

Talos: A More Effective and Efficient Adversarial Defense for GNN Models Based on the Global Homophily of Graphs

Duanyu Li, Huijun Wu*, Min Xie*, Xugang Wu, Zhenwei Wu and Wenzhe Zhang

School of Computer Science, National University of Defense Technology, China

Abstract. Graph neural network (GNN) models play a pivotal role in numerous tasks involving graph-related data analysis. Despite their efficacy, similar to other deep learning models, GNNs are susceptible to adversarial attacks. Even minor perturbations in graph data can induce substantial alterations in model predictions. While existing research has explored various adversarial defense techniques for GNNs, the challenge of defending against adversarial attacks on real-world scale graph data remains largely unresolved. On one hand, methods reliant on graph purification and preprocessing tend to excessively emphasize local graph information, leading to sub-optimal defensive outcomes. On the other hand, approaches rooted in graph structure learning entail significant time overheads, rendering them impractical for large-scale graphs. In this paper, we propose a new defense method named Talos, which enhances the global, rather than local, homophily of graphs as a defense. Experiments show that the proposed approach notably outperforms state-of-the-art defense approaches, while imposing little computational overhead.

1 Introduction

Graph Neural Networks (GNNs) extend deep learning to graph data and significantly outperform traditional methods in tasks like node classification[9], graph classification[16], and link prediction[24]. However, like deep learning models in computer vision [7] and natural language processing [1], GNNs are vulnerable to adversarial attacks, where small perturbations mislead the model into making incorrect predictions[8, 19, 23]. This lack of robustness makes GNNs challenging to apply in critical real-world scenarios.

Recent studies propose various defense methods against these attacks, often leveraging homophily in their design. GNNs typically aggregate similar neighbors to improve node representations[17, 27], effectively utilizing previously overlooked graph structure information. Adversarial attacks disrupt homophily by adding dissimilar edges or altering node features[17, 12, 14]. Defense methods use the homophily assumption differently; for example, GCN-Jaccard removes edges between dissimilar nodes, while Soft-Median/Medoid[6, 5] aggregates the median and medoid of neighboring features to mitigate adversarial effects. Pro-GNN[12] incorporates homophily constraints within its model architecture by penalizing edges between dissimilar nodes.

Although these defenses show promise, recent studies reveal they can be vulnerable to adaptive attacks[13]. We argue that this is partly because they focus on the local neighborhood of nodes, ignoring that graph homophily is a global concept involving multi-hop neighbors. To enhance GNN robustness, we should improve the overall homophily of graphs. Additionally, real-world graphs are large, and many existing defenses introduce significant computational overhead, making them impractical. For example, GNNGuard's[25] overhead increases with the number of node features, and GCN-SVD[3] and ProGNN[12] rely on expensive SVD decomposition. Efficiency should be prioritized in defense design.

We propose a more effective and efficient adversarial defense method, Talos, which enhances global homophily to defend against attacks. Talos preprocesses the graph only once and is computationally efficient. Our contributions are:

- Expanding defense to include multi-hop neighbors, ensuring comprehensive protection.
- Providing theoretical derivations to enhance the method's validity and effectiveness.
- Demonstrating that Talos is more effective and faster than state-of-the-art defenses.

The rest of the paper is organized as follows: Section 2 covers related work, Section 3 presents the methodology, Section 4 shows experimental results, and Section 5 concludes the paper.

2 Related Work

2.1 Graph Adversarial Attacks

Due to the discrete features and transductive learning settings, generating adversarial perturbations for GNN models presents unique challenges. To address these, Dai et al.[2] use reinforcement learning to craft generalized attacks. Nettack[29] uses a linear model as a surrogate, bypassing the non-linear components of GNNs. IG-FSGM and IG-JSMA[17] use integrated gradients to estimate gradients in discrete states, enabling precise measurement of perturbation impacts. Additionally, PGDAttackPGDAttack[20] optimizes attack strategies using gradient descent, disrupting GNNs by manipulating edges. However, these approaches primarily focus on local attacks targeting specific nodes. Metaattack[28] employs meta-learning to perform poisoning attacks during the training process, reducing accuracy in tasks like classification and clustering. Most existing methods are gradient-based, making them difficult to apply to large-scale

* Corresponding authors: Huijun Wu (wuhuijun@nudt.edu.cn) and Min Xie (xiemin@nudt.edu.cn)

graphs. PR-BCD[6] proposes a sparsity-aware attack method that modifies edges based on Randomized Block Coordinate Descent (R-BCD).

2.2 Graph Adversarial Defense and Challenges

To improve the robustness of GNN models, defense methods are actively explored. Wu et al.[17] found that perturbing the graph structure has a more significant impact than tampering with the features. Attacks tend to add edges rather than delete them. They also revealed that edges are mostly added between dissimilar nodes. Based on such observations, they proposed GCN-Jaccard, which utilizes Jaccard similarities to identify and eliminate the perturbed edges. This effectively purifies the graph data and enhances its robustness against attacks. Some other methods use a similar idea to filter out dissimilar edges. For example, SoftMedian [6] and SoftMedoid[5] take advantage of recent advancements in differentiable sorting to design robust aggregation functions in GNN models. As mid-frequency signals partly correspond to moderately similar edges, Mid-GCN[11] preserves the mid-frequency signals while abandoning the high/low-frequency signals to defend against adversarial attacks. GNNGuard[25] employs an attention mechanism to assign higher weights to edges between similar nodes. GRAND[4] uses random feature augmentations together with neighborhood augmentation to avoid the over-reliance on directly connected neighbors. However, the above methods focus solely on the local neighborhood of nodes, neglecting to consider the overall structure of graphs. This could potentially expose them to adaptive attacks by adversaries[13].as attackers can allocate more attacking budget to further neighbors of nodes to achieve adaptive attacks.

In addition to removing adversarial neighbors, some defense methods modify the model architecture. GCN-SVD[3] purifies the graph by replacing the perturbed graph with its low-rank approximation. Pro-GNN[12] adds a regularization term during training to generate a low-rank and sparse graph. Since most adversarial attacks target graph structures, Wu et al.[18] propose a co-training mechanism that uses different models to fully utilize both node features and graph structures. However, these methods often involve time-consuming processes, making them inefficient for handling larger, more complex real-world graphs. Moreover, existing adversarial defenses heavily rely on empirical observations and lack theoretical analysis, reducing their reliability in real-world scenarios.

We propose Talos, which considers the entire graph structure. Talos preprocesses the graph by enhancing homophily while leaving model training unchanged, making it more efficient and suitable for large-scale graphs.

3 Methodology

This section presents the design of Talos. The fundamental principle of Talos involves harnessing global homophily information to promptly identify the edge most likely added by the attack and eliminate it. We will provide a comprehensive explanation and derivation of the Talos approach.

3.1 Notations

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ be a graph, where \mathcal{V} represents the set of nodes containing n nodes $\{v_1, \dots, v_n\}$ with $|\mathcal{V}| = n$. \mathcal{E} denotes the set of edges, typically represented by an adjacency matrix $\mathbf{A} \in \{0, 1\}^{n \times n}$.

An entry $\mathbf{A}_{ij} = 1$ indicates the presence of an edge $(v_i, v_j) \in \mathcal{E}$, while 0 indicates its absence.

Additionally, we have the feature matrix \mathbf{X} , where the features of each node $v \in \mathcal{V}$ are represented as a d -dimensional feature vector $\mathbf{x}_v \in \mathbb{R}^d$. The feature matrix \mathbf{X} comprises these feature vectors, specifically $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)^T \in \mathbb{R}^{n \times m}$, where m denotes the number of features.

In the context of graph node classification tasks, each node is associated with a label y , and $\mathcal{Y}^{|\mathcal{V}|}$ represents the set of true labels for all nodes. The goal of Talos is to enhance the model's prediction capability after an attack.

3.2 Effectiveness of High-order Neighbor Information

Existing adversarial attacks, whether approximation-based or gradient-based, can perturb arbitrary edges in the graph. However, defenses like GCN-Jaccard often focus only on local neighborhood information, making them vulnerable to adaptive attacks that exploit this limitation. As shown in Figure 1, an attacker can carefully craft a perturbation on node A without directly manipulating its edges. By connecting nodes B and C, the attacker indirectly contaminates node B, leading to the misclassification of A during information propagation. If the defense strategy only considers A's first-order neighborhood and uses GCN-Jaccard, it may be ineffective.

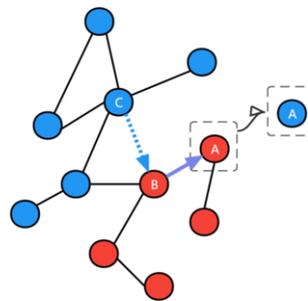


Figure 1. Indirect Attack

From this, we can deduce the importance of considering global information during defensive strategies. To substantiate this claim, we conducted a fundamental augmentation experiment on the GCN-Jaccard model by disconnecting all dissimilar second-order neighbor node pairs. The process is as follows:

By squaring the adjacency matrix, we identify all second-order neighbor node pairs (v_i, v_j) that satisfy $\mathbf{A}^2[i, j] > 0$ and calculate their similarity using the Jaccard coefficient. Subsequently, we select the node pairs with lower similarity and traverse all intermediate nodes (denoted as v_k). In each traversal, we remove the edge e_{ik} or e_{jk} with the lower similarity, effectively pruning all dissimilar second-order adjacent node pairs in a direct and intuitive manner. With this technique, we evaluate the accuracy of GCN models on the Citeseer dataset. As shown in Figure 2, utilizing higher-order neighbor information leads to significant potential for enhancing the overall defensive capability of the model.

3.3 Global Homophily Index

While improving robustness, high-order modifications to GCN-Jaccard may remove too many edges, leading to reduced accuracy. For instance, on the Cora dataset, setting the Jaccard threshold to

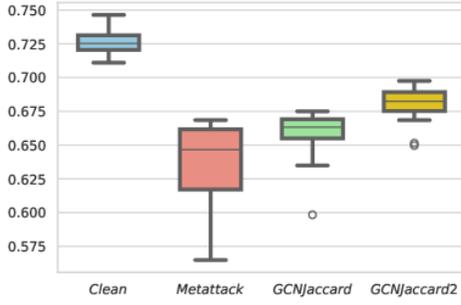


Figure 2. The augmentation experiment on the GCN-Jaccard model

0.01 results in 548 first-order dissimilar node pairs, but the number of third-order dissimilar pairs jumps to 44,799. This indicates that removing third-order dissimilar connections could cause significant edge loss. To address this, we propose a global homophily index that measures the information transmitted by all edges in the graph across all orders. By comparing the global homophily index before and after removing an edge, we can assess the edge’s importance in information transmission from low to high orders.

Consider the node pair (v_i, v_j) , which may have multiple intermediate nodes between them. In GNN models, the homophily information between these nodes is conveyed through the graph aggregation mechanism (where Jaccard similarity is used as the metric for measuring homophily). However, during the information transmission process, the message-passing mechanism can cause the loss of homophily information, which will be further strengthened by the number of paths. Consequently, our model construction takes into account both the **number of paths** and the **rate of information loss**.

3.3.1 The Number of Paths

We first consider how to obtain the number of paths n . Taking the k th power of the adjacency matrix can provide the number of k -hop paths between node pairs, which can be represented as \mathbf{A}_{ij}^k . Therefore, we can model the k th-order global homophily index $\text{Hom}^{(k)}$ as follows:

$$\text{Hom}^{(k)} = \langle \mathbf{A}^k, \mathcal{J} \rangle \quad (1)$$

where \mathbf{A} denotes the adjacency matrix, \mathcal{J} is the Jaccard similarity matrix, and $\langle \cdot, \cdot \rangle$ signifies the operation of matrix inner product.

3.3.2 The Rate of Information Loss

Zhu et al. [27] formally define the representation learned by a GNN for each node v at the k th iteration as follows:

$$\mathbf{r}_v^{(k)} = \text{COMBINE} \left(\mathbf{r}_v^{(k-1)}, \text{AGGR} \left(\left\{ \mathbf{r}_u^{(k-1)} : u \in \bar{N}(v) \right\} \right) \right) \quad (2)$$

Different GNN designs show distinct AGGR() and COMBINE() functions. AGGR() (e.g., Mean) computes the aggregation of neighbor embeddings, while COMBINE() integrates the self-embedding with the aggregated neighbor embeddings. The design of these two functions invariably introduces information loss. Hence, we denote the mean retention rate of the residual information following information aggregation as α (where $\alpha \in [0, 1]$). Subsequently, the global homophily index can be formulated as follows:

$$\text{Hom} = \sum_{k=0}^{\infty} \alpha^k \text{Hom}^{(k)} \quad (3)$$

After deriving Formulas 1 and 3, we can calculate Hom as follows.

$$\text{Hom} = \langle (I - \alpha \mathbf{A})^{-1}, \mathcal{J} \rangle \quad (4)$$

For the equation $\sum_{k=0}^{\infty} \alpha^k \mathbf{A}^k = (I - \alpha \mathbf{A})^{-1}$, certain conditions be met. Specifically, the convergence parameter α must satisfy the inequality $\alpha < \frac{1}{\rho(\mathbf{A})}$, where $\rho(\mathbf{A})$ denotes the spectral radius of the adjacency matrix \mathbf{A} . To simplify the subsequent exposition, we employ the unified notation $M = (I - \alpha \mathbf{A})^{-1}$.

3.4 Derivation and Analysis

We model the global homophily index to identify and eliminate edges that propagate the most misleading information. That is, we aim to find an edge whose removal would maximize ΔHom . This can be mathematically formulated as:

$$\underset{\Delta \mathbf{A}}{\text{argmax}} \Delta \text{Hom} \quad (5)$$

Suppose the edge e_{kl} is removed, resulting in a new adjacency matrix \mathbf{A}' . Here, $\Delta \mathbf{A} = \mathbf{A}' - \mathbf{A}$. For $\Delta \mathbf{A}$, $\Delta \mathbf{A}_{kl} = \Delta \mathbf{A}_{lk} = -1$, with all other values being 0. We further derived and optimized this result to achieve the final outcome.

$$\Delta \text{Hom} = \langle M' - M, \mathcal{J} \rangle = \alpha \langle M' \Delta \mathbf{A} M, \mathcal{J} \rangle \quad (6)$$

Where $M' = (I - \alpha \mathbf{A}')^{-1}$. Given this modification, we can determine the values at each position in the matrix $M' \Delta \mathbf{A} M$. The calculation process is detailed as follows:

$$[M' \Delta \mathbf{A} M]_{ij} = \sum_{s=1}^n \sum_{t=1}^n M'_{is} \Delta A_{st} M_{tj} = -M'_{il} M_{kj} - M'_{ik} M_{lj} \quad (7)$$

By substituting Equation (6) into Equation (5), we can perform detailed mathematical derivations and obtain the following results

$$\begin{aligned} \Delta \text{Hom} &= \alpha \langle M' \Delta \mathbf{A} M, \mathcal{J} \rangle = \alpha \sum_{i,j} [M' \Delta \mathbf{A} M]_{ij} \mathcal{J}_{ij} \\ &= -\alpha \sum_{i,j} (M'_{il} M_{kj} + M'_{ik} M_{lj}) \mathcal{J}_{ij} \\ &= -\alpha \sum_{i,j} (M'_{li} \mathcal{J}_{ij} M_{jk} + M'_{ki} \mathcal{J}_{ij} M_{jl}) \\ &= -2\alpha [M' \mathcal{J} M]_{kl} \end{aligned} \quad (8)$$

Further derivation leads to a conclusion.

$$\underset{\Delta \mathbf{A}}{\text{argmax}} \Delta \text{Hom} = \underset{\Delta \mathbf{A}}{\text{argmin}} \langle M' \mathcal{J} M, -\Delta \mathbf{A} \rangle \quad (9)$$

We have mathematically represented the impact on the global homophily index after the removal of an edge, providing mathematical support for Talos. In subsequent content, we will present its implementation and adopt various optimizations to ensure that the strategy is lightweight and efficient.

3.5 Optimizations

As shown in the previous section, Talos can be formulated as solving the problem of $\operatorname{argmin}_{\Delta \mathbf{A}} < M' \mathcal{J} M, -\Delta \mathbf{A} >$. However, there is a notable challenge to implement it. Whenever the decision is made to remove a specific edge, it necessitates traversing and computing the matrix $M' \mathcal{J} M$, which significantly increases the computational complexity. Fortunately, our primary objective is not to obtain precise numerical values but rather to effectively rank and evaluate edges with greater levels of harm. In light of this, we introduce an approximate optimization to the method, which enables substantial enhancement for computational efficiency.

3.5.1 Approximate Optimization Strategy

Due to the minor perturbation of state \mathbf{A} before and after defending a specific edge, we reasonably postulate the approximation between \mathbf{A} and \mathbf{A}' . Consequently, we can also approximate the similarity between M' and M . Grounded in this approximation, we can approximate the original formula as follows:

$$\operatorname{argmax}_{\Delta \mathbf{A}} \Delta \text{Hom} \approx \operatorname{argmin}_{\Delta \mathbf{A}} < M \mathcal{J} M, -\Delta \mathbf{A} > \quad (10)$$

This approximation allows us to circumvent the intricate computation involving M' , thereby simplifying the computation.

3.5.2 Efficient Batch Selection Strategy

If the recalculation of $M \mathcal{J} M$ is required every time an edge is removed, the computational cost becomes prohibitively high. However, in essence, precise calculations are not our primary need; rather, we seek to obtain a ranking of edges based on their homophily. Therefore, we can directly identify the edges corresponding to the positions of the smallest n values in the matrix $M \mathcal{J} M$ and focus our defense efforts on these edges. This approach eliminates the need for performing n iterations and repeatedly computing a new $M \mathcal{J} M$ to select the minimum value in each iteration.

3.5.3 Matrix Computation Strategy

Considering that node features in graph datasets are often discrete, as exemplified by common graph datasets such as CORA where the feature matrix \mathbf{X} contains elements with values limited to 0,1. Given this context, we can compute the Jaccard coefficient between node v_i and node v_j using the following approach:

$$\mathcal{J}_{ij} = \frac{\mathbf{x}_i \mathbf{x}_j^T}{1 - (1 - \mathbf{x}_i)(1 - \mathbf{x}_j)^T} = \frac{\mathbf{x}_i \mathbf{x}_j^T}{1 - \bar{\mathbf{x}}_i \bar{\mathbf{x}}_j^T} = \frac{\mathbf{x}_i \mathbf{x}_j^T}{\bar{\mathbf{x}}_i \bar{\mathbf{x}}_j^T} \quad (11)$$

This further leads to $\mathcal{J} = \frac{\mathbf{X} \mathbf{X}^T}{\bar{\mathbf{X}} \bar{\mathbf{X}}^T}$

By employing matrix operations, we can achieve significant computational savings, as the division operation here refers to element-wise division between corresponding positions in the matrices.

3.6 Overview of Talos

As shown in Algorithm 1, Talos can be seamlessly incorporated into existing GNNs. Given a graph $\mathcal{G}_p = (\mathcal{V}, \mathcal{E}_p, \mathbf{X})$ subjected to poisoning attacks, we can efficiently identify and remove suspected attacked edges, thus obtaining a clean graph structure data \mathcal{G}_c . Talos first calculates the Jaccard similarity matrix with reference to Eq. (11). Then it computes the similar gradient matrix $M \mathcal{J} M$, and we only need to select the smallest n values in the lower triangular part of this symmetric matrix and delete the corresponding edges.

Algorithm 1 Defense procedure of Talos

Require: Poisoned graph $\mathcal{G}_p = (\mathcal{V}, \mathcal{E}_p, \mathbf{X})$ with adjacency matrix \mathbf{A}_p ; Retention Rate α ; Number of edges to cut n

- 1: Compute the Jaccard similarity matrix \mathcal{J} using Eq. (11);
- 2: Evaluate $M \mathcal{J} M$, where $M = (1 - \mathbf{A}_p)^{-1}$;
- 3: Identify the n edge zeroes in \mathbf{A}_p corresponding to the smallest values in $M \mathcal{J} M$, resulting in matrix \mathbf{A}_c ;
- 4: Purifying the edge set of the graph structure into \mathcal{E}_c based on \mathbf{A}_c , we obtain the sanitized clean graph $\mathcal{G}_c = (\mathcal{V}, \mathcal{E}_c, \mathbf{X})$;
- 5: **return** \mathcal{G}_c

4 Experiments

In this section, we assess the efficacy of Talos against non-targeted adversarial attacks on graphs. Experiments were conducted on servers equipped with NVIDIA Tesla V100-SXM2 32GB GPUs. Through the experiments, we evaluate how Talos performs compared with the state-of-the-art defense methods. We also show that Talos can be applied to different types of GNN models.

4.1 Setup

4.1.1 Datasets

In this study, we primarily utilized three widely recognized datasets: Cora[21], Amazon Photo[22], and Amazon Computers. The selection of these three datasets was based on three main considerations:

Firstly, these datasets are commonly used to evaluate the performance of GNN algorithms.

Secondly, their scales cover the needs of small, medium, and large-scale graph-structured data, allowing for a more comprehensive evaluation of the algorithms.

Thirdly, under the constraints of our experimental conditions, these three datasets could undergo a thorough evaluation. As for the widely regarded large dataset, Pubmed, due to experimental constraints, we followed the work of Daniel Zügner [6] and conducted separate tests on it. In the experiments, we operated on the largest connected subgraph for each dataset. All datasets were randomly divided, with 10% of the data allocated for training, 10% for validation, and the remaining 80% for testing, thereby providing a rigorous environment for evaluating model performance.

4.1.2 Generating Adversarial Attacks

To evaluate the effectiveness of Talos, we first conducted attacks on graph data. We utilized two global attack methods: Metattack[28] and Projected Gradient Descent (PGD)[20]. During the implementation of the Metattack, we employed a simple GCN as a proxy model to guide the direction of the attack. This proxy model was configured with a two-layer structure, including a hidden layer with 64 nodes. To reduce the risk of overfitting and enhance the model's generalization ability, we set a dropout ratio of 0.5. The learning rate (lr) was set to 0.01, and the weight decay was set to 0. Additionally, the model was trained for 1000 epochs with an early stopping strategy to prevent overfitting. For Metattack itself, we utilized a meta-gradient approach with self-training and set the momentum parameter to 0.9. During the implementation of the PGDAttack, we conducted multiple attacks on datasets (Cora, Amazon Photo, and Amazon Computers) using these two adversarial attack methods. In these attacks, we perturbed the edges in the graph, with the proportion of attacked edges ranging from 5% to 25% in steps of 5%. After the attacks, we

used a GCN with the same settings as the Metattack proxy model for detection to evaluate the performance of the graph data after poisoning attacks. We have detailed the results of the attacks in Table 1, and the attacked graph data will serve as the foundation for further defense evaluation.

Table 1. Attack Result

Dataset	Ptb Rate(%)	Metattack(%)	PGD(%)
Cora	5	79.38±0.72	78.74±0.21
	10	75.31±0.50	75.80±0.77
	15	64.16±1.00	74.78±0.53
	20	55.23±1.01	73.01±0.63
	25	52.44±0.39	71.86±0.58
Photo	5	90.19±1.11	85.23±0.21
	10	84.18±1.06	82.06±0.30
	15	73.84±2.83	79.56±0.70
	20	59.54±0.43	78.42±0.45
	25	57.65±1.57	76.46±1.03
Computers	5	81.68±1.06	79.22±0.56
	10	74.20±1.84	74.40±0.74
	15	72.10±1.57	70.72±0.69
	20	71.06±1.40	68.58±0.89
	25	69.50±1.16	66.66±0.88

4.2 Defense Performance

In this section, we conduct a comprehensive evaluation of the node classification accuracy and execution time for various defense mechanisms against non-targeted attacks. Our Talos is compared with several state-of-the-art defense methods, including GCN-SVD, GCN-Jaccard, RobustGCN [26], GNNGuard, and Pro-GNN.

Specifically, in the case of the Cora, we meticulously followed the hyperparameter configurations outlined in the respective baseline defense methodologies. For the Photo and Computers, since the original defense algorithms did not have experimental results, we adjusted the hyperparameters for each method to ensure the fairness of the experiment. The precise hyperparameters are as follows: for GCN-SVD, $k = 100$; for GCN-Jaccard, threshold = 0.1; for RobustGCN, $\gamma = 0.2$. GNNGuard and Pro-GNN utilized their default settings. For the hyperparameters of Talos, validation sets were used for optimization. In addition, in order to ensure the fairness of the experiment, all the above algorithms adopt a two-layer GNN during the experiment, and the number of hidden units is 16.

4.2.1 Accuracy Evaluation

We initially assess the accuracy of node classification for various defense methods against non-target adversarial attacks. The adversarial attack dataset, generated in a prior section using Metattack and PGD, was employed for our experiments. Tables 2 present the mean accuracy along with the standard deviation under the two respective attack methods. Performance rankings are indicated in the upper right corner, with the top-performing results highlighted in bold. As observed from the tables, our approach consistently leads under diverse disturbance rates, demonstrating significant efficacy under higher levels of disturbance. Despite the decimdrule in classification accuracy across all defense methods as the attack intensity increased, it is worth noting that Talos tended to maintain a high classification accuracy and showed strong robustness under higher attack intensity.

In scenarios involving origin graphics and low percentage attacks, GCN’s accuracy is not substantially compromised. Defense strategies that employ structure purification (such as GCN-Jaccard) often damage the original structure and have the opposite effect on graphs. Through cumulative means, Talos can more accurately differentiate dissimilar edges, thus better filtering out attacked edges.

4.2.2 Execution Time

Execution time is a crucial metric for evaluating defense methods. However, comparing their practicality based solely on time complexity is challenging due to various parallel computing techniques and defensive strategies. For example, GCN-Jaccard focuses on graph structure purification and doesn’t include GNN fitting processes, while methods like RobustGCN do. To demonstrate practicality accurately, we measure execution time from the start of each retraining session to the completion of model fitting. Across six experiments with attack ratios from 0 to 25%, each group was conducted ten times, yielding an average running time recorded in Table 3.

To ensure fairness, both graph structure purification methods (GCN-Jaccard, GCN-SVD, and Talos) and methods employing robust graph neural networks (RobustGCN, GNNGuard, Pro-GNN) utilized similar configurations, including a two-layer GCN with 16 hidden units and a learning rate of 0.01. Talos operates significantly faster than almost any other defense method, as shown in the table. Although the experiment was conducted on a cloud computing platform where concurrent GPU usage may have influenced execution times, Talos closely resembles the processing speed of GCN, indicating its efficiency. Further investigation revealed that the average operation times for Talos across different datasets were 0.040s, 0.378s, and 2.582s for Cora, Photo, and Computers, respectively. Given its faster processing time compared to GCN, Talos emerges as an effective and practical strategy for purifying GNN data before defense.

4.2.3 Larger Scale Graph Application

Due to platform limitations, we were unable to perform Metattack on Pubmed, making subsequent defense experiments infeasible. To test whether Talos is more effective on larger graphs, we conducted experiments on Pubmed using PR-BCD to perturb the graph with 1%, 5%, and 10% attacks. We compared the results with baselines like Vanilla GCN, Vanilla GDC, GCNJaccard, and SoftMedian [6]. For fairness, all models used the same GCN structure with 64 hidden units. For SoftMedian, we used the SoftMedian GDC (T=1) method and also tested replacing GDC with Talos for preprocessing, combining it with SoftMedian. The accuracy results are shown in Figure 3.

Regarding experiment times, Vanilla GCN and Vanilla GDC took 1.63s and 5.80s, respectively. With defenses, GCNJaccard, SoftMedian GDC, SoftMedian Talos, and GCN Talos took 14.24s, 21.29s, 16.92s, and 11.61s, respectively. GCN Talos had the shortest time overall.

The experiments demonstrate that Talos offers the best defense under PR-BCD attacks on large graphs and has the lowest time consumption among the methods compared.

4.3 Universality of Talos Over GNNs

Talos only uses the information of the graph itself and the assumption of homophily property. Therefore, Talos is theoretically applicable to all kinds of GNNs. Although previous experiments have been tested on GCN to verify its effectiveness against other defense methods, it is

Table 2. Node Classification Performance under Metattack and PGD(%)

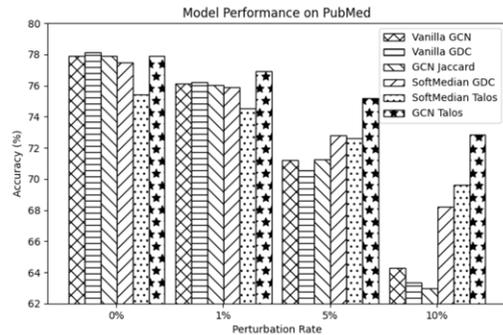
DataSet	Attack	Ptb Rate(%)	GCN	GCN-Jaccard	GCN-SVD	RobustGCN	GNNGuard	Pro-GNN	Talos
Cora	Metattack	0	83.11±0.41	80.89±0.10 ⁴	76.51±0.52 ⁶	81.49±0.33 ³	76.62±1.24 ⁵	81.66±2.50 ²	82.71±0.49¹
		5	79.38±0.72	79.31±0.45 ²	75.11±0.57 ⁵	77.65±0.83 ⁴	75.00±0.45 ⁶	79.08±1.60 ³	80.83±0.74¹
		10	75.31±0.50	77.08±0.07 ²	72.34±0.34 ⁶	73.37±1.12 ⁵	74.84±0.51 ⁴	76.64±2.15 ³	78.31±0.76¹
		15	64.16±1.00	72.90±0.38 ²	66.05±0.33 ⁵	63.61±1.31 ⁶	71.00±1.65 ⁴	72.21±1.53 ³	74.17±0.39¹
		20	55.23±1.01	65.88±1.39 ³	54.09±0.81 ⁶	55.31±1.22 ⁵	65.98±2.57 ²	58.47±1.12 ⁴	68.19±0.98¹
		25	52.44±0.39	60.96±0.68 ³	50.01±0.79 ⁶	50.60±0.46 ⁵	65.81±1.78¹	53.12±0.76 ⁴	65.28±1.16 ²
	PGD	0	82.87±0.45	80.94±0.28 ⁴	75.60±0.48 ⁶	81.95±0.23¹	76.81±0.36 ⁵	81.16±2.17 ²	81.06±0.91 ³
		5	78.74±0.21	79.03±0.43¹	76.73±0.71 ⁵	77.40±0.46 ⁴	76.07±0.61 ⁶	78.52±1.38 ³	78.81±0.47 ²
		10	75.80±0.77	76.24±0.30 ²	76.07±0.44 ³	75.90±0.84 ⁴	75.30±0.62 ⁶	75.57±0.48 ⁵	78.42±0.54¹
		15	74.78±0.53	74.56±0.73 ⁴	73.84±1.05 ⁶	74.78±0.51 ³	75.63±0.74 ²	74.01±0.58 ⁵	75.97±0.41¹
		20	73.01±0.63	73.01±0.66 ³	71.77±1.07 ⁶	72.38±0.73 ⁵	74.16±0.16 ²	72.39±0.47 ⁴	75.40±0.64¹
		25	71.86±0.58	72.39±0.56 ³	72.36±0.13 ⁴	71.94±0.29 ⁵	72.71±1.09 ²	71.79±0.11 ⁶	73.98±0.34¹
Photo	Metattack	0	93.80±0.21	93.09±0.07 ⁴	90.82±0.15 ⁵	93.35±0.21 ²	93.20±0.44 ³	90.36±1.39 ⁶	93.60±0.27¹
		5	90.19±1.11	92.29±0.22¹	87.73±0.24 ⁵	89.19±1.41 ³	88.70±2.54 ⁴	87.52±1.36 ⁶	92.17±0.38 ²
		10	84.18±1.06	90.20±0.42 ²	84.00±0.62 ⁴	84.00±0.35 ⁵	85.57±0.43 ³	83.78±3.03 ⁶	90.34±0.49¹
		15	73.84±2.83	84.28±1.16 ²	68.63±4.44 ⁶	71.73±3.19 ⁵	79.28±0.80 ³	77.82±2.04 ⁴	88.60±0.70¹
		20	59.54±0.43	82.87±1.52 ²	62.59±3.21 ⁵	60.53±3.45 ⁶	66.56±1.99 ⁴	76.54±2.27 ³	87.29±0.66¹
		25	57.65±1.57	79.71±2.88 ²	52.09±3.97 ⁶	64.55±2.46 ⁴	66.23±2.04 ³	56.70±2.75 ⁵	87.01±0.98¹
	PGD	0	93.64±0.25	93.21±0.08 ⁴	90.72±0.26 ⁵	93.52±0.23¹	93.51±0.15 ²	90.36±1.39 ⁶	93.50±0.33 ³
		5	85.23±0.21	86.72±0.12 ²	86.42±0.38 ³	85.88±0.11 ⁴	85.51±0.15 ⁵	84.11±0.57 ⁶	88.86±0.24¹
		10	82.06±0.30	83.39±0.19 ²	82.19±0.44 ⁵	83.12±0.05 ³	82.53±0.30 ⁴	81.54±0.46 ⁶	86.28±0.55¹
		15	79.56±0.70	80.37±0.69 ³	79.40±0.27 ⁶	81.04±0.08 ²	80.25±0.29 ⁴	79.62±0.47 ⁵	83.91±0.62¹
		20	78.42±0.45	78.96±0.45 ⁵	78.35±0.25 ⁶	79.50±0.08 ²	79.10±0.24 ³	79.09±0.39 ⁴	81.34±0.83¹
		25	76.46±1.03	77.78±0.04 ⁵	76.51±0.98 ⁶	78.33±0.42 ²	77.95±0.39 ⁴	77.98±0.28 ³	79.62±0.77¹
Computers	Metattack	0	88.30±0.43	87.44±0.42 ⁴	77.64±0.88 ⁵	88.71±0.32 ²	88.35±0.39 ³	*	88.98±0.39¹
		5	81.68±1.06	83.90±0.50 ²	73.10±0.72 ⁵	78.22±0.55 ⁴	82.18±0.66 ³	*	85.58±0.53¹
		10	74.20±1.84	78.60±1.14¹	67.43±0.76 ⁵	74.40±0.40 ⁴	75.63±0.59 ³	*	78.30±1.91 ²
		15	72.10±1.57	78.06±1.26 ²	64.78±1.10 ⁵	71.57±0.94 ⁴	74.21±0.55 ³	*	80.58±1.05¹
		20	71.06±1.40	77.54±0.85 ²	64.90±1.77 ⁵	71.11±1.00 ⁴	72.91±0.45 ³	*	80.49±0.67¹
		25	69.50±1.16	76.02±1.52 ²	63.92±1.44 ⁵	70.28±0.70 ⁴	71.95±0.54 ³	*	79.83±1.73¹
	PGD	0	88.30±0.43	87.26±0.38 ⁴	77.64±0.88 ⁵	88.69±0.31 ²	88.34±0.35 ³	*	88.98±0.39¹
		5	79.22±0.56	81.81±0.24 ²	75.07±0.78 ⁵	80.77±0.24 ³	80.41±0.28 ⁴	*	85.29±0.36¹
		10	74.40±0.74	77.60±0.65 ²	71.37±1.26 ⁵	77.38±0.18 ³	75.02±0.49 ⁴	*	82.95±0.69¹
		15	70.72±0.69	73.40±1.07 ³	67.87±1.07 ⁵	74.29±0.39 ²	71.45±0.47 ⁴	*	81.03±0.91¹
		20	68.58±0.89	70.30±1.41 ³	65.90±1.21 ⁵	71.76±0.31 ²	69.79±0.54 ⁴	*	79.07±0.86¹
		25	66.66±0.88	67.85±0.63 ⁴	64.13±1.23 ⁵	70.31±0.39 ²	67.86±0.61 ³	*	78.14±0.83¹

Table 3. Execution time(s)

DataSet	Cora	Photo	Computers
GCN	0.63±0.14	2.67±0.37	8.12±1.90
GCN-Jaccard	1.28±0.31	3.67±0.47	8.79±1.24
GCN-SVD	3.42±0.01	46.86±0.66	328.82±6.62
RobustGCN	5.66±0.02	8.13±0.04	37.71±0.06
GNNGuard	46.98±0.58	1239.13±148.32	1346.29±240.29
Pro-GNN	973.90±100.42	16289.10±204.53	Time Out
Talos	0.60±0.10	2.80±0.20	11.90±1.54

unclear whether it is also better on other GNNs. In order to verify the generality of Talos on GNNs, this section selects several common GNNs for testing, including GCN, GAT[15], and GraphSage[10]. The experimental results are shown in Table 4. Due to the large scale of the Computers dataset, on the experimental platform, Pro-GNN failed to complete the computation and produce results within the predetermined time frame. Therefore, we marked it as a timeout and excluded Pro-GNN from the comparative analysis of the Computers.

As can be seen from the Table, Talos achieves significant results in all three basic models, so Talos essentially purifies the graph structure regardless of the selected model. The defense of Talos over GNNs is universal.

**Figure 3.** Model Performance on Pubmed

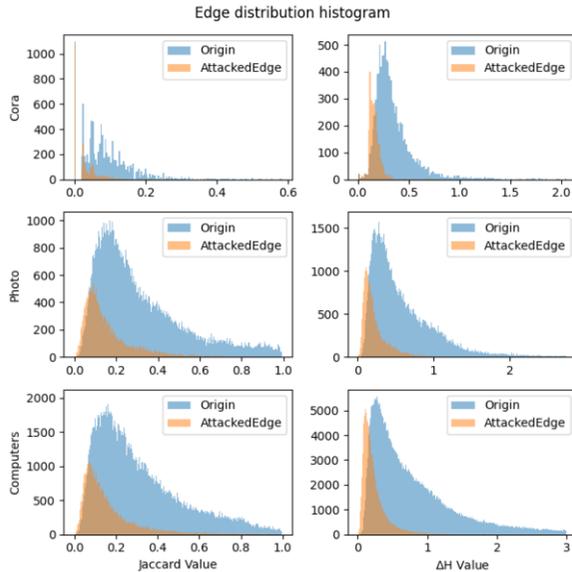
4.4 Why is Talos Effective

To explore why Talos is effective, we conducted a comparative analysis between Talos and GCN-Jaccard. Using the PGD Attack, we generated attacked graphs with a 25% attack rate. By plotting histograms of attack and clean edges for graphs processed by GCN-Jaccard and Talos, we compared their abilities to distinguish perturbed edges. Figure 4 illustrates that Talos, with its ΔHom metric, widens the distribution gap between attack and clean edges compared to GCN-

Table 4. Universality Performance under PGD(%)

DataSet	GNN Model	Defense Method	Clean Graph	Ptb5%	Ptb10%	Ptb15%	Ptb20%	Ptb25%
Cora	GCN	NoDefense	83.10±0.44	78.57±0.21	76.98±0.29	75.21±0.34	73.58±0.19	72.23±0.31
		GNNJaccard	81.53±0.33	78.74±0.67	76.75±0.85	75.16±0.99	73.49±0.65	72.86±0.66
		Talos	79.71±0.51	78.67±0.42	78.08±0.63	77.31±0.43	74.99±0.66	75.46±0.48
	GAT	NoDefense	83.78±0.84	80.31±0.61	77.61±0.43	75.54±0.74	73.83±0.90	72.40±0.75
		GNNJaccard	82.13±0.78	80.16±1.18	78.15±1.36	76.70±1.08	74.48±1.06	73.63±1.31
		Talos	83.10±1.20	79.85±1.01	79.71±1.37	78.79±0.82	77.55±0.95	76.45±0.94
GragphSAGE	NoDefense	82.98±0.50	79.68±0.60	77.55±0.50	76.35±0.65	74.36±0.61	73.49±0.38	
	GNNJaccard	81.52±0.51	80.18±0.55	78.59±0.40	77.21±0.62	75.55±0.64	75.16±0.50	
	Talos	82.96±0.39	80.42±0.44	79.10±0.83	78.72±0.47	77.68±0.35	77.17±0.48	
Photo	GCN	NoDefense	93.88±0.18	85.27±0.26	82.18±0.30	80.04±0.35	78.79±0.47	77.72±0.58
		GNNJaccard	93.29±0.16	87.17±0.31	83.75±0.40	80.81±0.26	79.28±0.25	77.91±0.43
		Talos	93.62±0.19	89.73±0.20	88.56±0.35	86.68±0.32	86.37±0.49	84.53±0.64
	GAT	NoDefense	94.10±0.22	85.02±0.21	82.64±0.44	80.38±0.59	78.77±1.26	77.89±0.37
		GNNJaccard	93.70±0.20	87.09±0.34	83.92±0.18	81.87±0.34	80.13±0.44	79.06±0.35
		Talos	93.99±0.25	90.20±0.29	88.98±0.37	87.27±0.29	86.49±0.39	85.99±0.51
GragphSAGE	NoDefense	94.41±0.21	88.06±0.25	85.19±0.49	83.78±0.49	82.61±0.45	82.03±0.48	
	GNNJaccard	94.18±0.27	90.32±0.28	87.64±0.26	85.94±0.46	84.44±0.33	82.89±0.68	
	Talos	94.16±0.19	91.43±0.21	90.54±0.33	89.26±0.56	88.89±0.34	88.50±0.30	
Computers	GCN	NoDefense	88.80±0.24	80.68±0.47	74.24±1.13	70.42±1.05	67.86±1.06	65.96±1.01
		GNNJaccard	88.36±0.39	82.76±0.44	79.15±0.51	75.00±1.53	70.57±1.49	68.30±1.14
		Talos	88.98±0.39	85.29±0.36	82.95±0.69	81.03±0.91	79.07±0.86	78.14±0.83
	GAT	NoDefense	89.22±0.19	80.89±0.46	77.05±1.61	71.65±2.58	69.82±2.50	66.74±2.85
		GNNJaccard	88.24±0.45	82.76±0.47	78.80±1.07	74.54±1.90	73.19±1.76	69.10±2.29
		Talos	89.41±0.48	85.22±0.42	83.66±0.58	82.58±0.63	81.72±0.48	81.14±0.50
GragphSAGE	NoDefense	88.50±0.71	81.96±1.24	78.10±1.33	75.66±1.68	74.07±1.49	71.36±1.41	
	GNNJaccard	87.10±0.66	83.60±0.84	81.03±0.77	79.02±1.40	77.24±0.98	75.36±2.01	
	Talos	87.63±0.55	84.99±0.99	83.07±1.35	81.53±1.78	80.88±0.99	81.19±1.21	

Jaccard. It also creates a clustering effect on attack edges, facilitating their identification and removal.

**Figure 4.** Edge distribution histogram under PGD

We also adopted two quantitative metrics to measure the distribution difference between attack edges and clean edges. Firstly, we used the Kolmogorov-Smirnov (KS) metric to quantify the difference between the two distributions. The KS metric is a measure in statistics used to gauge the disparity between two distributions. The range of the KS metric is from 0 to 1, with a higher KS value indicating a stronger discriminative capability of the model. Generally,

a KS metric between 0.2 and 0.5 is considered to have a good discriminative capability; above 0.5, the model’s discriminative capability is very strong. After calculation, for Jaccard, $\mathbf{KS}_{cora} = 0.46$, $\mathbf{KS}_{photo} = 0.39$, $\mathbf{KS}_{computers} = 0.38$. For Talos’ ΔHom , $\mathbf{KS}_{cora} = 0.66$, $\mathbf{KS}_{photo} = 0.51$, $\mathbf{KS}_{computers} = 0.57$. The results indicate that Talos has a very strong discriminative capability for attack edges and significantly surpasses Jaccard.

Secondly, we defined and calculated the separation rate (SR) for attack edges for both by $\mathbf{SR} = \frac{S_{\text{attacked}} - S_{\text{overlap}}}{S_{\text{attacked}}}$ where S_{attacked} is histogram area of attack edge distribution, S_{overlap} is area of overlap between the two histograms. After calculation, for Jaccard, $\mathbf{SR}_{cora} = 0$, $\mathbf{SR}_{photo} = 6.93\%$, $\mathbf{SR}_{computers} = 6.45\%$. For Talos’ ΔHom , $\mathbf{SR}_{cora} = 29.8\%$, $\mathbf{SR}_{photo} = 22.3\%$, $\mathbf{SR}_{computers} = 27.5\%$. It can be seen that Talos has a very strong discriminative capability for attack edges and significantly surpasses Jaccard.

In summary, Talos is more effective because it better distinguishes between attack edges and clean edges. During graph purification, it removes more attack edges while minimizing harm to the original graph, leading to better defensive results.

5 Conclusion

This paper introduces Talos, a generalized defense method that improves the robustness of GNN models. Talos enhances global homophily in graphs, effectively removing adversarial edges. We demonstrate that with proper approximations, Talos is highly efficient. Comparative experiments confirm Talos’ advantages in both time and accuracy, proving its applicability across different GNN models.

Due to page limits, additional supplementary materials are available on arXiv.¹

¹ Supplementary materials: <https://arxiv.org/abs/2406.03833>

Acknowledgements

This project is supported by the National Natural Science Foundation of China under grant No. 62306328 and the science and technology innovation Program of Hunan Province under grant No. 2023RC3021.

References

- [1] K.-W. Chang, H. He, R. Jia, and S. Singh. Robustness and adversarial examples in natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 22–26, 2021.
- [2] H. Dai, H. Li, T. Tian, X. Huang, L. Wang, J. Zhu, and L. Song. Adversarial attack on graph structured data. In *International conference on machine learning*, pages 1115–1124. PMLR, 2018.
- [3] N. Entezari, S. A. Al-Sayouri, A. Darvishzadeh, and E. E. Papalexakis. All you need is low (rank): Defending against adversarial attacks on graphs. In *WSDM*, pages 169–177. ACM, 2020.
- [4] W. Feng, J. Zhang, Y. Dong, Y. Han, H. Luan, Q. Xu, Q. Yang, E. Kharlamov, and J. Tang. Graph random neural networks for semi-supervised learning on graphs. *Advances in neural information processing systems*, 33:22092–22103, 2020.
- [5] S. Geisler, D. Zügner, and S. Günnemann. Reliable graph neural networks via robust aggregation. *Advances in neural information processing systems*, 33:13272–13284, 2020.
- [6] S. Geisler, T. Schmidt, H. Sirin, D. Zügner, A. Bojchevski, and S. Günnemann. Robustness of graph neural networks at scale. In M. Ranzato, A. Beygelzimer, Y. N. Dauphin, P. Liang, and J. W. Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 7637–7649, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/3ea2db50e62ceefceaf70a9d9a56a6f4-Abstract.html>.
- [7] I. J. Goodfellow, J. Shlens, and C. Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [8] S. Günnemann. Graph neural networks: Adversarial robustness. *Graph neural networks: foundations, frontiers, and applications*, pages 149–176, 2022.
- [9] W. Hamilton, Z. Ying, and J. Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.
- [10] W. L. Hamilton, R. Ying, and J. Leskovec. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 1025–1035, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
- [11] J. Huang, L. Du, X. Chen, Q. Fu, S. Han, and D. Zhang. Robust mid-pass filtering graph convolutional networks. In *Proceedings of the ACM Web Conference 2023*, pages 328–338, 2023.
- [12] W. Jin, Y. Ma, X. Liu, X. Tang, S. Wang, and J. Tang. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 66–74, 2020.
- [13] F. Mujkanovic, S. Geisler, S. Günnemann, and A. Bojchevski. Are defenses for graph neural networks robust? *Advances in Neural Information Processing Systems*, 35:8954–8968, 2022.
- [14] S. Tao, Q. Cao, H. Shen, J. Huang, Y. Wu, and X. Cheng. Single node injection attack against graph neural networks. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, pages 1794–1803, 2021.
- [15] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio. Graph attention networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rJXMpikCZ>.
- [16] Y. Wang, W. Wang, Y. Liang, Y. Cai, and B. Hooi. Mixup for node and graph classification. In *Proceedings of the Web Conference 2021*, pages 3663–3674, 2021.
- [17] H. Wu, C. Wang, Y. Tyshetskiy, A. Docherty, K. Lu, and L. Zhu. Adversarial examples for graph data: deep insights into attack and defense. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 4816–4823, 2019.
- [18] X.-G. Wu, H.-J. Wu, X. Zhou, X. Zhao, and K. Lu. Towards defense against adversarial attacks on graph neural networks via calibrated co-training. *Journal of Computer Science and Technology*, 37(5):1161–1175, 2022.
- [19] J. Xu, J. Chen, S. You, Z. Xiao, Y. Yang, and J. Lu. Robustness of deep learning models on graphs: A survey. *AI Open*, 2:69–78, 2021.
- [20] K. Xu, H. Chen, S. Liu, P. Chen, T. Weng, M. Hong, and X. Lin. Topology attack and defense for graph neural networks: An optimization perspective. In *IJCAI*, pages 3961–3967. ijcai.org, 2019.
- [21] X. Yang, Y. Liu, S. Zhou, S. Wang, W. Tu, Q. Zheng, X. Liu, L. Fang, and E. Zhu. Cluster-guided contrastive graph clustering network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 10834–10842, 2023.
- [22] X. Yang, C. Tan, Y. Liu, K. Liang, S. Wang, S. Zhou, J. Xia, S. Z. Li, X. Liu, and E. Zhu. Convert: Contrastive graph clustering with reliable augmentation. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 319–327, 2023.
- [23] H. Zhang, B. Wu, X. Yuan, S. Pan, H. Tong, and J. Pei. Trustworthy graph neural networks: Aspects, methods, and trends. *Proceedings of the IEEE*, 2024.
- [24] M. Zhang and Y. Chen. Link prediction based on graph neural networks. *Advances in neural information processing systems*, 31, 2018.
- [25] X. Zhang and M. Zitnik. GnnGuard: Defending graph neural networks against adversarial attacks. In *NeurIPS*, 2020.
- [26] D. Zhu, Z. Zhang, P. Cui, and W. Zhu. Robust graph convolutional networks against adversarial attacks. In *KDD*, pages 1399–1407. ACM, 2019.
- [27] J. Zhu, Y. Yan, L. Zhao, M. Heimann, L. Akoglu, and D. Koutra. Beyond homophily in graph neural networks: Current limitations and effective designs. *Advances in neural information processing systems*, 33:7793–7804, 2020.
- [28] D. Zügner and S. Günnemann. Adversarial attacks on graph neural networks via meta learning. In *ICLR (Poster)*. OpenReview.net, 2019.
- [29] D. Zügner, A. Akbarnejad, and S. Günnemann. Adversarial attacks on neural networks for graph data. In *KDD*, pages 2847–2856. ACM, 2018.