

Unsupervised Graph Poisoning Attack via Contrastive Loss Back-propagation

Sixiao Zhang[†]
zsx57575@gmail.com
University of Technology Sydney

Hongxu Chen^{*†}
hongxu.chen@uts.edu.au
University of Technology Sydney

Xiangguo Sun
sunxiangguo@seu.edu.cn
Southeast University

Yicong Li
Yicong.Li@student.uts.edu.au
University of Technology Sydney

Guandong Xu^{*}
guandong.xu@uts.edu.au
University of Technology Sydney

ABSTRACT

Graph contrastive learning is the state-of-the-art unsupervised graph representation learning framework and has shown comparable performance with supervised approaches. However, evaluating whether the graph contrastive learning is robust to adversarial attacks is still an open problem because most existing graph adversarial attacks are supervised models, which means they heavily rely on labels and can only be used to evaluate the graph contrastive learning in a specific scenario. For unsupervised graph representation methods such as graph contrastive learning, it is difficult to acquire labels in real-world scenarios, making traditional supervised graph attack methods difficult to be applied to test their robustness. In this paper, we propose a novel unsupervised gradient-based adversarial attack that does not rely on labels for graph contrastive learning. We compute the gradients of the adjacency matrices of the two views and flip the edges with gradient ascent to maximize the contrastive loss. In this way, we can fully use multiple views generated by the graph contrastive learning models and pick the most informative edges without knowing their labels, and therefore can promisingly support our model adapted to more kinds of downstream tasks. Extensive experiments show that our attack outperforms unsupervised baseline attacks and has comparable performance with supervised attacks in multiple downstream tasks including node classification and link prediction. We further show that our attack can be transferred to other graph representation models as well.

CCS CONCEPTS

• Information systems → Data mining.

^{*}Corresponding author.

[†]Both authors contributed equally to this research.

KEYWORDS

Graph Representation Learning, Graph Contrastive Learning, Adversarial Attack.

ACM Reference Format:

Sixiao Zhang, Hongxu Chen, Xiangguo Sun, Yicong Li, and Guandong Xu. 2022. Unsupervised Graph Poisoning Attack via Contrastive Loss Back-propagation. In *Proceedings of the ACM Web Conference 2022 (WWW '22)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3485447.3512179>

1 INTRODUCTION

Graph structured data can be commonly seen in our daily life, such as social networks [15, 34], biological networks [39], e-commercial networks [31], etc. Real-world graphs usually contain rich information and can be used for various tasks including recommendation, molecular structure classification, and community detection. To extract the information contained in graphs, researchers have been exploring various machine learning methods for graph learning, including DeepWalk [17], node2vec [10], matrix factorization [19], graph neural networks [20] and graph convolutional networks [14]. Graph representation learning plays an important role in these approaches, where the goal is to project nodes/graphs into a low-dimensional embedding space that preserves the structural and feature information [5, 6]. The quality of the embeddings directly influences the performance of downstream tasks such as link prediction, node/graph classification, community detection, etc.

However, recent studies have shown that graph learning models are vulnerable to adversarial attacks [12], where a small augmentation in the graph, such as adding/deleting edges/nodes and altering features, can lead to a large performance drop. The quality of the embeddings learned by graph representation models is also sensitive to such attacks [3, 33], and will in turn affects the performance of the downstream tasks. In addition, adversarial attacks in the graph domain are easy to conduct. For example, the attacker can attack a social network by simply creating some fake users and establishing connections with other people by following their accounts. The e-commercial networks can also be attacked by creating fake user profiles and writing fake reviews. Therefore, it is important to study the robustness of various graph learning models and develop models that are robust to adversarial attacks.

Recently, researchers have been exploring the robustness of supervised and semi-supervised graph representation models

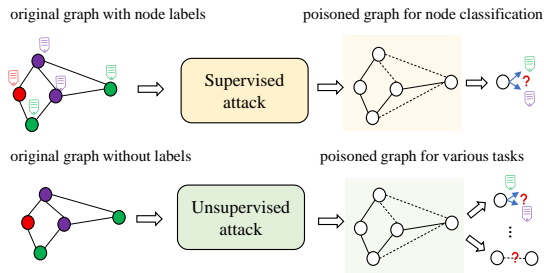


Figure 1: An illustration of the difference between supervised attacks and unsupervised attacks. Supervised attacks require labels and are targeted at certain downstream tasks, whereas unsupervised attacks do not require labels and are targeted at multiple downstream tasks.

[9, 13, 35, 36, 40], but the robustness of unsupervised graph representation models still remains an open challenge. Due to the complexity and scale-freeness of many real-world graphs, the labels are hard to acquire in most cases [25]. It is also impractical to do manual annotation for unlabeled large graphs. Thus, unsupervised graph representation learning becomes an important branch in graph representation learning. However, most state-of-the-art graph representation learning methods are supervised or semi-supervised [14, 24]. Traditional unsupervised methods, such as DeepWalk [17] and node2vec [10], have a relatively lower performance on downstream tasks compared with supervised methods. However, unsupervised graph representation learning has again attracted the attention of researchers in recent years because of the success of contrastive learning [18, 28, 30, 37, 38]. It has emerged as a new state-of-the-art unsupervised learning framework and has shown comparable performance with supervised and semi-supervised baselines. It generates different views of the original graph using stochastic augmentations and adopts a special contrastive loss to learn embeddings by comparing these views, which also makes graph contrastive learning models more robust to adversarial attacks compared with other graph representation methods [30]. However, such conclusions are obtained using existing graph adversarial attacks, which are mostly supervised attacks [9, 27, 40, 41], i.e. they need labels to conduct the attack as shown in Fig. 1. In real-world scenarios, it is difficult to attack contrastive learning with supervised attacks because the labels are hard to acquire. Therefore, this raises a new challenge that how to attack unsupervised graph representation learning such as contrastive learning without knowing the labels? Previous work by Bojchevski et al. [2] has shown that unsupervised graph learning methods based on random walks such as DeepWalk can be attacked in an unsupervised way by approximating the influence of a single augmentation in an optimization perspective. However, their method is limited to random walk and we will show in our experiments that it can not perform well on graph contrastive learning.

Attacking graph contrastive learning in an unsupervised way is a non-trivial task. Our goal is to poison the graph so that the overall performance on various downstream tasks is degraded. But we can’t use the downstream task to measure the quality of the

embeddings. Therefore, how to explicitly measure the quality of the embeddings is the first problem needed to be solved. In addition, the stochastic augmentation process provides robustness against adversarial attacks, so the attack needs to be insensitive to stochastic augmentations. Besides, a powerful graph attack model should let the poisoned graph be deterrent as much as possible, which means our generated poisoned graphs should be also transferable to the other graph representation models.

To solve the above challenges, we propose a novel unsupervised adversarial attack based on gradient ascent on graph contrastive learning for node embeddings, which is called Contrastive Loss Gradient Attack (CLGA). To the best of our knowledge, this is the first work aiming at attacking graph contrastive learning in an unsupervised manner. Specifically, we compute the gradient of the contrastive loss w.r.t. the adjacency matrix, and flip the edges with the largest gradients, causing the contrastive loss to be damaged. We show by extensive experiments that CLGA outperforms other existing unsupervised attack baselines and has comparable performance with some of the supervised attack methods. We also show that CLGA can be transferred to other graph representation models. To guarantee reproducibility, we open the source code at ¹. In summary, our contributions are as follows:

- We propose CLGA, a gradient-based unsupervised attack method targeting graph contrastive learning. Unlike most supervised attack models, CLGA does not rely on labels and can degrade the quality of the learned embeddings and thus affect the performance of various downstream tasks.
- We show by extensive experiments that CLGA outperforms unsupervised attack baselines and has comparable performance with some of the supervised attack methods on three benchmark datasets and on both node classification and link prediction tasks.
- We also show that CLGA can be transferred to other graph representation models such as GCN and DeepWalk.
- We visualize the learned embeddings to show how CLGA influences the quality of them.

In the following content, we first introduce the related work and preliminaries for graph adversarial attack and graph contrastive learning. Next we introduce how CLGA works. Moreover, we show the experiment results on three benchmark datasets and compare CLGA with several state-of-the-art baseline attacks. At last, we provide the conclusion of the paper.

2 RELATED WORK

2.1 Graph Adversarial Attack

The robustness of various machine learning models is a hot topic in the research community. A small perturbation in the training data might lead to a big change in the model performance. How to attack models and how to defend against such attacks are two main streams of this field. In the graph domain, people have already proposed many successful adversarial attack methods. RL-S2V [9] adopts reinforcement learning to attack graph neural networks by adding/deleting edges that have a positive effect on the attacker’s

¹<https://github.com/RinneSz/CLGA>

goal. Later, NIPA [21] and ReWatt [16] also use reinforcement learning to attack GNNs but only by injecting nodes and rewiring edges respectively. Some works try to model the attack as an optimization problem. PGD and MinMax [27] optimize the negative cross-entropy loss using the gradient. Nettack [40] iteratively selects augmentations by calculating the score/loss of each possible augmentation and chooses the one that maximizes the loss. Besides, because of the strong ability of gradients to locate the most important instances, gradient-based attack methods have emerged as new state-of-the-arts, such as Mettack [41] and FGA [7], where they both use the gradient of the classification loss w.r.t. the adjacency matrix to select the edges to augment. However, all the attack methods mentioned above are supervised attack methods, which means that they require labels to conduct the attack, either by using the ground-truth labels or by using the predictions of the target model as labels. They cannot be directly used to attack an unsupervised model that aims at learning pre-trained embeddings. Bojchevski et al. [2] have proposed an unsupervised optimization-based attack method to attack graph embedding methods based on random walks such as DeepWalk without using the labels. But we will show that it does not work very well for the state-of-the-art contrastive learning framework.

2.2 Graph Contrastive Learning

Recently, because of the great success of contrastive learning in computer vision [8, 11], researchers have begun to explore contrastive learning in the graph domain and have achieved competitive performance compared with supervised and semi-supervised models. DGI [25] maximizes the mutual information between patch representations and corresponding high-level summaries of graphs. CSSL [32] first defines four types of basic augmentation operations, edge deletion, node deletion, edge insertion and node insertion, wherein all of them choose nodes/edges randomly. Then they randomly sample a sequence of operations to apply to the graph. GraphCL [30] maximizes the agreement between two views by uniform perturbations including node dropping, edge perturbation, attribute masking and subgraphs. GCC [18] uses random walk on the r-ego network to generate different subgraphs, and learns embeddings by comparing these subgraphs. GRACE [37] learns node embeddings by randomly removing edges and masking features. GCA [38] further improves GRACE by designing new augmentation schemes that are aimed to keep important structures and attributes unchanged, while perturbing possibly unimportant links and features. These works differ in the specific augmentation strategy, such as adding/deleting edges/nodes, masking features, and extracting subgraphs. The loss functions used in these works share the same idea. They treat the embeddings of the same instance in different augmented views as positive pairs and treat the embeddings of different instances as negative pairs. Later, bootstrapped graph contrastive learning models have been proposed by Thakoor et al [22] and Che et al [4], which do not require negative pairs. Such a method refers to two neural networks, the online network and the target network. Two augmented views are the input to the two networks respectively. The online network is trained to predict the target network output, and the target network is updated by an exponential moving average of the online network. After training, the encoder of the online

network is used to compute the representations of the downstream tasks.

3 PRELIMINARIES

3.1 Graph Contrastive Learning on Node Embeddings

A graph G is defined as $G = (V, E)$, where V is the set of vertices (nodes) and E is the set of edges. It can also be represented by the adjacency matrix $A \in \mathbb{R}^{N \times N}$, where N is the number of nodes. Some graphs also have features $X \in \mathbb{R}^{N \times d}$ associated with nodes, where d is the number of features. The goal of graph contrastive learning is to learn an encoder $f(A, X)$ that outputs embeddings of each node. Such embeddings can then be used for downstream tasks including node classification and link prediction.

In this paper, we focus on node-level contrastive learning. So we will introduce how node-level graph contrastive learning works following the state-of-the-art GCA model [38]. A general graph contrastive learning framework consists of three steps. First, two stochastic augmentation processes t_1 and t_2 are applied to the original graph to obtain two different views. Typical augmentation strategies include edge dropping/insertion, feature masking, sub-graph extracting, etc. Second, the two views are fed into a shared encoder $f(A, X)$ to obtain node embeddings. Each node will have two embeddings corresponding to the two views. Third, a contrastive loss is applied to the embeddings to push positive node pairs close to each other and push negative node pairs away from each other. Specifically, the loss for the i -th node in the first view is

$$l(e_i^1, e_i^2) = -\log \frac{e^{\beta(e_i^1, e_i^2)/\tau}}{e^{\beta(e_i^1, e_i^2)/\tau} + \sum_{j \neq i} (e^{\beta(e_i^1, e_j^1)/\tau} + e^{\beta(e_i^1, e_j^2)/\tau})} \quad (1)$$

where e_i^1 and e_i^2 denote the embeddings of the i -th node in the first and second view respectively. β is a similarity function, e.g. cosine similarity. τ is a temperature parameter. Note that the above loss is non-symmetric for e_i^1 and e_i^2 . So we need to add it up with its counterpart $l(e_i^2, e_i^1)$. The final loss \mathcal{L} is the sum of l for all nodes, which is

$$\mathcal{L} = \sum_{i=1}^N l(e_i^1, e_i^2) + l(e_i^2, e_i^1) \quad (2)$$

By minimizing \mathcal{L} , the model treats the embeddings of the same node in the two views as positive pairs and the embeddings of different nodes in either the same view or different views as negative pairs. Because of the stochastic augmentation procedure, the learned embeddings are more robust to the small perturbations in the graph compared with other conventional graph representation models.

3.2 Graph Adversarial Attack

Adversarial attack aims to add noise to the data to degrade the performance of the target model. For graph data, the common noise includes edge-level, node-level, and graph-level augmentations, such as adding/deleting edges/nodes, augmenting features, and adding fake graphs for graph classification.

According to the goal of the attacker, adversarial attacks can be divided into two categories:

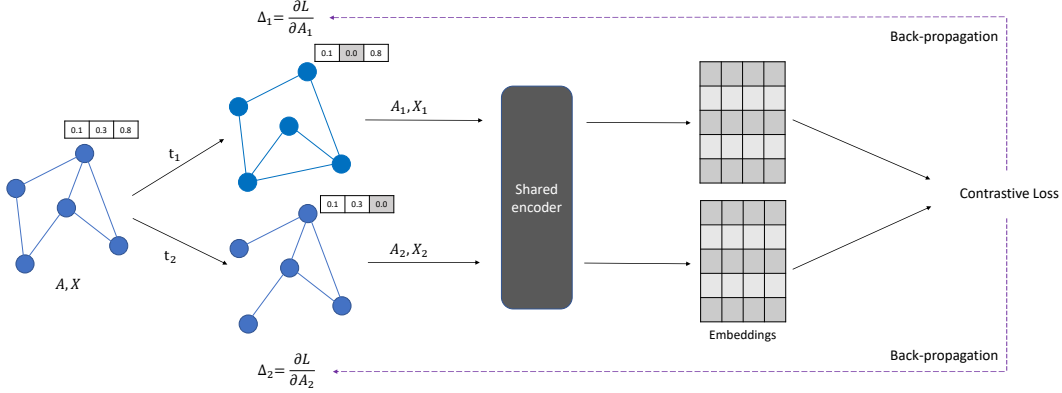


Figure 2: CLGA framework. The clean adjacency matrix A and feature matrix X are augmented by two stochastic augmentations t_1 and t_2 to obtain two views A_1, X_1 and A_2, X_2 . The two views are then fed into a shared encoder to obtain the node embeddings. The contrastive loss is computed using the embeddings. We then back-propagate the contrastive loss to obtain the gradients of the two views’ adjacency matrices. The gradients are used for selecting the edges that will be flipped in our attack.

- **Untargeted Attack:** force the model to have bad overall performance on all given instances.
- **Targeted Attack:** force the target model to have bad performance on a subset of instances.

A successful untargeted attack will make the target model produce a biased prediction for every query, and a targeted attack will make the target model produce a biased prediction for only certain queries. As a result, when testing the performance of the model on a test set, a successful untargeted attack will make the model have a worse overall accuracy compared with the model trained on clean data. A successful targeted attack will only make the model produce wrong predictions on a specific subset of the test set, while still producing correct predictions for other data, and thus the overall accuracy may not be significantly reduced.

Besides, according to the capacity of the attacker, adversarial attacks can be divided into two categories:

- **Poisoning Attack:** attack happens before the target model is trained. The training data of the target model is poisoned and is being used to train the target model.
- **Evasion Attack:** attack happens after the model is trained. The model is trained on clean training data and is fixed.

Poisoning attacks are a stronger attack since the attacker can affect the learning process of the target model. For evasion attacks, since the target model is fixed, what we can do is to perturb the data to cause misclassifications. For example, the attacker can modify the graph of a protein molecule to make it be misclassified as another type of protein.

For node representation learning on graphs, evasion attack is hard to conduct due to the transductive setting where all nodes are seen by the target model during training. Therefore, in this paper, we focus on the untargeted poisoning attack since this is best suited for attacking node representation learning models. Targeted poisoning attacks are a special case for untargeted poisoning attacks and can be achieved by restricting the target nodes.

3.3 Graph Convolutional Networks

The graph convolutional network proposed by Kipf et al. [14] is usually used as the encoder f in graph contrastive learning because of its state-of-the-art performance in modeling graph-structured data. Given the adjacency matrix A and feature matrix X , the output X' of a typical GCN layer is computed by

$$X' = \tilde{D}^{-1/2} \tilde{A} \tilde{D}^{-1/2} X \Theta \quad (3)$$

where $\tilde{A} = A + I$ denotes the adjacency matrix with self-loops. \tilde{D} is a diagonal degree matrix and $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$. Θ is the parameter matrix of the GCN layer. Activation functions such as ReLU and Sigmoid are applied at the output of each layer. In conventional GCN models, Θ is the only learnable parameter in Eq. 3 and is updated by gradient descent. However, to poison the graph, we need to know the gradient on the adjacency matrix A and luckily it is also easy to be deduced by Eq. 3. We will introduce how we attack graph contrastive learning using the gradient of the adjacency matrix in the next section.

4 CONTRASTIVE LOSS GRADIENT ATTACK

4.1 Overview

We want to design an untargeted poisoning attack, where the goal is to poison the graph data such that the overall quality of the embeddings learned by graph contrastive learning is degraded, which leads to worse performance in multiple downstream tasks. The contrastive loss used in contrastive learning is a natural measurement of the embedding quality. Therefore, we choose to poison the graph by maximizing the contrastive loss. In our attack, we only augment edges and do not augment features, because features are auxiliary information and not all the graphs have features. We formulate our problem as:

$$\begin{aligned} \max_{\hat{A}} \quad & \mathcal{L}(f_{\theta'}(A_1, X_1), f_{\theta'}(A_2, X_2)) \\ \text{s.t.} \quad & \theta' = \arg \min_{\theta} \mathcal{L}(f_{\theta}(A_1, X_1), f_{\theta}(A_2, X_2)), \\ & (A_1, X_1) = t_1(\hat{A}, X), (A_2, X_2) = t_2(\hat{A}, X), \|A - \hat{A}\| = \sigma. \end{aligned} \quad (4)$$

here \mathcal{L} is the contrastive loss stated in Eq. 2. f is the encoder and θ is the set of parameters of the encoder. A and X are the adjacency matrix and feature matrix of the clean graph. \hat{A} is the poisoned adjacency matrix. t_1 and t_2 are two stochastic augmentations. A_1 and X_1 are the adjacency matrix and feature matrix of the first view, and A_2 and X_2 are those of the second view. The last constraint indicates that the number of augmented edges is bounded by a given threshold σ .

This is a bi-level optimization problem and is hard to be solved directly. Inspired by [41], we propose to use meta-gradients, i.e. gradients w.r.t. the adjacency matrix. We back-propagate the contrastive loss to obtain the gradients of the adjacency matrix and update the adjacency matrix to maximize the loss.

4.2 Gradient based Attack

To solve Eq. 4, we propose to use gradient ascent on the adjacency matrix. The key idea is to flip the edges with the largest gradient values and the correct gradient directions, and thus the contrastive loss will be maximized. For example, if an observed edge (which means the corresponding entry is 1 in the adjacency matrix) has a negative gradient, then deleting it (from 1 to 0 in the adjacency matrix) will be very likely to increase the loss. Similarly, for an unobserved edge that has a large positive gradient, adding it to the graph will also increase the loss. An illustration of how we acquire the gradients is shown in Fig. 2. Specifically, if we use a differentiable encoder $f(A, X)$, e.g. graph convolutional network (GCN), we can easily obtain the gradients of the two views' adjacency matrix A_1 and A_2 :

$$\Delta_1 = \frac{\partial \mathcal{L}}{\partial A_1} = \frac{\partial \mathcal{L}}{\partial f(A_1, X_1)} \cdot \frac{\partial f(A_1, X_1)}{\partial A_1} \quad (5)$$

$$\Delta_2 = \frac{\partial \mathcal{L}}{\partial A_2} = \frac{\partial \mathcal{L}}{\partial f(A_2, X_2)} \cdot \frac{\partial f(A_2, X_2)}{\partial A_2} \quad (6)$$

Ideally, the most informative edges usually contribute largely to the graph learning model because their loss gradients on the clean adjacency matrix $\Delta = \frac{\partial \mathcal{L}}{\partial A}$ usually have larger absolute values. To find these edges, we need to differentiate through the stochastic augmentations t_1 and t_2 to obtain Δ as shown in Fig. 2. Unfortunately, the stochastic augmentation process t might contain some indifferentiable policy such as adding/deleting nodes and extracting subgraphs, making the problem even harder. To solve this problem, we do not back-propagate through t and only use Δ_1 and Δ_2 to help select which edges to flip. Another problem is that we cannot directly use Δ_1 or Δ_2 to select edges because the adjacency matrices of the two views are different from the clean adjacency matrix due to the stochastic augmentations. That means the edges having the largest gradients in these two views (Δ_1 and Δ_2) might not be the truly important ones in the original graph. With the above discussion, we can find that the core challenge is how to combine these two gradient matrices Δ_1 and Δ_2 to find the most informative edges and in the meanwhile, alleviate the bias caused by the stochastic augmentations (t_1 and t_2) as much as possible.

But what exactly is the form of the bias caused by the stochastic augmentations? We answer this question with a small example. As we know, conventional graph contrastive learning learns embeddings that are insensitive to the change of the graph by updating the parameters of the encoder. It tries to eliminate the difference

between the two views. If we assume that the adjacency matrices of the two views are learnable parameters, how will the adjacency matrices be like? Apparently they will become identical. Consider that if the two views are already identical, then every edge will have a zero gradient. If we now apply a stochastic augmentation and one single edge is changed, where it exists in the first view but does not exist in the second view, it will have a negative gradient in the first view and a positive gradient in the second view, pushing the corresponding entries 1 and 0 closer to each other. In this case, this augmented edge will have a much larger gradient compared with other edges, but can we say that it is more informative than other edges? Of course not, because the large gradient is actually caused by the stochastic augmentation. Instead, if we add the two gradients up, where one is positive and another is negative, the result will be close to 0 and we will thus acquire a relatively accurate measurement of how important this edge is by successfully alleviating the bias introduced by the stochastic augmentation.

With the above example, we propose to solve the problem by using two little tricks. First, we add Δ_1 and Δ_2 up to alleviate the bias introduced by the stochastic augmentations:

$$\Delta_0 = \Delta_1 + \Delta_2 \quad (7)$$

In addition, we add up gradients of K random stochastic augmentation pairs to further alleviate the bias caused by some rare cases:

$$\Delta' = \sum_{k=1}^K \Delta_0^k = \Delta_1^k + \Delta_2^k \quad (8)$$

where for each k , Δ_1^k and Δ_2^k is obtained by a random stochastic augmentation pair t_1^k and t_2^k . We use Δ' to select the edges that have the largest absolute gradients and correct gradient directions to flip.

Specifically, if for an edge, the gradients in the two views are both positive, it suggests that even if the stochastic augmentations are applied, the two views both prefer a smaller value in the corresponding entry which will lead to a smaller contrastive loss. Instead, increasing the value in this entry by adding this edge (if it does not exist) is very likely to increase the contrastive loss. This also applies to the case where the gradients are both negative. In our method, we add the two gradients to obtain a larger absolute value and this edge will be ranked at a higher position. If one of the gradients is positive and another is negative, it is usually caused by the stochastic augmentation on itself or its neighborhood. In this case, the information contained in its gradients is more about how to compensate for the heavily augmented neighborhood, or how to alleviate the changes caused by the stochastic augmentations. We are unable to tell how important this edge itself is w.r.t. the contrastive loss, so we hope that this edge will be ranked in a lower place when we are picking which edges to flip. We can achieve this by adding the two gradients up to get a new value that has a smaller absolute value than before, and it will thus have a lower ranking when we rank all the edges.

To further improve the attack, in each iteration we only pick one edge to flip. Specifically, in the first iteration, we train the contrastive model with the clean graph. Then we compute Δ' and select one edge with the largest gradient and correct direction. We flip the selected edge and use the new adjacency matrix (which

differs from the clean adjacency only on this single edge) to retrain the contrastive model in the next iteration. So if we are going to flip 100 edges, we will need to retrain the model 100 times. This iterative process helps us to better locate informative edges.

The overall algorithm is shown in algorithm 1.

Algorithm 1: CLGA

Input: Clean adjacency matrix A , feature matrix X , differentiable encoder f , stochastic augmentation set T , augmentation threshold σ , number of iterations K ;

Output: Poisoned graph \hat{A} ;

```

1:  $i = 0; \hat{A} = A;$ 
2: while  $i < \sigma$  do
3:   Train  $f$  with  $\hat{A}$  and  $X$ ;
4:    $\Delta' = 0;$ 
5:   for  $k = 1$  to  $K$  do
6:     Sample two stochastic augmentations  $t_1^k, t_2^k \in T$ ;
7:     Obtain two views  $(A_1^k, X_1^k) = t_1^k(\hat{A}, X)$ ,
        $(A_2^k, X_2^k) = t_2^k(\hat{A}, X)$ ;
8:     Forward propagate  $(A_1^k, X_1^k), (A_2^k, X_2^k)$  through  $f$  and
       compute contrastive loss  $\mathcal{L}$ ;
9:     Obtain the gradients of  $A_1^k$  and  $A_2^k$  w.r.t. the contrastive
       loss,  $\Delta_1^k = \frac{\partial \mathcal{L}}{\partial A_1^k}, \Delta_2^k = \frac{\partial \mathcal{L}}{\partial A_2^k}$ ;
10:     $\Delta' = \Delta' + \Delta_1^k + \Delta_2^k;$ 
11:   end for
12:   Flip one edge with both the largest absolute gradient in  $\Delta'$ 
       and the correct direction, i.e., if the index of the edge is
        $[m, n]$ , then it should satisfy either  $\hat{A}[m, n] = 1$ ,
        $\Delta'[m, n] < 0$  or  $\hat{A}[m, n] = 0, \Delta'[m, n] > 0$ ;
13:    $\hat{A}[m, n] = 1 - \hat{A}[m, n];$ 
14:   Freeze the chosen edge and avoid being flipped again in
       next iterations;
15:    $i = i + 1;$ 
16: end while

```

4.3 Complexity Analysis

4.3.1 Time Complexity. The time complexity of CLGA itself is low since we only need to rank the gradients and select the desired edges. What is truly time-consuming is the training of the target model, because we need to retrain it in each iteration. Therefore, for complex models whose time complexity is large, it will take a relatively long time for the poisoned graph to be generated. However, we can alleviate this issue by doing some approximations at the price of attack performance. For example, we can select multiple edges in each iteration, or we can train the target model with pre-training or early-stopping because we do not necessarily need the target model to be fully converged as long as the order of the ranking is correct.

4.3.2 Space Complexity. Since we are computing the gradients of the whole adjacency matrix, the memory requirement is $O(N^2)$, where N is the number of nodes. This is expensive for graphs with a large number of nodes. However, we can also alleviate this issue by only considering a subset of nodes and edges. For real-world scale-free graphs, it is impractical for the attackers to be able to

Dataset	# Nodes	# Edges	# Features	# Classes
Cora	2708	5278	1433	7
CiteSeer	3327	4552	3703	6
PolBlogs	1490	16715	None	2

Table 1: Dataset statistics.

modify the whole graph, where they can only manipulate a small subset of nodes and edges. In this case, we only need to compute the gradients of the subset, which only has a small memory cost.

5 EXPERIMENTS

In our experiments, we aim to answer the following research questions:

- Can CLGA successfully degrade the performance of various downstream tasks of graph contrastive learning?
- Does CLGA outperform other unsupervised and supervised attacks?
- Is the poisoned graph generated by CLGA able to degrade the performance of other graph representation models?

To answer these questions, we compare CLGA with five state-of-the-art representative graph untargeted poisoning attacks on three benchmark datasets. We evaluate the quality of the embeddings by two downstream tasks: node classification and link prediction. We further show the transferability of CLGA by evaluating other graph representation models on the obtained poisoned graph.

5.1 Setup

5.1.1 Datasets. We use three popular public available benchmark datasets. The two of them are citation networks Cora and CiteSeer from Yang et al. [29]. We also use PolBlog dataset from Adamic et al. [1], which is a graph of political blogs. There are features associated with each node in Cora and CiteSeer. However, there are no features for PolBlog. The statistics of the three datasets are shown in Table 1.

5.1.2 Baselines. We compare CLGA with five baseline untargeted poisoning attacks, including PGD [27], DICE [26], MinMax [27], Metattack [41], and the unsupervised node embedding attack proposed by Bojchevski et al. [2]. Among these baselines, only Bojchevski et al. [2] is unsupervised and does not need labels, which is the same as our method. Since the labels are used as extra knowledge in the other four supervised attacks, they are expected to have a better performance compared with unsupervised ones. But we will show that CLGA has comparable performance and can even outperform some of the supervised baselines in some cases.

5.1.3 Experimental Settings. For Metattack, MinMax and PGD, we use a 2-layer GCN as the surrogate model to acquire the poisoned graphs. For all the attack methods, we first generate the poisoned graphs and then feed them to the state-of-the-art contrastive learning framework GCA proposed by Zhu et al. [38]. Because a 2-layer GCN is used as the encoder in GCA and PolBlog does not have features, we randomly initialize a 32-dimensional feature for each node in PolBlog. We set the augmentation threshold σ to be 1%/5%/10% of the number of edges, i.e. we are allowed to modify at most 1%/5%/10% edges. For instance, if the number of edges in the clean

	Attack	Cora			CiteSeer			PolBlog		
		1%	5%	10%	1%	5%	10%	1%	5%	10%
Supervised	Metattack	0.7586	0.6928	0.6168	0.5920	0.3986	0.2952	0.8208	0.8039	0.8011
	PGD	0.7680	0.7592	0.7402	0.6098	0.6198	0.6056	0.8100	0.8010	0.7987
	MinMax	0.7624	0.7218	0.6174	0.6302	0.5254	0.5618	0.8016	0.7913	0.7986
	DICE	0.7712	0.7642	0.7240	0.6256	0.5774	0.5246	0.8107	0.7847	0.7394
Unsupervised	Bojchevski et al. [2]	0.7490	0.7710	0.7670	0.6442	0.6448	0.6608	0.8187	0.8042	0.7892
	CLGA	0.7316	0.7188	0.6814	0.6368	0.5906	0.5368	0.8088	0.7944	0.7726

Table 2: Node classification accuracy of logistic regression model trained after GCA. 1%/5%/10% denote the maximum number of edges allowed to be augmented. The boldfaced ones are the best in either supervised or unsupervised approaches.

	Attack	Cora			CiteSeer			PolBlog		
		1%	5%	10%	1%	5%	10%	1%	5%	10%
Supervised	Metattack	0.9010	0.8733	0.8500	0.9109	0.8853	0.8544	0.8617	0.8585	0.8635
	PGD	0.9143	0.9073	0.9073	0.9169	0.9248	0.9057	0.8605	0.8584	0.8625
	MinMax	0.9116	0.9004	0.8944	0.9145	<u>0.8890</u>	0.8981	0.9145	0.8890	0.8981
	DICE	0.9046	0.8828	0.8593	0.9137	0.8918	0.8679	0.8551	0.8450	0.8352
Unsupervised	Bojchevski et al. [2]	0.9164	0.9099	0.9101	0.9239	0.9168	0.9196	0.8593	<u>0.8543</u>	0.8587
	CLGA	<u>0.9012</u>	<u>0.8741</u>	0.8420	<u>0.9114</u>	0.8911	<u>0.8610</u>	<u>0.8584</u>	0.8598	<u>0.8563</u>

Table 3: Link prediction AUC of the MLP trained after GCA. 1%/5%/10% denote the maximum number of edges allowed to be augmented. The boldfaced ones are the best and the underlined ones are the second best among all approaches.

graph is 1000, then under 1% threshold, we can add or delete at most 10 edges. For a fair comparison, we fix the hyperparameters of GCA across experiments as suggested by Zhu et al. [38], including the temperature τ and stochastic augmentation rates, where $\tau = 0.4$ and the edge dropping rates are 0.3 and 0.4 and the feature dropping rates are 0.1 and 0.0 for two views respectively. We use Adam optimizer and set the learning rate to be 0.01.

For the downstream node classification task, we train a simple logistic regression model using the learned embeddings and report the classification accuracy. When training, we use the public split from [29] for Cora and CiteSeer. For PolBlog, we randomly split the nodes into 10%/10%/80% train/test/val set. We run each experiment 10 times and report the average.

For downstream link prediction task, we use a 2-layer MLP as the projection head to project the learned embeddings into a new latent space. We train the MLP with negative sampling and margin loss, i.e.,

$$l = - \sum_{i,j,k} \log(\sigma(\beta(e_i, e_j) - \beta(e_i, e_k))) \quad (9)$$

where e_i is the projected embedding of node i , β is the cosine similarity function, σ is the sigmoid function. Node i and node j is a positive pair, e.g. they are linked by an observed edge. Node i and node k is a randomly sampled negative pair, e.g. they are not linked in the training set. For all three datasets, we split the edges into 70%/20%/10% train/test/val set, and only the edges in the training set are used to train the contrastive model. We report the area under curve (AUC) score. We also run each experiment 10 times and report the average.

5.2 Comparing Node Classification Performance

Table 2 shows the accuracy of the downstream logistic regression model trained on the embeddings learned by GCA under each attack.

We can observe that, CLGA significantly outperforms the unsupervised baseline Bojchevski et al. [2], and can even outperform some of the supervised baselines in certain situations. For example, on Cora, CLGA outperforms PGD and DICE in all cases, and outperforms Metattack and MinMax at 1% and 1%/5% augmentation rates respectively.

As the augmentation rate increases from 1% to 10%, the node classification accuracy continuously drops under Metattack, DICE, and CLGA. However, for the other three attacks, a larger augmentation rate does not guarantee a lower node classification accuracy. This suggests that PGD, MinMax, and Bojchevski et al. [2] fail to capture the edges that are truly important to graph contrastive learning framework, or in other words, graph contrastive learning is robust to such attacks.

We can also observe that Metattack has the best overall performance and is very effective on Cora and CiteSeer. This is because that Metattack is originally designed for attacking an end-to-end GCN model for node classification. By utilizing the labels and the gradients, it can accurately locate the most informative edges that affect the node classification accuracy. However, such an approach is not guaranteed to work well on other downstream tasks, such as link prediction. Instead, we will show that our CLGA has a good performance not only on node prediction, but also on link prediction.

5.3 Comparing Link Prediction Performance

Table 3 shows the AUC scores of the three datasets under different attacks. We have several observations and conclusions regarding the results.

First, CLGA outperforms Bojchevski et al. [2] except for 5% augmentation rate on PolBlog, showing that CLGA is more effective than Bojchevski et al. [2]. Second, CLGA has comparable performance with supervised attacks. In most cases, CLGA is the second best among all six attacks. This suggests that, even without labels,

Attack	Node Classification						Link Prediction					
	DeepWalk			GCN			DeepWalk			GCN		
	1%	5%	10%	1%	5%	10%	1%	5%	10%	1%	5%	10%
Metattack	<u>0.6766</u>	0.6370	0.5554	0.7830	0.7300	0.6330	0.8777	<u>0.8629</u>	0.8319	0.8655	0.8378	<u>0.8182</u>
Bojchevski et al. [2]	0.6826	<u>0.6436</u>	<u>0.6276</u>	0.8240	0.7900	0.7950	0.8852	0.8873	0.8810	0.8675	0.8634	0.8542
CLGA	0.6720	0.6459	0.6376	<u>0.7860</u>	<u>0.7770</u>	<u>0.7690</u>	<u>0.8799</u>	0.8612	<u>0.8390</u>	<u>0.8667</u>	<u>0.8546</u>	0.8146

Table 4: Transferability of attacks on Cora for DeepWalk and GCN. We report the classification accuracy for node classification and AUC score for link prediction. 1%/5%/10% denote the maximum number of edges allowed to be augmented. The boldfaced ones are the best and the underlined ones are the second best among the three attacks.




Figure 3: Visualization of node embeddings learned by GCA under different attacks on Cora. Different Colors represent different classes of nodes.

CLGA can effectively reduce the downstream link prediction performance for graph contrastive learning and achieve comparable or even better performance than supervised attacks.

Together with the observations in the node classification scenario, our unsupervised attack CLGA is shown to be able to reduce both the downstream node classification performance and link prediction performance of graph contrastive learning, and has comparable performance with state-of-the-art supervised untargeted poisoning attacks.

5.4 Transferability

Although CLGA has a good performance on graph contrastive learning, it is still unknown how it works for other graph representation models. Here we further compare the transferability of CLGA with Bojchevski et al. [2] and Metattack. We use the poisoned graphs generated by the three attacks to train DeepWalk and a 2-layer GCN. We evaluate the quality of the embeddings on both node classification and link prediction. For node classification, the learned embeddings are used to train a logistic regression model. For link prediction, we use a 2-layer MLP as a projection head and train it with the loss in Eq. 9. We report the accuracy for node classification and AUC score for link prediction in Table 4. Other experimental settings are kept consistent with previous experiments. Note that even though the attack proposed by Bojchevski et al. [2] is originally designed for methods based on random walks such as DeepWalk, our CLGA can still achieve comparable performance with it on DeepWalk for the node classification task, and even outperforms it on the link prediction task. Moreover, CLGA outperforms Bojchevski et al. [2] on GCN. And CLGA even has comparable performance with Metattack at 1% augmentation rate for node classification, and at 1%/10% augmentation rates for link prediction. These observations suggest that CLGA can be transferred to other graph representation models.

5.5 Visualization

We visualize the learned embeddings by GCA under Metattack, Bojchevski et al. [2], and CLGA, to show how our attack influences the distribution of embeddings. The t-SNE [23] results are shown in Fig. 3. We can observe that the distributions of embeddings by Metattack and CLGA are denser than Bojchevski et al. [2] and clean embeddings, especially in the center, which is considered to be the main reason why Metattack and CLGA outperform others. Nodes in the center are hard to be well classified because such nodes have a similar distance to each cluster and we can't tell which class it belongs to with high confidence. This suggests that CLGA can reduce the quality of the embeddings by pushing embeddings to the center in Fig. 3, or in other words, pushing them close to the intersections/decision boundaries of the clusters, which makes the embeddings harder to be classified.

6 CONCLUSION

In this paper, we introduce Contrastive Loss Gradient Attack (CLGA), an unsupervised untargeted poisoning attack for attacking graph contrastive learning. This is the first work to attack graph contrastive learning in an unsupervised manner without using labels. The quality of the learned embeddings is damaged and the performance of various downstream tasks is degraded. We show by extensive experiments that our CLGA outperforms unsupervised baselines and has comparable and even better performance with supervised baselines. We also show that CLGA can be transferred to other graph representation models such as DeepWalk and GCN.

ACKNOWLEDGMENTS

The work has been supported by Australian Research Council under grants DP220103717, DP200101374, LP170100891, and LE220100078.

REFERENCES

- [1] Lada A Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*. 36–43.
- [2] Aleksandar Bojchevski and Stephan Günnemann. 2019. Adversarial attacks on node embeddings via graph poisoning. In *International Conference on Machine Learning*. PMLR, 695–704.
- [3] Heng Chang, Yu Rong, Tingyang Xu, Wenbing Huang, Honglei Zhang, Peng Cui, Wenwu Zhu, and Junzhou Huang. 2019. The General Black-box Attack Method for Graph Neural Networks. *arXiv preprint arXiv:1908.01297* (2019).
- [4] Feihu Che, Guohua Yang, Dawei Zhang, Jianhua Tao, Pengpeng Shao, and Tong Liu. 2020. Self-supervised Graph Representation Learning via Bootstrapping. *arXiv preprint arXiv:2011.05126* (2020).
- [5] Hongxu Chen, Hongzhi Yin, Tong Chen, Quoc Viet Hung Nguyen, Wen-Chih Peng, and Xue Li. 2019. Exploiting centrality information with graph convolutions for network representation learning. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 590–601.
- [6] Hongxu Chen, Hongzhi Yin, Weiqing Wang, Hao Wang, Quoc Viet Hung Nguyen, and Xue Li. 2018. PME: projected metric embedding on heterogeneous networks for link prediction. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1177–1186.
- [7] Jinyin Chen, Yangyang Wu, Xuanheng Xu, Yixian Chen, Haibin Zheng, and Qi Xuan. 2018. Fast gradient attack on network embedding. *arXiv preprint arXiv:1809.02797* (2018).
- [8] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PMLR, 1597–1607.
- [9] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial attack on graph structured data. In *International conference on machine learning*. PMLR, 1115–1124.
- [10] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. 855–864.
- [11] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9729–9738.
- [12] Wei Jin, Yaxin Li, Han Xu, Yiqi Wang, and Jiliang Tang. 2020. Adversarial attacks and defenses on graphs: A review and empirical study. *arXiv preprint arXiv:2003.00653* (2020).
- [13] Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. 2020. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 66–74.
- [14] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proceedings of ICLR*.
- [15] Hao Ma, Haixuan Yang, Michael R Lyu, and Irwin King. 2008. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings of the 17th ACM conference on Information and knowledge management*. 931–940.
- [16] Yao Ma, Suhang Wang, Tyler Derr, Lingfei Wu, and Jiliang Tang. 2019. Attacking graph convolutional networks via rewiring. *arXiv preprint arXiv:1906.03750* (2019).
- [17] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 701–710.
- [18] Jiezhong Qiu, Qibin Chen, Yuxiao Dong, Jing Zhang, Hongxia Yang, Ming Ding, Kuansan Wang, and Jie Tang. 2020. Gcc: Graph contrastive coding for graph neural network pre-training. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1150–1160.
- [19] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).
- [20] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* 20, 1 (2008), 61–80.
- [21] Yiwei Sun, Suhang Wang, Xianfeng Tang, Tsung-Yu Hsieh, and Vasant Honavar. 2019. Node injection attacks on graphs via reinforcement learning. *arXiv preprint arXiv:1909.06543* (2019).
- [22] Shantanu Thakoor, Corentin Tallec, Mohammad Gheshlaghi Azar, Rémi Munos, Petar Veličković, and Michal Valko. 2021. Bootstrapped Representation Learning on Graphs. *arXiv preprint arXiv:2102.06514* (2021).
- [23] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [24] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. In *Proceedings of ICLR*.
- [25] Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. 2019. Deep graph infomax. In *Proceedings of ICLR*.
- [26] Marcin Wanek, Tomasz P Michalak, Michael J Wooldridge, and Talal Rahwan. 2018. Hiding individuals and communities in a social network. *Nature Human Behaviour* 2, 2 (2018), 139–147.
- [27] Kaidi Xu, Hongge Chen, Sijia Liu, Pin-Yu Chen, Tsui-Wei Weng, Mingyi Hong, and Xue Lin. 2019. Topology attack and defense for graph neural networks: An optimization perspective. *arXiv preprint arXiv:1906.04214* (2019).
- [28] Haoran Yang, Hongxu Chen, Lin Li, Philip S Yu, and Guandong Xu. 2021. Hyper Meta-Path Contrastive Learning for Multi-Behavior Recommendation. *arXiv preprint arXiv:2109.02859* (2021).
- [29] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. 2016. Revisiting semi-supervised learning with graph embeddings. In *International conference on machine learning*. PMLR, 40–48.
- [30] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. 2020. Graph contrastive learning with augmentations. *Advances in Neural Information Processing Systems* 33 (2020).
- [31] Hanqing Zeng, Hongkuan Zhou, Ajitesh Srivastava, Rajgopal Kannan, and Viktor Prasanna. 2019. Graphsaint: Graph sampling based inductive learning method. *arXiv preprint arXiv:1907.04931* (2019).
- [32] Jiaqi Zeng and Pengtao Xie. 2020. Contrastive self-supervised learning for graph classification. *arXiv preprint arXiv:2009.05923* (2020).
- [33] Hengtong Zhang, Tianhang Zheng, Jing Gao, Chenglin Miao, Lu Su, Yaliang Li, and Kui Ren. 2019. Towards data poisoning attack against knowledge graph embedding.
- [34] Sixiao Zhang, Hongxu Chen, Xiao Ming, Lizhen Cui, Hongzhi Yin, and Guandong Xu. 2021. Where are we in embedding spaces? A Comprehensive Analysis on Network Embedding Approaches for Recommender Systems. *arXiv preprint arXiv:2105.08908* (2021).
- [35] Shijie Zhang, Hongzhi Yin, Tong Chen, Zi Huang, Lizhen Cui, and Xiangliang Zhang. 2021. Graph Embedding for Recommendation against Attribute Inference Attacks. In *Proceedings of the Web Conference 2021*. 3002–3014.
- [36] Dingyuan Zhu, Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2019. Robust graph convolutional networks against adversarial attacks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1399–1407.
- [37] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2020. Deep graph contrastive representation learning. *arXiv preprint arXiv:2006.04131* (2020).
- [38] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2020. Graph Contrastive Learning with Adaptive Augmentation. *arXiv preprint arXiv:2010.14945* (2020).
- [39] Marinka Zitnik and Jure Leskovec. 2017. Predicting multicellular function through multi-layer tissue networks. *Bioinformatics* 33, 14 (2017), i190–i198.
- [40] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2847–2856.
- [41] Daniel Zügner and Stephan Günnemann. 2019. Adversarial attacks on graph neural networks via meta learning. *arXiv preprint arXiv:1902.08412* (2019).