the poisoning attack graph perturbations with our framework.

Algorithm 1 and Algorithm 2 in Appendix show two instances of applying our CR inspired attack framework to the PGD evasion attack and Minmax [40] poisoning graph, respectively. To save time, we calculate nodes’ certified perturbation sizes per INT iterations. Then, comparing with PGD, the computational overhead of our CR-PGD is calculating the node’s certified perturbation size with a set of N sampled noises every INT iterations, which only involves making predictions on N noisy matrices and is efficient. Note that the predictions are independent and can be also parallelized. Comparing with Minmax, the computational overhead of our CR-Minmax is to independently train (a small number of) N models every INT iterations, which can be implemented in parallel.

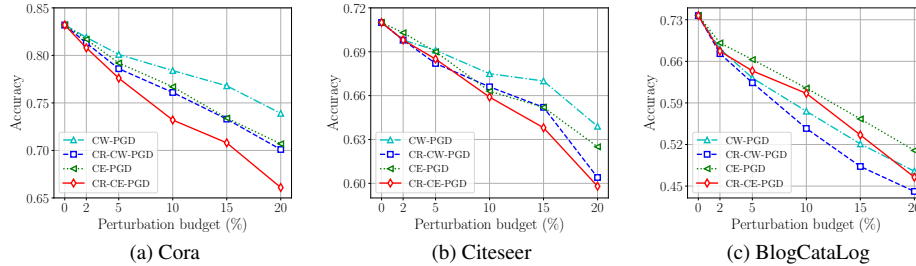


Figure 1. Evasion attack accuracy vs. perturbation budget.

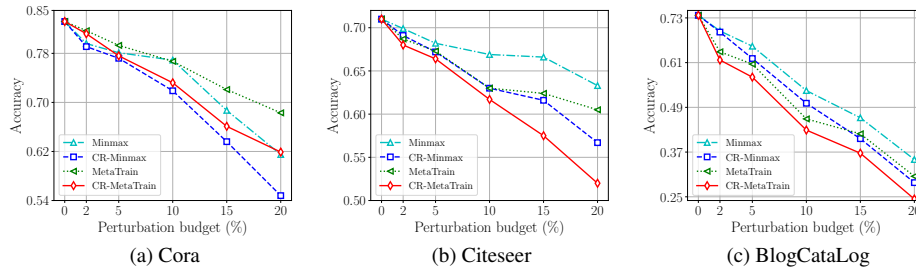


Figure 2. Poisoning attack accuracy vs. perturbation budget.

5. Experiments

5.1. Setup

Datasets and GNN models. Following [40, 51], we evaluate our attacks on benchmark graph datasets, i.e., Cora, Citeseer [29], and BlogCatalogs [27]. Table 3 in Appendix shows basic statistics of these graphs. We choose Graph Convolutional Network (GCN) [14] as the targeted GNN model, also following [40, 51].

Base attack methods. For graph evasion attacks, we choose the PGD attack [40]² that uses the cross-entropy loss and CW loss [3] as the base attack methods, and denote the two attacks as CE-PGD and CW-PGD, respectively. For graph poisoning attacks, we choose the Minmax attack [40] and MetaTrain attack [51]³ as the base attack methods. We apply our CR inspired attack framework to these evasion and poisoning attacks and denote them as CR-CE-PGD, CR-CW-PGD, CR-Minmax, and CR-MetaTrain, respectively. All attacks are implemented in PyTorch and run on a Linux server with 96 core 3.0GHz CPU, 768GB RAM, and 8 Nvidia A100 GPUs.

Training and testing. Following [51], we split the datasets into 10% training nodes, 10% validation nodes, and 80% testing nodes. The validation nodes are used to tune the hyperparameters, and the testing nodes are used to evaluate the attack performance. We repeat all attacks on 5 different splits of the training/evaluation/testing nodes and report the mean attack accuracy on testing nodes, i.e., fraction of testing nodes are misclassified after the attack.

Parameter settings. Without otherwise mentioned, we set the perturbation budget Δ as 20% of the total number of edges in a graph (before attack). We set the parameter

$\beta = 0.999$ in the noise distribution Equation 7, the confidence level $1 - \alpha = 0.9$, the number of samples N in Monte Carlo sampling to calculate node’s certified perturbation size is set to be 200 and 20 in evasion attacks and poisoning attacks, respectively, and $a = 1$ in Equation 13. The number of iterations T is 100 and 10, and the interval is set to be $INT = 10$ and $INT = 2$ in evasion attacks and poisoning attacks, respectively. The other hyperparameters in CE-PGD, CW-PGD, Minmax, and MetaTrain are selected based on their source code, and we set equal values in our CR inspired attack counterparts. We also study the impact of the important hyperparameters that could affect our attack performance: Δ , N , $1 - \alpha$, β , and a . When studying the impact of a hyperparameter, we fix the other hyperparameters to be their default values.

5.2. Attack Results

Our attack framework is effective. Figure 1 and Figure 2 show the evasion attack accuracy and poisoning attack accuracy of the base attacks and those with our attack framework vs. perturbation budget, respectively. We can observe that *our certified robustness inspired attack framework can enhance the base attack performance in all datasets*. For instance, when attacking GCN on Cora and the perturbation ratio is 20%, our CR-CE-PGD and CR-CW-PGD have a relative 7.0% and 5.6% gain over the CE-PGD and CW-PGD evasion attacks. Moreover, CR-Minmax and CR-MetaTrain have a relative 12.2% and 10.3% gain over the Minmax and MetaTrain poisoning attacks. These results demonstrate that the node’s certified robustness can indeed guide our attack framework to find the more vulnerable region in the graph to be perturbed, which helps to better allocate the perturbation budget, and thus makes the base attacks with our attack framework misclassify more nodes.

²https://github.com/KaidiXu/GCN_ADV_Train

³<https://www.kdd.in.tum.de/gnn-meta-attack>

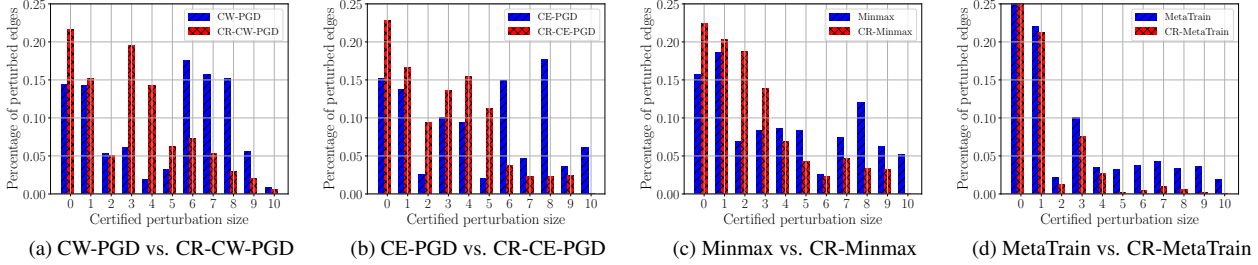


Figure 3. Distribution of the perturbed edges vs. node’s certified perturbation size on Citeseer.

To further understand the effectiveness of our framework, we visualize the distribution of the perturbed edges vs. node’s certified perturbation size. Specifically, we first obtain the perturbed edges via the base attacks and our CR inspired attacks, and calculate testing/training nodes’ certified perturbation sizes for evasion/poisoning attacks, respectively. Then we plot the distribution of the perturbed edges vs node’s certified perturbation size. Specifically, if a perturbed edge is connected with a testing/training node in the evasion/poisoning attack, then we map this perturbed edge to this node’s certified perturbation size. Our intuition is that a perturbed edge affects its connected node the most. Figure 3 shows the results on Citeseer (We observe that the conclusions on the other datasets are similar). We can see that a majority number of the perturbed edges connect with testing/training nodes that have relatively smaller certified perturbation sizes in our CR inspired attacks. In contrast, a significant number of the perturbed edges in the base attacks connect with nodes with relatively larger certified perturbation sizes. Hence, under a fixed perturbation budget, our attacks can misclassify more nodes.

Comparing with other weight design strategies. Recall that our weight design is based on node’s certified robustness: nodes less provably robust to graph perturbations are assigned larger weights, in order to enlarge these node attack losses. Here, we consider three other possible strategies to design node weights that aim to *empirically* capture this property: 1) **Random**, where we uniformly assign node weights between $[0, 1]$ at random; 2) **Node degree**, where a node with a smaller degree might be less robust to graph perturbations, and we set a larger weight. Following our weight design, we set $w_{\text{deg}}(u) = \frac{1}{1+\exp(a \cdot \text{deg}(u))}$; 3) **Node centrality** [25], where a node with a smaller centrality might be less robust to graph perturbations, and we set a larger weight. Similarly, we set $w_{\text{cen}}(u) = \frac{1}{1+\exp(a \cdot \text{cen}(u))}$. As a baseline, we also consider no node weights.

Table 1 shows the attack results by applying these weight design strategies to the existing graph evasion and poisoning attacks. We have the following observations: 1) **Random** obtains the attack performance even worse than **No weight**’s. This indicates an inappropriate weight design can be harmful to the attack. 2) Node **Degree** and **Centrality** perform slightly better than **No weight**. One possible rea-

Table 1. Attack performance with different weight design.

Dataset	Method	CW-PGD	CE-PGD	Minmax	MetaTrain
Cora	No weight	0.74	0.71	0.62	0.68
	Random	0.77	0.75	0.65	0.72
	Degree	0.72	0.70	0.61	0.66
	Centrality	0.73	0.70	0.60	0.66
	Ours	0.70	0.66	0.55	0.62
Citeseer	No weight	0.64	0.63	0.63	0.61
	Random	0.66	0.66	0.68	0.64
	Degree	0.64	0.61	0.60	0.59
	Centrality	0.64	0.62	0.60	0.58
	Ours	0.60	0.60	0.57	0.52
B.C.Log	No weight	0.48	0.51	0.35	0.31
	Random	0.54	0.55	0.40	0.35
	Degree	0.46	0.50	0.32	0.28
	Centrality	0.47	0.49	0.32	0.27
	Ours	0.44	0.46	0.29	0.24

Table 2. Attack performance with different a .

Dataset	a	CR-CWPGD	CR-CEPGD	CR-Minmax	CR-MetaTrain
Cora	0.5	0.70	0.67	0.55	0.62
	1	0.70	0.66	0.55	0.62
	2	0.70	0.66	0.54	0.62
Citeseer	0.5	0.60	0.61	0.58	0.54
	1	0.60	0.60	0.57	0.52
	2	0.60	0.59	0.57	0.53
B.C.Log	0.5	0.44	0.47	0.31	0.25
	1	0.44	0.46	0.29	0.24
	2	0.44	0.46	0.29	0.24

son is that nodes with larger degree and centrality are empirically more robust to perturbations, which are also observed in previous works, e.g., [34, 50]. 3) Our weight design strategy performs the best. This is because our weight design *intrinsically* captures nodes’ certified robustness and thus yields more effective attacks.

Ablation study. In this experiment, we study the impact of hyperparameters: β in Equation 7, confidence level $1 - \alpha$ in Equation 9, and N in Equation 9, a in Equation 13, as well the running time vs. N . Figure 4 shows the results of β , $1 - \alpha$, and N , and running time vs. N on our attack. We observe that: 1) Our attack is not sensitive to β ; 2) Our attack slightly becomes worse as the confidence level $1 - \alpha$ increases. Such an observation can guide an attacker to set a relatively small β in practice. 3) Our attack becomes better as N increases, but already works well with a relatively smaller N . From this observation, an attacker can choose a small N in practice to save the time and cost when performing the attack. 4) Running time does not increase too much with N on the evasion attacks and is linear to N on poisoning attacks, consistent with our analysis in Section 4.3.

Table 2 shows the impact of a . We see the performances are stable across different a . This is largely because our weight design already ensures the node weight is inversely and exponentially to the node’s certified perturbation size.

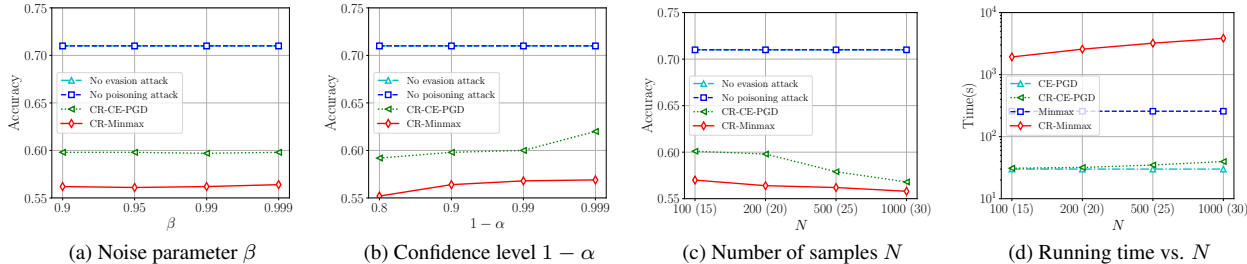


Figure 4. Impact of (a) β , (b) $1 - \alpha$, (c) N (# in bracket in x-axis is for poisoning attacks), and (d) running time vs. N on Citeseer. Note that (c): “No evasion attack” and “No poisoning attack” curves are overlapped; (d) $INT=10$ (2) for our evasion (poisoning) attacks.

6. Discussion

Evaluations on other GNNs. We mainly follow existing attacks [40, 51], which only evaluate GCN. Here, we also test SGC [38] on Cora and results show our CR-based GNN attacks also have a 6%-12% gain over the base attacks. This validates our strategy is generic to design better attacks.

Transferability between different GNNs. We evaluate the transferability of the graph perturbations generated by our 4 CR-based attacks on GCN to SGC on Cora, when the attack budget is 15. Accuracy on SGC under the 4 attacks are: 73%, 76%, 66%, and 67%, while accuracy on GCN are 71%, 73%, 63%, and 65%, respectively. This implies a promising transferability between GCN and SGC.

Defenses against our attacks. Almost all existing empirical defenses [10, 11, 13, 39, 45, 48, 49] are ineffective to adaptive attacks [24]. We adopt adversarial training [22], which is the only known effective empirical defense. Specifically, we first generate graph perturbations for target nodes via our attack and use the perturbed graph to retrain GNN with true node labels. The trained GNN is used for evaluation. We test on Cora and show this defense is effective to some extent, but has a nonnegligible utility loss. For instance, when budget=15, the accuracy under the CR-CW-PGD (CR-CE-PGD) attack increases from 73% (71%) to 76% (73%), but the normal accuracy reduces from 84% to 73% (72%).

7. Related Work

Attacking graph neural networks (GNNs). We classify the existing attacks to GNNs as evasion attacks [8, 20, 21, 23, 36, 39, 40, 50] and poisoning attacks [8, 19, 30, 31, 40, 46, 50, 51]. E.g., Xu et al. [40] proposed an untargeted PGD graph evasion attack to the GCN. The PGD attack leverages first-order optimization and generates discrete graph perturbations via convexly relaxing the binary graph structure, and obtains the state-of-the-art attack performance. Regarding graph poisoning attacks, Zugner et al. [51] proposed a graph poisoning attack, called Metattack, that perturbs the whole graph based on meta-learning. Our attack framework can be seamlessly plugged into these graph evasion and poisoning attacks and enhance their attack performance.

Attacking other graph-based methods. Besides attacking GNNs, other adversarial attacks against graph data include attacking graph-based clustering [6], graph-based col-

lective classification [32, 34], graph embedding [1, 4, 5, 9], community detection [18], graph matching [47], etc. For instance, Chen et al. [6] proposed a practical attack against spectral clustering, which is a well-known graph-based clustering method. Wang and Gong [34] designed an attack to the collective classification method, called linearized belief propagation, by modifying the graph structure.

Certified robustness and randomized smoothing. Randomized smoothing [7, 15–17, 42, 43] was the first method to obtain certified robustness of large models and achieved state-of-the-art performance. For instance, Cohen et al. [7] leveraged the Neyman-Pearson Lemma [26] to obtain a tight l_2 certified robustness for randomized smoothing with Gaussian noise on normally trained image models. Salman et al. [28] improved the certified robustness by combining the design of an adaptive attack against smoothed soft image classifiers and adversarial training on the attacked classifiers. [12], [35], and [2] applied randomized smoothing in the graph domain and derived certified robustness for community detection and node/graph classifications methods against graph perturbations. In this paper, we use randomized smoothing to design better attacks against GNNs.

8. Conclusion

We study graph evasion and poisoning attacks to GNNs and propose a novel attack framework motivated by certified robustness. We are the first work that uses certified robustness for an attack purpose. In particular, we first derive the node’s certified perturbation size, by extending randomized smoothing from the classifier perspective to a general function perspective. Based on it, we design certified robustness inspired node weights, which can be seamlessly plugged into the existing graph perturbation attacks’ loss and produce our certified robustness inspired attack loss and attack framework. Evaluations on multiple datasets demonstrate that existing attacks’ performance can be significantly enhanced by applying our attack framework.

Acknowledgments. This work was supported by Wang’s startup funding, the Cisco Research Award, and the National Science Foundation under grant No. 2216926. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the funding agencies.

References

- [1] Aleksandar Bojchevski and Stephan Günnemann. Adversarial attacks on node embeddings via graph poisoning. In *ICML*, 2019. 8
- [2] Aleksandar Bojchevski, Johannes Klicpera, and Stephan Günnemann. Efficient robustness certificates for discrete data: Sparsity-aware randomized smoothing for graphs, images and more. In *ICML*, 2020. 8
- [3] Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE S&P*, 2017. 6
- [4] Heng Chang, Yu Rong, Tingyang Xu, Wenbing Huang, Honglei Zhang, Peng Cui, Wenwu Zhu, and Junzhou Huang. A restricted black-box adversarial framework towards attacking graph embedding models. In *AAAI*, 2020. 8
- [5] Jinyin Chen, Yangyang Wu, Xuanheng Xu, Yixian Chen, Haibin Zheng, and Qi Xuan. Fast gradient attack on network embedding. *arXiv*, 2018. 8
- [6] Yizheng Chen, Yacin Nadjji, Athanasios Kountouras, and more. Practical attacks against graph-based clustering. In *CCS*, 2017. 8
- [7] Jeremy M Cohen, Elan Rosenfeld, and J Zico Kolter. Certified adversarial robustness via randomized smoothing. In *ICML*, 2019. 1, 3, 8
- [8] Hanjun Dai, Hui Li, Tian Tian, and more. Adversarial attack on graph structured data. In *ICML*, 2018. 1, 8
- [9] Quanyu Dai, Qiang Li, Jian Tang, and Dan Wang. Adversarial network embedding. In *AAAI*, 2018. 8
- [10] Negin Entezari, Saba A Al-Sayouri, Amirali Darvishzadeh, and Evangelos E Papalexakis. All you need is low (rank) defending against adversarial attacks on graphs. In *WSDM*, 2020. 8
- [11] Simon Geisler, Tobias Schmidt, Hakan Şirin, Daniel Zügner, Aleksandar Bojchevski, and Stephan Günnemann. Robustness of graph neural networks at scale. In *NeurIPS*, 2021. 8
- [12] Jinyuan Jia, Binghui Wang, Xiaoyu Cao, and Neil Gong. Certified robustness of community detection against adversarial structural perturbation via randomized smoothing. In *WWW*, 2020. 8
- [13] Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. Graph structure learning for robust graph neural networks. In *KDD*, 2020. 8
- [14] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*, 2017. 1, 6
- [15] Mathias Lecuyer, Vaggelis Atlidakis, Roxana Geambasu, Daniel Hsu, and Suman Jana. Certified robustness to adversarial examples with differential privacy. In *IEEE S & P*, 2019. 3, 8
- [16] Alexander Levine and Soheil Feizi. Robustness certificates for sparse adversarial attacks by randomized ablation. In *AAAI*, 2020. 8
- [17] Bai Li, Changyou Chen, Wenlin Wang, and Lawrence Carin. Certified adversarial robustness with additive noise. 2019. 3, 8
- [18] Jia Li, Honglei Zhang, Zhichao Han, and et al. Adversarial attack on community detection by hiding individuals. In *WWW*, 2020. 8
- [19] Xuanqing Liu, Si Si, Xiaojin Zhu, Yang Li, and Chojui Hsieh. A unified framework for data poisoning attack to graph-based semi-supervised learning. In *NeurIPS*, 2019. 1, 8
- [20] Jiaqi Ma, Shuangrui Ding, and Qiaozhu Mei. Towards more practical adversarial attacks on graph neural networks. In *NeurIPS*, 2020. 1, 8
- [21] Yao Ma, Suhang Wang, Tyler Derr, Lingfei Wu, and Jiliang Tang. Attacking graph convolutional networks via rewiring. *arXiv*, 2019. 8
- [22] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. In *ICLR*, 2018. 8
- [23] Jiaming Mu, Binghui Wang, Qi Li, Kun Sun, Mingwei Xu, and Zhuotao Liu. A hard label black-box adversarial attack to graph neural networks. In *CCS*, 2021. 1, 8
- [24] Felix Mujkanovic, Simon Geisler, Stephan Günnemann, and Aleksandar Bojchevski. Are defenses for graph neural networks robust? In *NeurIPS*, 2022. 8
- [25] Mark Newman. *Networks*. Oxford university press, 2018. 7
- [26] Jerzy Neyman and Egon Sharpe Pearson. IX. on the problem of the most efficient tests of statistical hypotheses. 1933. 8
- [27] Ryan Rossi and Nesreen Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015. 6
- [28] Hadi Salman, Jerry Li, Ilya Razenshteyn, and more. Provably robust deep learning via adversarially trained smoothed classifiers. In *NeurIPS*, 2019. 1, 8
- [29] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008. 6
- [30] Yiwei Sun, Suhang Wang, Xianfeng Tang, Tsung-Yu Hsieh, and Vasant Honavar. Adversarial attacks on

- graph neural networks via node injections: A hierarchical reinforcement learning approach. In *The Web Conference*, 2020. 1, 8
- [31] Tsubasa Takahashi. Indirect adversarial attacks via poisoning neighbors for graph convolutional networks. In *2019 IEEE International Conference on Big Data (Big Data)*, 2019. 8
- [32] MohamadAli Torkamani and Daniel Lowd. Convex adversarial collective classification. In *ICML*, 2013. 8
- [33] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *ICLR*, 2018. 1
- [34] Binghui Wang and Neil Gong. Attacking graph-based classification via manipulating the graph structure. In *CCS*, 2019. 1, 7, 8
- [35] Binghui Wang, Jinyuan Jia, Xiaoyu Cao, and Neil Gong. Certified robustness of graph neural networks against adversarial structural perturbation. In *KDD*, 2021. 1, 3, 4, 8, 11
- [36] Binghui Wang, Youqi Li, and Pan Zhou. Bandits for black-box attacks to graph neural networks with structure perturbation. In *CVPR*, 2022. 1, 8
- [37] Binghui Wang, Tianxiang Zhou, Minhua Lin, Pan Zhou, Ang Li, Meng Pang, Cai Fu, Hai Li, and Yiran Chen. Evasion attacks to graph neural networks via influence function. *arXiv*, 2020. 1
- [38] Felix Wu, Tianyi Zhang, Amauri Holanda de Souza Jr, Christopher Fifty, Tao Yu, and Kilian Q Weinberger. Simplifying graph convolutional networks. In *ICML*, 2019. 1, 8
- [39] Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. Adversarial examples on graph data: Deep insights into attack and defense. In *IJCAI*, 2019. 1, 8
- [40] Kaidi Xu, Hongge Chen, Sijia Liu, and more. Topology attack and defense for graph neural networks: An optimization perspective. In *IJCAI*, 2019. 1, 2, 3, 5, 6, 8
- [41] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *ICLR*, 2019. 1
- [42] Greg Yang, Tony Duan, J Edward Hu, Hadi Salman, Ilya Razenshteyn, and Jerry Li. Randomized smoothing of all shapes and sizes. In *ICML*, 2020. 8
- [43] Runtian Zhai, Chen Dan, Di He, Huan Zhang, Boqing Gong, Pradeep Ravikumar, Cho-Jui Hsieh, and Liwei Wang. Macer: Attack-free and scalable robust training via maximizing certified radius. In *ICLR*, 2020. 8
- [44] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. In *NIPS*, 2018. 1
- [45] Xiang Zhang and Marinka Zitnik. Gnn-guard: Defending graph neural networks against adversarial attacks. In *NeurIPS*, 2020. 8
- [46] Zaixi Zhang, Jinyuan Jia, Binghui Wang, and Neil Gong. Backdoor attacks to graph neural networks. 2021. 8
- [47] Zijie Zhang, Zeru Zhang, Yang Zhou, Yelong Shen, Ruoming Jin, and Dejing Dou. Adversarial attacks on deep graph matching. 2020. 8
- [48] Dingyuan Zhu, Ziwei Zhang, Peng Cui, and Wenwu Zhu. Robust graph convolutional networks against adversarial attacks. In *KDD*, 2019. 8
- [49] Jun Zhuang and Mohammad Al Hasan. Defending graph convolutional networks against dynamic graph perturbations via bayesian self-supervision. In *AAAI*, 2022. 8
- [50] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. Adversarial attacks on neural networks for graph data. In *KDD*, 2018. 1, 7, 8
- [51] Daniel Zügner and Stephan Günnemann. Adversarial attacks on graph neural networks via meta learning. In *ICLR*, 2019. 1, 2, 3, 6, 8