

---

# GRAPHPATCHER: Mitigating Degree Bias for Graph Neural Networks via Test-time Augmentation

---

Mingxuan Ju<sup>1</sup>, Tong Zhao<sup>2</sup>, Wenhao Yu<sup>1</sup>, Neil Shah<sup>2</sup>, Yanfang Ye<sup>1</sup>

<sup>1</sup>University of Notre Dame, <sup>2</sup>Snap Inc.

<sup>1</sup>{mju2, wyu1, yye7}@nd.edu; <sup>2</sup>{tzhao, nshah}@snap.com

## Abstract

Recent studies have shown that graph neural networks (GNNs) exhibit strong biases towards the node degree: they usually perform satisfactorily on high-degree nodes with rich neighbor information but struggle with low-degree nodes. Existing works tackle this problem by deriving either designated GNN architectures or training strategies specifically for low-degree nodes. Though effective, these approaches unintentionally create an artificial out-of-distribution scenario, where models mainly or even only observe low-degree nodes during the training, leading to a downgraded performance for high-degree nodes that GNNs originally perform well at. In light of this, we propose a test-time augmentation framework, namely GRAPHPATCHER, to enhance test-time generalization of any GNNs on low-degree nodes. Specifically, GRAPHPATCHER iteratively generates virtual nodes to patch artificially created low-degree nodes via corruptions, aiming at progressively reconstructing target GNN’s predictions over a sequence of increasingly corrupted nodes. Through this scheme, GRAPHPATCHER not only learns how to enhance low-degree nodes (when the neighborhoods are heavily corrupted) but also preserves the original superior performance of GNNs on high-degree nodes (when lightly corrupted). Additionally, GRAPHPATCHER is model-agnostic and can also mitigate the degree bias for either self-supervised or supervised GNNs. Comprehensive experiments are conducted over seven benchmark datasets and GRAPHPATCHER consistently enhances common GNNs’ overall performance by up to 3.6% and low-degree performance by up to 6.5%, significantly outperforming state-of-the-art baselines. The source code is publicly available at <https://github.com/jumxglhf/GraphPatcher>.

## 1 Introduction

Graph Neural Networks (GNNs) have gained significant popularity as a powerful approach for learning representations of graphs, achieving state-of-the-art performance on various predictive tasks, such as node classification [22, 38, 9], link prediction [50, 53], and graph classification [43, 47, 11]. These tasks further form the archetypes of many real-world applications, such as recommendation systems [45, 3], predicative user behavior models [31, 52], and molecular property prediction [51, 48].

While existing GNNs are highly proficient at capturing information from rich neighborhoods (i.e., high-degree nodes), recent studies [13, 25, 36, 55] have revealed a significant performance degradation of GNNs when dealing with nodes that have sparse neighborhoods (i.e., low-degree nodes). This observation can be attributed to the fact that GNNs make predictions based on the distribution of node neighborhoods [27]. According to this line of theory, GNNs struggle with low-degree nodes due to the limited amount of available neighborhood information, which may not be able to precisely depict the learned distributions. Empirically, as shown in Figure 1, the classification accuracy of GCN [22] proportionally decays as the node degree decreases, resulting in a performance gap of  $\sim 20\%$  accuracy. Furthermore, the sub-optimal performance of GNNs on low-degree nodes can be

## Node Degree vs. Accuracy across SoTA

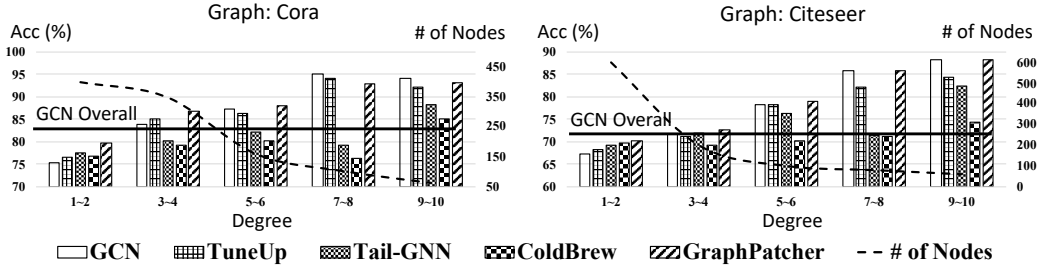


Figure 1: The classification accuracy of GCN and SoTA frameworks that mitigate degree biases.

aggravated by the power-law degree distribution commonly observed in real-world graphs, where the number of low-degree nodes significantly exceeds that of high-degree nodes [36].

To bridge this gap, several frameworks have been proposed to specifically improve GNNs’ performance on low-degree nodes [36, 25, 13, 55, 49]. These frameworks either introduce designated architectures or training strategies specifically for low-degree nodes. For examples, Tail-GNN [25] enhances latent representations of low-degree nodes by incorporating high-degree structural information; whereas Coldbrew [55] retrieves a set of existing nodes as virtual neighbors for low-degree nodes. However, these approaches suffer from two significant drawbacks. Firstly, while benefiting low-degree nodes, they inadvertently create an artificial out-of-distribution scenario during training [42], where models primarily observe low-degree nodes, leading to a downgraded performance for high-degree nodes that GNNs originally perform well on. Secondly, deploying these frameworks often requires changing model architectures, which can be impractical in real-world scenarios where the original models are well-trained due to the expensive re-training cost (on large-scale graphs) and the shared usage of it across different functionalities in production.

In light of these drawbacks, we propose a test-time augmentation framework for GNNs, namely GRAPHPATCHER. Given a well-trained GNN, GRAPHPATCHER mitigates the degree bias by patching corrupted ego-graphs with multiple generated virtual neighbors. Notably, GRAPHPATCHER not only enhances the performance of low-degree nodes but also maintains (sometimes improves) GNNs performance on high-degree nodes. This behavior is empirically important because practitioners can universally apply GRAPHPATCHER to all nodes without, like previous works, manually discovering a degree threshold that differentiates the low- and high-degree nodes. To achieve so, we first generate a sequence of ego-graphs corrupted with increasing strengths. Then, GRAPHPATCHER recursively generates multiple virtual nodes to patch the mostly corrupted graph, such that the frozen GNN gives similar predictions for the patched graph and the corresponding corrupted ego-graph in the sequence. Through this scheme, GRAPHPATCHER not only learns how to patch low-degree nodes (i.e., heavily corrupted) but also maintains GNNs original superior performance on high-degree nodes (i.e., lightly corrupted). As a test-time augmentation framework, GRAPHPATCHER is parameterized in parallel with the target GNN. Hence, GRAPHPATCHER is model-agnostic and requires no updates on the target GNN, enabling practitioners to easily utilize it as a plug-and-play module to existing well-established infrastructures. Overall, our contributions are summarized as:

- We study a more practical setting of degree biases on graphs, where both the performances on low- and high-degree nodes are considered. In this case, a good framework is required to not only improve the performance over low-degree nodes but also maintain the original superior performance over high-degree nodes. We evaluate existing frameworks in this setting and observe that many of them trade off performance on high-degree nodes for that on low-degree nodes.
- To mitigate degree biases, we propose GRAPHPATCHER, a novel test-time augmentation framework for graphs. Given a well-trained GNN, GRAPHPATCHER iteratively generates multiple virtual nodes and uses them to patch the original ego-graphs. These patched ego-graphs not only improve GNNs’ performance on low-degree nodes but also maintains that over high-degree nodes. Moreover, GRAPHPATCHER is applied at the testing time for GNNs, a plug-and-play module that is easily applicable to existing well-established infrastructures.
- We conduct extensive evaluation of GRAPHPATCHER along with six state-of-the-art frameworks that mitigate degree biases on seven benchmark datasets. GRAPHPATCHER consistently enhances the overall performance by up to 3.6% and low-degree performance by up to 6.5% of multiple GNNs, significantly outperforming state-of-the-art baselines.

## 2 Related Works

**Graph Neural Networks.** Graph Neural Networks (GNNs) have become one of the most popular paradigms for learning representations over graphs [22, 38, 9, 43, 23, 17, 4]. GNNs aim at mapping the input nodes into low-dimensional vectors, which can be further utilized to conduct either graph-level or node-level tasks. Most GNNs explore a layer-wise message passing scheme, where a node iteratively extracts information from its first-order neighbors, and information from multi-hop neighbors can be captured by stacked layers. They achieved state-of-the-art performance on various tasks, such as node classification [22, 44, 12, 35], link prediction [50, 53, 8], node clustering [2, 37], etc. These tasks further form the archetypes of many real-world applications, such as recommendation systems [45, 3], predictive user behavior models [31, 52], question answering [18], and molecular property prediction [51, 48, 7, 24].

**Degree Bias underlying GNNs.** Recent studies have shown that GNNs exhibit strong biases towards the node degree: they usually perform satisfactorily over high-degree nodes with rich neighbor information but suffer over low-degree nodes [13, 25, 36, 55]. Existing frameworks that mitigate degree biases derive either designated architectures or training strategies specifically for low-degree nodes. For instance, Tail-GNN [25] enhances low-degree nodes’ latent representations by injecting high-degree structural information learned from high-degree nodes; Coldbrew [55] retrieves a set of existing nodes as virtual neighbors for low-degree nodes; TuneUp [13] fine-tunes the well-trained GNNs with pseudo labels and heavily corrupted graphs. Though effective for low-degree nodes, they unintentionally create an artificial out-of-distribution scenario [42], where models only observe low-degree nodes during the training, leading to downgraded performance for high-degree nodes that GNNs originally perform well at.

**Test-time Augmentation.** While data augmentations during the training phase have become one of the essential ingredients for training machine learning models [54], the augmentation applied during the testing time is far less studied, especially for the graph learning community. It has been moderately researched in the computer vision field, aimed at improving performance or mitigating uncertainties [34, 21, 41, 1]. They usually corrupt the same sample by different augmentation approaches and aggregate the model’s predictions on all corrupted samples. Whereas in the graph community, GTrans [15] proposes a test-time enhancement framework, where the node feature and graph topology are modified at the test time to mitigate potential out-of-distribution scenarios.

## 3 Methodology

### 3.1 Preliminary

In this work, we specifically focus on the node classification task. Let  $G = (V, E)$  denote a graph, where  $V$  is the set of  $|V| = N$  nodes and  $E \subseteq V \times V$  is the set of  $|E|$  edges between nodes.  $\mathbf{X} \in \mathbb{R}^{N \times d}$  represents the feature matrix, with  $i$ -th row representing node  $v_i$ ’s  $d$ -dimensional feature vector.  $\mathbf{Y} \subseteq \{0, 1\}^{N \times C}$  denotes the label matrix, where  $C$  is the number of total classes. And  $\mathbf{Y}^{(L)}$  denotes the label matrix for training nodes. We denote the ego-graph of node  $v_i$  is defined as  $\mathcal{G}(v_i) = (V_i, E_i)$  with  $V_i = \mathcal{N}_k(v_i)$ , where  $\mathcal{N}_k(v_i)$  stands for all nodes within the  $k$ -hop neighborhood of  $v_i$  including itself and  $E_i$  refers to the edges in-between  $\mathcal{N}_k(v_i)$ . A well-trained GNN  $f_g(\cdot; \theta) : G \rightarrow \mathbb{R}^{N \times C}$  parameterized by  $\theta$  takes  $G$  as input and maps every node in  $G$  to a  $C$ -dimensional class distribution. Formally, we define test-time node patching as the following:

**Definition 1** (Test-time Node Patching). *Given a GNN  $f_g(\cdot; \theta)$  and a graph  $G$ , a test-time node patching framework  $f(\cdot; \phi) : G \rightarrow G$  takes  $G$  and outputs the patched graph  $\hat{G}$  with generated nodes and edges, such that the performance of  $f_g$  over nodes in  $G$  is enhanced when  $\hat{G}$  is utilized:*

$$\arg \min_{\phi} \mathcal{L}\left(f_g(f(G; \phi); \theta^*), \mathbf{Y}\right), \quad \text{where } \theta^* = \arg \min_{\theta} \mathcal{L}(f_g(G; \theta), \mathbf{Y}^{(L)}), \quad (1)$$

where  $\mathcal{L}$  refers to the loss function evaluating the GNN (e.g., cross-entropy or accuracy).

In this work, we aim at mitigating the degree bias via test-time node patching. To achieve so, two challenges need to be addressed: (1) how to optimize and formulate  $f(\cdot; \phi)$ , such that the graphs patched by  $f(\cdot; \phi)$  enhance the performance of  $f_g(\cdot; \theta^*)$  over low-degree nodes; and (2) how to derive a unified learning scheme that allows  $f(\cdot; \phi)$  to not only improve low-degree nodes but also maintain the GNN’s original superiority over high-degree nodes.

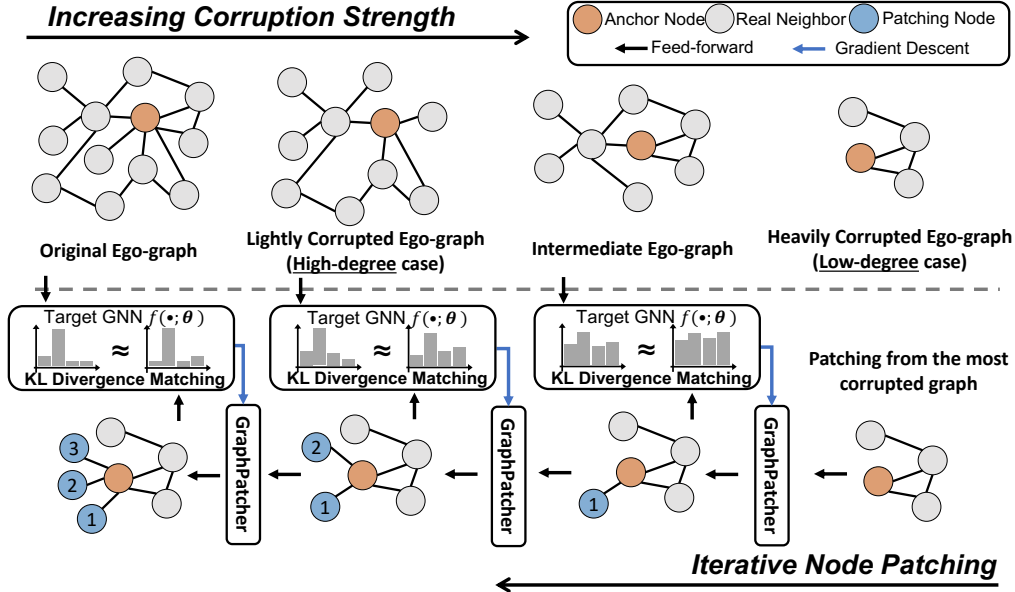


Figure 2: GRAPHPATCHER is presented ego-graphs corrupted by increasing strengths (i.e., the top half of the figure). From the most corrupted graph, it iteratively generates patching nodes to the anchor node, such that the target GNN behaves similarly given the currently patched graph or the corrupted graph next in the hierarchy (i.e., the bottom half of the figure).

### 3.2 The Proposed Framework: GRAPHPATCHER

Our proposed GRAPHPATCHER is a test-time augmentation framework for GNNs to mitigate their degree biases. As shown in Figure 2, GRAPHPATCHER is presented a sequence of ego-graphs corrupted by increasing strengths. Starting from the most corrupted graphs, GRAPHPATCHER iteratively generates patching nodes to augment the anchor nodes. Compared with the corrupted graphs next in the hierarchy, the patched graphs should allow the target GNN to deliver similar outputs. Through this scheme, GRAPHPATCHER not only learns how to patch low-degree nodes while preserving the superior performance over high-degree nodes.

#### 3.2.1 Patching Ego-graphs via Prediction Reconstruction

In order to patch low-degree nodes, a straightforward approach is to corrupt high-degree nodes into low-degree nodes, and allowing the learning model to patch the corrupted nodes to restore their original properties [25, 13]. However, patching low-degree nodes not only affects their own representations but also those of their neighbors due to the message-passing mechanism of GNNs as well as the non-i.i.d. property of nodes in a graph. Besides, modeling over the entire graphs requires the learning model to consider all potential circumstances, whose overheads grow quadratically w.r.t. the number of nodes. Consequently, it becomes challenging to simultaneously determine both features and neighbors of the patching nodes given the entire graph.

To reduce the complexity of the optimization process, instead of working over the entire graph, we conduct node patching over ego-graphs and regard each ego-graph as an i.i.d. sample of the anchor node [56, 19]. For each node  $v_i$ , we have  $f_g(G; \theta)[v_i] = f_g(\mathcal{G}(v_i); \theta)[v_i]$  if  $k$  equals to the number of layers in  $f_g(\cdot; \theta)$ . To further simplify the optimization process, we directly wire the generated virtual nodes to the anchor node (i.e., the generated virtual nodes are the first-order neighbors of the anchor node). This implementation is simple yet effective, because we no longer consider the location to place the patching node: any modification that affects the latent representation of the anchor node can be achieved by patching nodes (with different features) directly to the anchor nodes.

We start explaining GRAPHPATCHER by the most basic case where we only conduct node patching once. Specifically, given the a trained GNN  $f_g(\cdot; \theta^*)$ , an anchor node  $v_i$ , and a corruption function  $\mathcal{T}(\cdot; t)$  with strength  $t$  (i.e., first-order neighbor dropping with probability  $t$  to simulate a low-degree

scenario), that is,  $\mathcal{G}'(v_i) = (V'(v_i), E'(v_i)) = \mathcal{T}(\mathcal{G}(v_i), t)$ . GRAPHPATCHER  $f(\cdot; \phi)$  takes the corrupted ego-graph  $\mathcal{G}'(v_i)$  as input and outputs the augmented ego-graph  $\hat{\mathcal{G}}(v_i)$  with a patching node  $v_p$  and its feature  $\mathbf{x}_p$ , which is directly connected to  $v_i$ . That is,

$$\hat{\mathcal{G}}(v_i) = f(\mathcal{G}'(v_i); \phi), \text{ where } \hat{V} = V'(v_i) \cup \{v_p\}, \hat{E} = E'(v_i) \cup \{e_{(i,p)}\}, \quad (2)$$

where  $e_{(i,p)}$  refers to the edge connecting  $v_i$  and  $v_p$  and  $V'(v_i)$  and  $E'(v_i)$  refer to the nodes and edges in  $\mathcal{G}'(v_i)$ , respectively. To optimize  $f(\cdot; \phi)$  such that  $f_g(\cdot; \theta^*)$  gives similar predictions to  $\hat{\mathcal{G}}(v_i)$  and  $\mathcal{G}(v_i)$ , we minimize the Kullback–Leibler divergence between the frozen GNN’s predictions on these two ego-graphs, which is defined as:

$$\arg \min_{\phi} \sum_{v_i \in V_{\text{tr}}} \text{KL-Div} \left( f_g(\mathcal{G}(v_i); \theta^*)[v_i], f_g(f(\mathcal{G}'(v_i); \phi); \theta^*)[v_i] \right), \quad (3)$$

where  $\text{KL-Div}(\mathbf{y}_1, \mathbf{y}_2) = (\mathbf{y}_1 + \epsilon) \cdot (\log(\mathbf{y}_2 + \epsilon) - \log(\mathbf{y}_1 + \epsilon))^1$  with  $\epsilon > 0$  and  $V_{\text{tr}}$  refers to the set of anchor nodes for training. Intuitively, the reconstruction process above enforces GRAPHPATCHER to remedy the corrupted neighborhood caused by  $\mathcal{T}(\cdot; t)$  via adding a patching node directly to the anchor node. It is philosophically similar to the existing works (e.g., TuneUp [13] and Tail-GNN [25]), where models gain better generalization over low-degree nodes via the corrupted high-degree nodes. Empirically, we observe that this branch of approaches can effectively enhance performance over low-degree nodes. Though promising, according to our empirical studies, it falls short on the high-degree node that original GNNs perform well at. This phenomenon may be attributed to the unintentionally created out-of-distribution scenario [42], wherein models primarily encounter nodes with low degrees during the training. Consequently, the performance of GNNs, which is typically proficient with high-degree nodes, is adversely affected and downgraded.

### 3.2.2 Iterative Patching to Mitigate Degree Bias

In this work, we emphasize that: *mitigating degree bias should not focus specifically on the low-degree nodes: trading off performance on high-degree nodes for that on low-degree nodes simply creates a new bias towards high-degree nodes*. Therefore, besides enhancing the performance on low-degree nodes, maintaining GNN’s original superiority on high-degree nodes is equally critical. This behavior is empirically desirable because practitioners can universally apply GRAPHPATCHER to all nodes without, like previous works do, manually discovering the degree threshold that differentiates the low- and high-degree nodes. Furthermore, the fact that these frameworks are applicable only to low-degree nodes indicates a lack of robustness: further remedying a neighborhood that is informative enough to deliver a good classification result should not jeopardize the performance.

To mitigate the degree bias, we propose a novel training scheme for GRAPHPATCHER such that it observes both low- and high-degree nodes simultaneously during the optimization. Specifically, given a node  $v_i$ , we firstly create a sequence of  $M$  corrupted ego-graphs of  $v_i$ , denoted as  $\mathcal{S}(v_i) = [\mathcal{G}'(v_i)_m = \mathcal{T}(\mathcal{G}(v_i), t_m)]_{m=1}^M$ , with decreasing corruption strength (i.e.,  $\forall m, n \in \{1, \dots, M\}$ ,  $t_m > t_n$  if  $m < n$ ). Instead of the one-step patching to match the prediction on the original ego-graph as described in Section 3.2.1, GRAPHPATCHER traverses  $\mathcal{S}(v_i)$  and recursively patches the corrupted ego-graph to match the target GNN’s prediction on the ego-graph next in the sequence. As also illustrated in Figure 2, this optimization process is formulated as:

$$\arg \min_{\phi} \sum_{v_i \in V_{\text{tr}}} \sum_{m=1}^{M-1} \text{KL-Div} \left( f_g(\mathcal{G}'(v_i)_{m+1}; \theta^*)[v_i], f_g(\hat{\mathcal{G}}(v_i)_m; \theta^*)[v_i] \right), \quad (4)$$

$$\text{s.t. } \hat{\mathcal{G}}(v_i)_m = f(\hat{\mathcal{G}}(v_i)_{m-1}; \phi),$$

where  $\hat{\mathcal{G}}(v_i)_m = (\hat{V}_m, \hat{E}_m)$  with  $\hat{V}_m = V'_1(v_i) \cup \{v_p\}_{p=1}^m$ ,  $\hat{E}_m = E'_1(v_i) \cup \{e_{(i,p)}\}_{p=1}^m$ , and  $\hat{\mathcal{G}}(v_i)_0 = \mathcal{G}'(v_i)_1$ .

The one-step patching described in Section 3.2.1 remedies low-degree anchor nodes directly to the distributions of high-degree nodes. During this process, the model does not observe distributions of high-degree nodes and hence delivers sub-optimal performance. Therefore, we design GRAPHPATCHER to be an iterative multi-step framework. At each step, it takes the previously patched

<sup>1</sup>KL divergence used here is equal to the regularized cross-entropy. It is strongly convex and Lipschitz continuous due to the incorporation of  $\epsilon$ . These two properties are required for the derivation of Theorem 1.



ego-graph as input and further remedies the partially patched ego-graph to match the GNN’s prediction on the ego-graph next in the sequence. This scheme enables GRAPHPATCHER to learn to patch low-degree nodes in early steps when the ego-graphs are heavily corrupted (e.g., low-degree case in [Figure 2](#)) and maintain the original performance in later steps when ego-graphs are lightly corrupted (e.g., high-degree case in [Figure 2](#)). Specifically, at the  $m$ -th patching step, the currently patched ego-graph  $\hat{\mathcal{G}}(v_i)_m$  reflects the neighbor distribution of ego-graphs corrupted by a specific strength of  $t_{m+1}$ . GRAPHPATCHER takes  $\hat{\mathcal{G}}(v_i)_m$  as input and further generates another patching node  $v_{m+1}$  to approach the neighbor distribution of ego-graphs corrupted by a slightly weaker strength of  $t_{m+2}$ . This process iterates until GRAPHPATCHER traverses  $\mathcal{S}(v_i)$ . Intuitively, the incorporation of  $v_{m+1}$  enriches the neighbor distribution by an amount of  $t_{m+2} - t_{m+1}$  corruption strength. This optimization scheme allows GRAPHPATCHER to observe neighbor distributions with varying corruption strengths and makes our proposal applicable to both low- and high-degree nodes.

However, the target distribution at each step (i.e.,  $f_g(\mathcal{G}'(v_i)_{m+1}; \theta)[v_i]$  in [Equation \(4\)](#)) is not deterministic due to the stochastic nature of the corruption function  $\mathcal{T}$ . Given an ego-graph  $\mathcal{G}(v_i)$  and a corruption strength  $t$ , one can at most generate  $\binom{|V_i|}{(1-t)|V_i|}$  different corrupted ego-graphs. With a large corruption strength (e.g., ego-graphs early in the sequence  $\mathcal{S}(v_i)$ ), two corrupted ego-graphs generated by the same exact priors might exhibit completely different topologies. Such differences could bring high variance to the supervision signal and instability to the optimization process. To alleviate the issue above, at each step we sample  $L$  ego-graphs with the same corruption strength and let GRAPHPATCHER approximate multiple predictions over them, formulated as:

$$\mathcal{L}_{\text{patch}} = \sum_{v_i \in V_{\text{tr}}} \sum_{m=1}^{M-1} \sum_{l=1}^L \text{KL-Div} \left( f_g(\mathcal{G}'(v_i)_{m+1}^l; \theta^*)[v_i], f_g(\hat{\mathcal{G}}(v_i)_m; \theta^*)[v_i] \right), \quad (5)$$

where  $\hat{\mathcal{G}}(v_i)_m = f(\hat{\mathcal{G}}(v_i)_{m-1}; \phi)$  and  $\mathcal{G}'(v_i)_{m+1}^l$  refers to one of the  $L$  target corrupted ego-graphs that GRAPHPATCHER aims to approximate at the  $m$ -th step. This approach allows GRAPHPATCHER to patch the anchor node towards a well-approximated region where its high-degree counterparts should locate, instead of one point randomly sampled from this region.

With  $M - 1$  virtual nodes patched to the ego-graph, we further ask GRAPHPATCHER to generate a last patching node to  $\hat{\mathcal{G}}(v_i)_{M-1}$  and enforce the resulted graph  $\hat{\mathcal{G}}'(v_i)_M$  to match the GNN’s prediction on the original ego-graph. The last patching node could be regarded as a slack variable to complement minor differences between the original and the least corrupted ego-graphs, formulated as:

$$\mathcal{L}_{\text{recon}} = \sum_{v_i \in V_{\text{tr}}} \text{KL-Div} \left( f_g(\mathcal{G}(v_i); \theta^*)[v_i], f_g(\hat{\mathcal{G}}'(v_i)_M; \theta^*)[v_i] \right), \quad (6)$$

where  $\hat{\mathcal{G}}'(v_i)_M = f(\hat{\mathcal{G}}(v_i)_{M-1}; \phi)$ .  $\mathcal{L}_{\text{recon}}$  ([Equation \(6\)](#)) also prevents GRAPHPATCHER from overfitting to the low-degree nodes and enforces GRAPHPATCHER to maintain the target GNN’s performance over high-degree nodes, since only marginal distribution modification should be expected with this last patching node. Hence, GRAPHPATCHER is optimized by a linear combination of the above two objectives (i.e.,  $\arg \min_{\phi} \mathcal{L}_{\text{patch}} + \mathcal{L}_{\text{recon}}$ ).

### 3.2.3 Theoretical Analysis

As shown in [Equation \(5\)](#), one of the important factors that contribute to the success of GRAPHPATCHER is sampling multiple ego-graphs with the same corruption strength. The following theorem shows that the error is bounded w.r.t. the number of sampled ego-graphs  $L$ .

**Theorem 1.** *Assuming the parameters of GRAPHPATCHER are initialized from the set  $P_{\beta} = \{\phi : \|\phi - \mathcal{N}(\mathbf{0}_{|\phi|}; \mathbf{1}_{|\phi|})\|_F < \beta\}$  where  $\beta > 0$ , with probability at least  $1 - \delta$  for all  $\phi \in P_{\beta}$ , the error (i.e.,  $\mathbb{E}(\mathcal{L}_{\text{patch}}) - \mathcal{L}_{\text{patch}}$ ) is bounded by  $\mathcal{O}(\beta \sqrt{\frac{|\phi|}{L}} + \sqrt{\frac{\log(1/\delta)}{L}})$ .*

The proof of [Theorem 1](#) is provided in [Appendix C](#). From the above theorem, we note that without the sampling strategy (i.e.,  $L = 1$ ), the generalization error depends only on the number of parameters (i.e.,  $|\phi|$ ) given the same objective function, which could lead to high variance to the supervision signal and instability to the optimization process. According to this theorem and our empirical observation, an affordable value of  $L$  (e.g.,  $L = 10$ ) delivers stable results across datasets.

## 4 Experiments

### 4.1 Experimental Setting

**Datasets.** We conduct comprehensive experiments on seven real-world benchmark datasets that are broadly utilized by the graph community, including Cora, Citeseer, Pubmed, Wiki.CS, Amazon-Photo, Coauthor-CS, ogbn-arxiv, Actor, and Chameleon [44, 28, 12, 33]. This list of datasets covers graphs with distinctive characteristics (i.e., graphs with different domains and dimensions) to fully evaluate the effectiveness of GRAPHPATCHER. The detail of these datasets can be found in [Appendix A](#).

**Baselines.** We compare GRAPHPATCHER with six state-of-the-art graph learning frameworks from three branches. The first branch specifically aims at enhancing the performance on low-degree nodes, including TAIL-GNN [25], COLBBREW [55], and TUNEUP [13]. The second branch consists of frameworks that focus on handling out-of-distribution scenarios, including EERM [42] and GTRANS [15]. We list this branch of frameworks as baselines because the sub-optimal performance of GNNs over low-degree nodes could be regarded as an out-of-distribution scenario. As GRAPHPATCHER is a test-time augmentation framework, the last branch of baseline includes DROPEdge, which is a data augmentation framework employed during training.

**Evaluation Protocol.** We evaluate all models using the node classification task [22, 38], quantified by the accuracy score. For datasets with publicly available (i.e., ogbn-arxiv, Cora, Citeseer, and Pubmed), we employ their the provided splits for the model training and testing. Whereas for other datasets, we create a random 10%/10%/80% split for the training/validation/testing split, to simulate a semi-supervised learning setting. All reported performance is averaged over 10 independent runs with different random seeds. Both mean values and standard deviations for the performances of all models are reported. Besides mitigating the degree bias for supervised GNNs, GRAPHPATCHER is also applicable to self-supervised GNNs. To evaluate the model performance for them, we apply GRAPHPATCHER and TUNEUP to state-of-the-art self-supervised GNNs including DGI [39], GRACE [57], and PARETOGNN [20]. We only compare our proposal with TUNEUP since other frameworks require specific model architectures and hence do not apply to self-supervised GNNs.

**Hyper-parameters.** We use the optimal settings on all baselines given by the authors for the shared datasets and a simple two-layer GCN [22] as the backbone model architecture for all applicable baselines. Hyper-parameters we tune for GRAPHPATCHER include learning rate, hidden dimension, the augmentation strength at each step, and the total amount of patching steps with details described in [Appendix B](#). Besides, all of our models are trained on a single RTX3090 with 24GB VRAM; additional hardware information can also be found in the appendix.

### 4.2 Performance Comparison with Baselines

We compare GRAPHPATCHER with six state-of-the-art frameworks that mitigate the degree bias problem and the performances of all models are shown in [Table 1](#). Firstly we notice that the problem of degree bias is quite serious across datasets for GCN. The performances on low-degree nodes are  $\sim 10\%$  lower than those over high-degree nodes. Comparing GCN with COLBBREW, TAIL-GNN, and TUNEUP, we can observe that frameworks that focus specifically on low-degree nodes can usually enhance GNN’s performance over the lower percentile (e.g., 1.2% accuracy gain on Cora by TUNEUP, 0.74% on Citeseer by COLBBREW, 1.38% on Pubmed by TAIL-GNN, etc.). However, these frameworks fall short on the high-degree nodes and sometimes perform worse than the vanilla GCN (e.g., -2.7% accuracy degradation on Cora by TAIL-GNN, -11.42% on Wiki.CS by COLBBREW, and -2.5% on Amazon Photo by TUNEUP). This phenomenon could result from that they unintentionally create an artificial out-of-distribution scenario, where they only observe low-degree nodes during the training, leading to downgraded performance for high-degree nodes that GNNs originally perform well at. Comparing GCN with GTRANS and EERM, we observe that they deliver similar performances as the vanilla GCN does, indicating that frameworks targeting out-of-distribution scenarios cannot mitigate degree biases. Comparing GRAPHPATCHER with all baselines, we notice that our proposed GRAPHPATCHER consistently improves the low-degree performance with an average improvement gain of 2.23 accuracy score. Besides, unlike other frameworks that have downgraded performance over high-degree nodes, GRAPHPATCHER can maintain GCN’s original high-degree superiority, due to our iterative node patching. On average, GRAPHPATCHER improves GCN’s overall performance by a 1.4 accuracy score across datasets.

Table 1: Performance (%) of all models over nodes with different degrees. Lower and upper percentile indicate the set of nodes whose degree is ranked in the lower and upper 33% population respectively. A two-layer GCN is used as the backbone model for all applicable baselines. **Bold** indicates the best performance and underline indicates the runner-up, with standard deviations as subscripts.

Method	Cora	Citeseer	Pubmed	Wiki_CS	Am.Photo	Co_CS	Arxiv	Chameleon	Actor
ACCURACY ON LOW-DEGREE NODES (LOWER PERCENTILE)									
GCN	73.27 $\pm$ 0.01	64.86 $\pm$ 0.92	76.88 $\pm$ 0.40	72.98 $\pm$ 0.50	75.59 $\pm$ 0.43	84.59 $\pm$ 0.45	63.15 $\pm$ 0.13	54.05 $\pm$ 0.18	27.30 $\pm$ 0.52
COLBBREW	73.82 $\pm$ 0.98	65.60 $\pm$ 0.08	77.72 $\pm$ 0.63	73.98 $\pm$ 0.52	76.18 $\pm$ 0.80	85.56 $\pm$ 0.69	63.02 $\pm$ 0.21	53.41 $\pm$ 0.22	27.88 $\pm$ 0.13
TAIL-GNN	71.17 $\pm$ 0.80	57.66 $\pm$ 0.83	75.38 $\pm$ 0.89	<b>74.36</b> $\pm$ 0.18	77.22 $\pm$ 0.94	85.13 $\pm$ 0.60	OOM	53.48 $\pm$ 0.04	27.80 $\pm$ 0.62
TUNEUP	74.47 $\pm$ 0.34	65.17 $\pm$ 0.22	77.18 $\pm$ 0.39	72.60 $\pm$ 0.75	76.08 $\pm$ 0.62	84.68 $\pm$ 0.50	63.34 $\pm$ 0.32	53.87 $\pm$ 0.43	27.94 $\pm$ 0.14
EERM	73.40 $\pm$ 0.06	64.27 $\pm$ 0.33	76.30 $\pm$ 0.20	73.12 $\pm$ 0.68	75.15 $\pm$ 0.59	84.82 $\pm$ 0.74	63.20 $\pm$ 0.11	54.11 $\pm$ 0.32	27.48 $\pm$ 0.39
GTRANS	73.16 $\pm$ 0.66	64.95 $\pm$ 0.83	77.05 $\pm$ 1.00	72.15 $\pm$ 0.50	75.55 $\pm$ 0.55	84.74 $\pm$ 0.06	62.88 $\pm$ 0.14	54.29 $\pm$ 0.14	27.53 $\pm$ 0.21
DROPEdge	73.57 $\pm$ 0.97	65.47 $\pm$ 0.27	75.68 $\pm$ 0.82	73.94 $\pm$ 0.20	76.49 $\pm$ 0.03	84.31 $\pm$ 0.33	61.33 $\pm$ 0.33	54.12 $\pm$ 0.41	27.39 $\pm$ 0.24
GRAPHPATCHER	<b>78.08</b> $\pm$ 0.06	<b>67.27</b> $\pm$ 0.20	<b>78.98</b> $\pm$ 0.21	74.04 $\pm$ 0.86	<b>77.84</b> $\pm$ 0.36	<b>86.76</b> $\pm$ 0.84	<b>64.01</b> $\pm$ 0.12	<b>54.48</b> $\pm$ 0.71	<b>29.27</b> $\pm$ 0.57
ACCURACY ON HIGH-DEGREE NODES (UPPER PERCENTILE)									
GCN	86.83 $\pm$ 0.17	77.25 $\pm$ 1.00	80.84 $\pm$ 0.76	83.40 $\pm$ 0.70	84.07 $\pm$ 0.71	90.20 $\pm$ 0.37	80.46 $\pm$ 0.18	54.11 $\pm$ 0.73	27.41 $\pm$ 0.29
COLBBREW	84.80 $\pm$ 0.04	75.33 $\pm$ 0.84	78.66 $\pm$ 0.38	71.98 $\pm$ 0.95	77.07 $\pm$ 0.14	82.16 $\pm$ 0.39	70.57 $\pm$ 0.36	53.72 $\pm$ 0.48	26.67 $\pm$ 0.29
TAIL-GNN	84.13 $\pm$ 0.48	74.85 $\pm$ 0.30	78.74 $\pm$ 0.34	78.91 $\pm$ 0.97	80.32 $\pm$ 0.60	86.75 $\pm$ 0.90	OOM	54.53 $\pm$ 0.12	27.13 $\pm$ 0.44
TUNEUP	87.13 $\pm$ 0.67	76.95 $\pm$ 0.63	81.74 $\pm$ 0.49	83.11 $\pm$ 0.53	81.57 $\pm$ 0.07	90.65 $\pm$ 0.86	80.09 $\pm$ 0.51	54.25 $\pm$ 0.59	26.64 $\pm$ 0.71
EERM	85.89 $\pm$ 0.09	76.32 $\pm$ 0.23	79.98 $\pm$ 0.06	82.98 $\pm$ 0.07	84.32 $\pm$ 0.96	90.17 $\pm$ 0.11	80.37 $\pm$ 0.12	54.41 $\pm$ 0.71	27.39 $\pm$ 0.14
GTRANS	86.32 $\pm$ 0.34	76.60 $\pm$ 0.44	80.56 $\pm$ 0.92	83.42 $\pm$ 0.04	83.95 $\pm$ 0.99	89.99 $\pm$ 0.10	<b>80.77</b> $\pm$ 0.26	54.21 $\pm$ 0.19	27.29 $\pm$ 0.12
DROPEdge	86.53 $\pm$ 0.99	76.35 $\pm$ 0.17	81.44 $\pm$ 0.51	83.37 $\pm$ 0.43	<b>84.97</b> $\pm$ 0.56	89.28 $\pm$ 0.08	80.64 $\pm$ 0.36	54.17 $\pm$ 0.11	27.38 $\pm$ 0.21
GRAPHPATCHER	<b>88.02</b> $\pm$ 0.11	76.65 $\pm$ 0.18	<b>83.83</b> $\pm$ 0.79	<b>83.49</b> $\pm$ 0.22	84.17 $\pm$ 0.97	<b>90.59</b> $\pm$ 0.46	80.61 $\pm$ 0.25	54.20 $\pm$ 0.21	<b>27.43</b> $\pm$ 0.62
OVERALL PERFORMANCE									
GCN	81.22 $\pm$ 0.40	70.51 $\pm$ 0.46	79.14 $\pm$ 0.31	77.30 $\pm$ 0.41	80.38 $\pm$ 0.86	88.16 $\pm$ 0.66	71.73 $\pm$ 0.14	52.83 $\pm$ 0.35	27.20 $\pm$ 0.57
COLBBREW	80.70 $\pm$ 0.86	70.10 $\pm$ 0.55	78.66 $\pm$ 0.93	73.82 $\pm$ 0.69	78.24 $\pm$ 0.62	85.80 $\pm$ 0.79	63.55 $\pm$ 0.48	52.12 $\pm$ 0.53	26.75 $\pm$ 0.32
TAIL-GNN	79.44 $\pm$ 0.64	65.80 $\pm$ 0.04	76.14 $\pm$ 0.25	74.66 $\pm$ 0.18	80.68 $\pm$ 0.58	87.02 $\pm$ 0.33	OOM	52.46 $\pm$ 0.12	27.62 $\pm$ 0.47
TUNEUP	82.11 $\pm$ 0.39	70.92 $\pm$ 0.02	79.91 $\pm$ 0.26	76.93 $\pm$ 0.81	79.74 $\pm$ 0.28	88.46 $\pm$ 0.97	71.51 $\pm$ 0.30	52.89 $\pm$ 0.41	27.32 $\pm$ 0.64
EERM	81.47 $\pm$ 0.19	70.08 $\pm$ 0.19	78.65 $\pm$ 0.43	77.29 $\pm$ 0.96	79.79 $\pm$ 0.61	88.07 $\pm$ 0.30	71.70 $\pm$ 0.18	52.93 $\pm$ 0.24	27.65 $\pm$ 0.29
GTRANS	80.79 $\pm$ 0.51	69.51 $\pm$ 0.93	78.67 $\pm$ 0.93	76.39 $\pm$ 0.27	80.02 $\pm$ 0.80	88.06 $\pm$ 0.96	71.77 $\pm$ 0.19	52.67 $\pm$ 0.35	26.93 $\pm$ 0.41
DROPEdge	81.10 $\pm$ 0.31	71.10 $\pm$ 0.86	78.90 $\pm$ 0.68	77.49 $\pm$ 0.78	81.11 $\pm$ 0.72	87.56 $\pm$ 0.64	71.82 $\pm$ 0.33	52.89 $\pm$ 0.22	27.28 $\pm$ 0.32
GRAPHPATCHER	<b>84.17</b> $\pm$ 0.54	<b>71.65</b> $\pm$ 0.05	<b>81.13</b> $\pm$ 0.68	<b>78.12</b> $\pm$ 0.57	<b>81.23</b> $\pm$ 0.32	<b>89.44</b> $\pm$ 0.79	<b>72.31</b> $\pm$ 0.22	<b>53.21</b> $\pm$ 0.39	<b>28.34</b> $\pm$ 0.24

Table 2: Performance (%) of GRAPHPATCHER and TUNEUP for different GNN architectures.

Method	Cora	Citeseer	Pubmed	Wiki_CS	Am.Photo	Co_CS	Arxiv
ACCURACY ON LOW-DEGREE NODES (LOWER PERCENTILE)							
GCN	73.27 $\pm$ 0.01	64.86 $\pm$ 0.92	76.88 $\pm$ 0.40	72.98 $\pm$ 0.50	75.59 $\pm$ 0.43	84.59 $\pm$ 0.45	63.15 $\pm$ 0.13
+TUNEUP	74.47 $\pm$ 0.34	65.17 $\pm$ 0.22	77.18 $\pm$ 0.39	72.60 $\pm$ 0.75	76.08 $\pm$ 0.62	84.68 $\pm$ 0.50	63.34 $\pm$ 0.32
+GRAPHPATCHER	<b>78.08</b> $\pm$ 0.06	<b>67.27</b> $\pm$ 0.20	<b>78.98</b> $\pm$ 0.21	<b>74.04</b> $\pm$ 0.86	<b>77.84</b> $\pm$ 0.36	<b>86.76</b> $\pm$ 0.84	<b>64.01</b> $\pm$ 0.12
G-SAGE	70.57 $\pm$ 0.84	67.44 $\pm$ 0.11	76.58 $\pm$ 0.36	<b>61.83</b> $\pm$ 0.89	76.32 $\pm$ 0.33	74.53 $\pm$ 0.69	61.64 $\pm$ 0.62
+TUNEUP	71.47 $\pm$ 0.11	67.44 $\pm$ 0.21	77.78 $\pm$ 0.92	58.46 $\pm$ 0.77	<b>78.48</b> $\pm$ 0.89	74.88 $\pm$ 0.76	62.43 $\pm$ 0.59
+GRAPHPATCHER	<b>72.33</b> $\pm$ 0.21	<b>67.89</b> $\pm$ 0.33	<b>78.93</b> $\pm$ 0.05	61.46 $\pm$ 0.40	78.11 $\pm$ 0.13	<b>75.14</b> $\pm$ 0.21	<b>62.92</b> $\pm$ 0.28
GAT	73.27 $\pm$ 0.51	69.07 $\pm$ 0.11	72.37 $\pm$ 0.27	73.72 $\pm$ 0.64	79.66 $\pm$ 0.58	86.93 $\pm$ 0.50	63.54 $\pm$ 0.20
+TUNEUP	76.58 $\pm$ 0.07	66.67 $\pm$ 0.33	72.07 $\pm$ 0.81	72.08 $\pm$ 0.20	<b>81.08</b> $\pm$ 0.25	86.72 $\pm$ 0.84	63.71 $\pm$ 0.31
+GRAPHPATCHER	<b>76.88</b> $\pm$ 0.32	<b>70.87</b> $\pm$ 0.78	<b>74.47</b> $\pm$ 0.63	<b>74.26</b> $\pm$ 0.72	80.05 $\pm$ 0.19	<b>89.50</b> $\pm$ 0.93	<b>64.12</b> $\pm$ 0.14
ACCURACY ON HIGH-DEGREE NODES (UPPER PERCENTILE)							
GCN	86.83 $\pm$ 0.17	77.25 $\pm$ 1.00	80.84 $\pm$ 0.76	83.40 $\pm$ 0.70	84.07 $\pm$ 0.71	90.20 $\pm$ 0.37	80.46 $\pm$ 0.18
+TUNEUP	87.13 $\pm$ 0.67	76.95 $\pm$ 0.63	81.74 $\pm$ 0.49	83.11 $\pm$ 0.53	81.57 $\pm$ 0.07	90.65 $\pm$ 0.86	80.09 $\pm$ 0.51
+GRAPHPATCHER	<b>88.02</b> $\pm$ 0.11	76.65 $\pm$ 0.18	<b>83.83</b> $\pm$ 0.79	<b>83.49</b> $\pm$ 0.22	<b>84.17</b> $\pm$ 0.97	<b>90.59</b> $\pm$ 0.46	<b>80.61</b> $\pm$ 0.25
G-SAGE	82.04 $\pm$ 0.01	72.46 $\pm$ 0.70	80.24 $\pm$ 0.12	60.83 $\pm$ 0.18	77.30 $\pm$ 0.56	69.24 $\pm$ 0.55	78.66 $\pm$ 0.08
+TUNEUP	80.84 $\pm$ 0.36	<b>73.95</b> $\pm$ 0.36	81.14 $\pm$ 0.41	60.90 $\pm$ 0.05	<b>79.36</b> $\pm$ 0.89	70.12 $\pm$ 0.25	79.26 $\pm$ 0.51
+GRAPHPATCHER	<b>82.14</b> $\pm$ 0.48	73.22 $\pm$ 0.25	<b>81.66</b> $\pm$ 0.46	<b>61.02</b> $\pm$ 0.44	78.57 $\pm$ 0.14	<b>70.53</b> $\pm$ 0.68	<b>79.91</b> $\pm$ 0.31
GAT	85.33 $\pm$ 0.36	76.65 $\pm$ 0.80	81.14 $\pm$ 0.20	82.21 $\pm$ 0.44	87.84 $\pm$ 0.23	91.33 $\pm$ 0.81	81.37 $\pm$ 0.16
+TUNEUP	86.23 $\pm$ 0.47	<b>76.65</b> $\pm$ 0.33	80.84 $\pm$ 0.02	81.73 $\pm$ 0.36	<b>89.02</b> $\pm$ 0.78	<b>92.00</b> $\pm$ 0.04	81.44 $\pm$ 0.11
+GRAPHPATCHER	<b>86.53</b> $\pm$ 0.37	76.35 $\pm$ 0.39	<b>81.14</b> $\pm$ 0.89	<b>82.34</b> $\pm$ 0.08	87.84 $\pm$ 0.56	91.61 $\pm$ 0.20	<b>81.49</b> $\pm$ 0.15

We further apply GRAPHPATCHER to other GNN architectures (i.e., GraphSAGE [9] and GAT [38]) and compare its performance to TUNEUP. We only compare with TUNEUP since other baselines explore specific model architectures that do not allow a different backbone. From Table 2, we can observe that the issue of degree bias still exists on GAT and GraphSAGE with a performance gap between low- and high-degree nodes around  $\sim 10\%$ . Both TUNEUP and GRAPHPATCHER can improve the performance over low-degree nodes. Specifically, TUNEUP on average improves 0.27 low-degree accuracy for GraphSAGE and 0.40 for GAT across datasets; whereas GRAPHPATCHER improves 1.13 for GraphSAGE and 1.66 for GAT, outperforming TUNEUP by a large margin.



### 4.3 Performance of GRAPHPATCHER for Self-supervised GNNs

To fully demonstrate the effectiveness of GRAPHPATCHER, we also apply our proposal to self-supervised GNNs, as shown in Table 3. We can observe that self-supervised learning can mitigate degree bias by itself, proved by smaller gaps between low- and high-degree nodes than those of semi-supervised GNNs. Combined with GRAPHPATCHER, the degree biases can be further without sacrificing GNN’s original superiority over high-degree nodes. On average, GRAPHPATCHER can enhance the low-degree performance of these three self-supervised GNNs by 1.78, 0.74, and 1.36 accuracy scores respectively.

Table 3: Effectiveness for self-supervised GNNs.

Method	Cora	Pubmed	Wiki .CS
LOW-DEGREE NODES (LOWER PERCENTILE)			
DGI	78.47 $\pm$ 0.37	75.63 $\pm$ 0.82	75.86 $\pm$ 0.61
+GRAPHPATCHER	79.95 $\pm$ 0.53	78.04 $\pm$ 0.97	77.31 $\pm$ 0.91
GRACE	77.81 $\pm$ 0.73	77.80 $\pm$ 0.65	74.31 $\pm$ 0.63
+GRAPHPATCHER	78.53 $\pm$ 0.82	78.49 $\pm$ 0.16	75.12 $\pm$ 0.34
PARETOGNN	78.85 $\pm$ 0.71	78.32 $\pm$ 0.33	74.17 $\pm$ 0.18
+GRAPHPATCHER	79.91 $\pm$ 0.62	79.11 $\pm$ 0.89	76.41 $\pm$ 0.22
HIGH-DEGREE NODES (UPPER PERCENTILE)			
DGI	86.83 $\pm$ 0.82	81.14 $\pm$ 0.28	81.09 $\pm$ 0.81
+GRAPHPATCHER	86.91 $\pm$ 0.10	82.31 $\pm$ 0.53	80.95 $\pm$ 0.19
GRACE	85.03 $\pm$ 0.05	78.74 $\pm$ 0.84	83.91 $\pm$ 0.56
+GRAPHPATCHER	85.12 $\pm$ 0.25	79.58 $\pm$ 0.31	84.12 $\pm$ 0.22
PARETOGNN	87.03 $\pm$ 0.84	80.89 $\pm$ 0.84	81.57 $\pm$ 0.84
+GRAPHPATCHER	87.32 $\pm$ 0.27	80.55 $\pm$ 0.32	81.78 $\pm$ 0.51

### 4.4 Effectiveness of GRAPHPATCHER for Enhancing SoTA Method

We apply GRAPHPATCHER to GRAND [5], a strong GNN that utilizes a random propagation strategy to perform graph data augmentation and significantly improve the node classification performance. The performance improvement brought by GRAPHPATCHER is shown in Table 4. We observe that GRAPHPATCHER can still consistently improve the node classification for GRAND. Specifically, on low-degree nodes, GRAPHPATCHER can improve 1.40, 2.23, and 4.20 accuracy score on Cora, Citeseer, and Pubmed, respectively. Overall, GRAPHPATCHER further enhances the SoTA performance on these three datasets, with an outstanding accuracy score of 85.90, 76.10, and 84.20. The significant gain from GRAPHPATCHER indicates that the effectiveness brought by the test-time augmentation is not overlapped with the data augmentation during the training.

Table 4: Effectiveness for SoTA.

Method	Cora	Citeseer	Pubmed
LOW-DEGREE NODES (LOWER PERCENTILE)			
GRAND	80.18 $\pm$ 0.64	70.57 $\pm$ 0.68	80.48 $\pm$ 0.14
+GRAPHPATCHER	81.58 $\pm$ 0.45	72.73 $\pm$ 0.29	84.68 $\pm$ 0.29
HIGH-DEGREE NODES (UPPER PERCENTILE)			
GRAND	88.32 $\pm$ 0.75	79.64 $\pm$ 0.86	83.53 $\pm$ 0.52
+GRAPHPATCHER	88.92 $\pm$ 0.18	79.54 $\pm$ 0.13	84.43 $\pm$ 0.21
OVERALL PERFORMANCE			
GRAND	85.22 $\pm$ 0.80	74.90 $\pm$ 0.77	82.30 $\pm$ 0.41
+GRAPHPATCHER	85.90 $\pm$ 0.44	76.10 $\pm$ 0.38	84.20 $\pm$ 0.26

### 4.5 Performance w.r.t. the Number of Patching Nodes

To investigate the necessity of patching multiple nodes, we conduct experiments over the number of patching nodes at the test time. As shown in Figure 3, we notice that the overall performance gradually increments as the number of patching nodes increases, demonstrating that multiple patching nodes are required to remedy the incomplete neighborhood of low-degree nodes. Besides, we discover that the performance of GRAPHPATCHER saturates with around four nodes patched, which aligns with our training procedure, where the length of the ego-graph sequence is at most five. Experiments concerning the number of patching nodes during the optimization and the number of sampled ego-graphs per corruption strength (i.e.,  $M$  and  $L$  in Equation (5)) can be found in Appendix B.

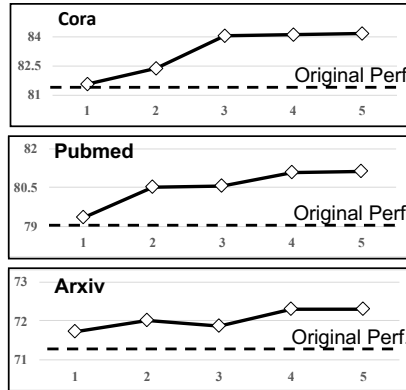


Figure 3: Overall perf. (y-axis) w.r.t. the number of patching nodes (x-axis).

## 5 Discussion w.r.t. Diffusion Models

Both diffusion models and GRAPHPATCHER conduct multiple corruptions to training samples with increasing strengths and generate examples in an iterative fashion. This scheme is conceptually inspired by heat diffusion from physics. However, the motivations behind them are different, where diffusion models focus on the generation quality (i.e., fidelity to the original data distribution) but ours

aims at the results brought by our generated nodes (i.e., the performance improvement). Specifically, diffusion models [10, 32] aim at learning the probability distribution of the data and accordingly generating examples following the learned distribution. Their goal is to generate samples that follow the original data distribution, agnostic of any other factor like the target GNN we have in our scenario. Whereas for GRAPHPATCHER, we aim at generating nodes to ego-nets such that the target GNN models deliver better predictions when the node degree is low. We mostly care about performance improvement and the generated node may be very different from the original nodes in the graph.

## 6 Discussion w.r.t. Generation Methods for Graph

Most graph generation frameworks (including those using diffusion models) explore iterative generation schemes to synthesize real graphs [58, 46, 6, 30, 16, 40]. They improve the generation quality and focus on applications such as molecule design, protein design, and program synthesis. Though GRAPHPATCHER also generates patching nodes for ego-graphs, ours is a different research direction than these methods. We do not focus on whether or not the generated patching nodes are faithful to the original data distribution, as long as the low-degree performance is enhanced and the high-degree performance is maintained. Another relevant work named GPT-GNN [14] explores an iterative node generation for pre-training, which also falls under the category of maintaining the original data distribution. In summary, GRAPHPATCHER is relevant to these frameworks in the sense that it generates nodes to add to ego-graphs. However, our proposal is motivated by a different reason and we aim at the performance improvement brought by generated nodes in downstream tasks.

## 7 Conclusion

We study the problem of degree bias underlying GNNs and accordingly propose a test-time augmentation framework, namely GRAPHPATCHER. GRAPHPATCHER iteratively patches ego-graphs with its generated virtual nodes to remedy the incomplete neighborhood. Through our designated optimization scheme, GRAPHPATCHER not only patches low-degree nodes but also maintains GNN’s original superior performance over high-degree nodes. Comprehensive experiments are conducted over seven benchmark datasets and our proposal can consistently enhance GNN’s overall performance by up to 3.6% and low-degree performance by up to 6.5%, outperforming all baselines by a large margin. Besides, GRAPHPATCHER can also mitigate the degree bias issue for self-supervised GNNs. When applied to graph learning methods with state-of-the-art performance (i.e., GRAND), GRAPHPATCHER can further improve the SoTA performance by a large margin, indicating that the effectiveness brought by the test-time augmentation is not overlapped with existing inductive biases.

## Limitation and Broader Impact

One limitation is the additional overhead entailed by generating ego-graphs. To address this limitation, we generate all ego-graphs before the optimization to avoid duplicated computations. This operation takes more hard-disk storage, which is relatively cheap compared with computational resources. Furthermore, we observe no ethical concern entailed by our proposal, but we note that both ethical or unethical applications based on graphs may benefit from the effectiveness of our work. Care should be taken to ensure socially positive and beneficial results of machine learning algorithms.

## Acknowledgement

We appreciate Shifu Hou from University of Notre Dame for valuable discussions and suggestions. We would also like to thank anonymous reviewers for their constructive suggestions and comments (i.e., experiments over heterophilic datasets, connections to diffusion models, and discussion w.r.t. iterative generation models for graphs). This work is partially supported by the NSF under grants IIS-2334193, IIS-2321504, IIS-2203262, IIS-2214376, IIS-2217239, OAC-2218762, CNS-2203261, and CMMI-2146076. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of any funding agencies.

## References

- [1] Murat Seckin Ayhan and Philipp Berens. Test-time data augmentation for estimation of heteroscedastic aleatoric uncertainty in deep neural networks. In *Medical Imaging with Deep Learning*, 2018.
- [2] Filippo Maria Bianchi, Daniele Grattarola, and Cesare Alippi. Spectral clustering with graph neural networks for graph pooling. In *Procs. of ICML*, 2020.
- [3] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. Graph neural networks for social recommendation. In *Procs. of WWW*, 2019.
- [4] Yujie Fan, Mingxuan Ju, Chuxu Zhang, and Yanfang Ye. Heterogeneous temporal graph neural network. In *Procs. of SDM*, 2022.
- [5] Wenzheng Feng, Jie Zhang, Yuxiao Dong, Yu Han, Huanbo Luan, Qian Xu, Qiang Yang, Evgeny Kharlamov, and Jie Tang. Graph random neural networks for semi-supervised learning on graphs. *Procs. of NeurIPS*, 2020.
- [6] Nikhil Goyal, Harsh Vardhan Jain, and Sayan Ranu. Graphgen: A scalable approach to domain-agnostic labeled graph generation. In *Procs. of WWW*, 2020.
- [7] Zhichun Guo, Chuxu Zhang, Wenhao Yu, John Herr, Olaf Wiest, Meng Jiang, and Nitesh V Chawla. Few-shot graph learning for molecular property prediction. In *Procs. of WWW*, 2021.
- [8] Zhichun Guo, William Shiao, Shichang Zhang, Yozen Liu, Nitesh V Chawla, Neil Shah, and Tong Zhao. Linkless link prediction via relational distillation. In *Procs. of ICML*, 2023.
- [9] William L Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Procs. of NeurIPS*, 2017.
- [10] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Procs. of NeurIPS*, 2020.
- [11] Weihua Hu, Bowen Liu, Joseph Gomes, Marinka Zitnik, Percy Liang, Vijay Pande, and Jure Leskovec. Strategies for pre-training graph neural networks. In *Procs. of ICLR*, 2019.
- [12] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. In *Procs. of NeurIPS*, 2020.
- [13] Weihua Hu, Kaidi Cao, Kexin Huang, Edward W Huang, Karthik Subbian, and Jure Leskovec. Tuneup: A training strategy for improving generalization of graph neural networks. *arXiv*, 2022.
- [14] Ziniu Hu, Yuxiao Dong, Kuansan Wang, Kai-Wei Chang, and Yizhou Sun. Gpt-gnn: Generative pre-training of graph neural networks. In *Procs. of KDD*, 2020.
- [15] Wei Jin, Tong Zhao, Jiayuan Ding, Yozen Liu, Jiliang Tang, and Neil Shah. Empowering graph representation learning with test-time graph transformation. In *Procs. of ICLR*, 2023.
- [16] Jaehyeong Jo, Seul Lee, and Sung Ju Hwang. Score-based generative modeling of graphs via the system of stochastic differential equations. In *Procs. of ICML*, 2022.
- [17] Mingxuan Ju, Shifu Hou, Yujie Fan, Jianan Zhao, Yanfang Ye, and Liang Zhao. Adaptive kernel graph neural network. In *Procs. of AAAI*, 2022.
- [18] Mingxuan Ju, Wenhao Yu, Tong Zhao, Chuxu Zhang, and Yanfang Ye. Grape: Knowledge graph enhanced passage reader for open-domain question answering. In *Findings of EMNLP*, 2022.
- [19] Mingxuan Ju, Yujie Fan, Chuxu Zhang, and Yanfang Ye. Let graph be the go board: gradient-free node injection attack for graph neural networks via reinforcement learning. In *Procs. of AAAI*, 2023.

- [20] Mingxuan Ju, Tong Zhao, Qianlong Wen, Wenhao Yu, Neil Shah, Yanfang Ye, and Chuxu Zhang. Multi-task self-supervised graph neural networks enable stronger task generalization. In *Procs. of ICLR*, 2023.
- [21] Ildoo Kim, Younghoon Kim, and Sungwoong Kim. Learning loss for test-time augmentation. *Procs. of NeurIPS*, 2020.
- [22] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *Procs. of ICLR*, 2016.
- [23] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph neural networks meet personalized pagerank. In *Procs. of ICLR*, 2019.
- [24] Gang Liu, Tong Zhao, Eric Inae, Tengfei Luo, and Meng Jiang. Semi-supervised graph imbalanced regression. In *Procs. of KDD*, 2023.
- [25] Zemin Liu, Trung-Kien Nguyen, and Yuan Fang. Tail-gnn: Tail-node graph neural networks. In *Procs. of SIGKDD*, 2021.
- [26] Philip M Long and Hanie Sedghi. Generalization bounds for deep convolutional neural networks. In *Procs. of ICLR*, 2020.
- [27] Yao Ma, Xiaorui Liu, Neil Shah, and Jiliang Tang. Is homophily a necessity for graph neural networks? In *Procs. of ICLR*, 2022.
- [28] Julian McAuley, Rahul Pandey, and Jure Leskovec. Inferring networks of substitutable and complementary products. In *Procs. of SIGKDD*, 2015.
- [29] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018.
- [30] Chenhao Niu, Yang Song, Jiaming Song, Shengjia Zhao, Aditya Grover, and Stefano Ermon. Permutation invariant graph generation via score-based generative modeling. In *Procs. of AISTATS*, 2020.
- [31] Aditya Pal, Chantat Eksombatchai, Yitong Zhou, Bo Zhao, Charles Rosenberg, and Jure Leskovec. Pinnerage: Multi-modal user embedding framework for recommendations at pinterest. In *Procs. of SIGKDD*, 2020.
- [32] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Procs. of CVPR*, 2022.
- [33] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 2008.
- [34] Divya Shanmugam, Davis Blalock, Guha Balakrishnan, and John Guttag. Better aggregation in test-time augmentation. In *Procs. of CVPR*, 2021.
- [35] William Shiao, Uday Singh Saini, Yozen Liu, Tong Zhao, Neil Shah, and Evangelos E Papalexakis. Carl-g: Clustering-accelerated representation learning on graphs. *Procs. of KDD*, 2023.
- [36] Xianfeng Tang, Huaxiu Yao, Yiwei Sun, Yiqi Wang, Jiliang Tang, Charu Aggarwal, Prasenjit Mitra, and Suhang Wang. Investigating and mitigating degree-related biases in graph convolutional networks. In *Procs. of CIKM*, 2020.
- [37] Anton Tsitsulin, John Palowitch, Bryan Perozzi, and Emmanuel Müller. Graph clustering with graph neural networks. *arXiv*, 2020.
- [38] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. In *Procs. of ICLR*, 2017.
- [39] Petar Velickovic, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. In *Procs. of ICLR*, 2019.

- [40] Clement Vignac, Igor Krawczuk, Antoine Siraudin, Bohan Wang, Volkan Cevher, and Pascal Frossard. Digress: Discrete denoising diffusion for graph generation. In *Procs. of ICLR*, 2022.
- [41] Guotai Wang, Wenqi Li, Michael Aertsen, Jan Deprest, Sébastien Ourselin, and Tom Vercauteren. Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing*, 2019.
- [42] Qitian Wu, Hengrui Zhang, Junchi Yan, and David Wipf. Handling distribution shifts on graphs: An invariance perspective. In *Procs. of ICLR*, 2022.
- [43] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *Procs. of ICLR*, 2018.
- [44] Zhilin Yang, William Cohen, and Ruslan Salakhudinov. Revisiting semi-supervised learning with graph embeddings. In *Procs. of ICML*, 2016.
- [45] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. Graph convolutional neural networks for web-scale recommender systems. In *Procs. of SIGKDD*, 2018.
- [46] Jiaxuan You, Rex Ying, Xiang Ren, William Hamilton, and Jure Leskovec. Graphrnn: Generating realistic graphs with deep auto-regressive models. In *Procs. of ICML*, 2018.
- [47] Yuning You, Tianlong Chen, Yongduo Sui, Ting Chen, Zhangyang Wang, and Yang Shen. Graph contrastive learning with augmentations. In *Procs. of NeurIPS*, 2020.
- [48] Yuning You, Tianlong Chen, Yang Shen, and Zhangyang Wang. Graph contrastive learning automated. In *Procs. of ICML*, 2021.
- [49] Sukwon Yun, Kibum Kim, Kanghoon Yoon, and Chanyoung Park. Lte4g: Long-tail experts for graph neural networks. In *Procs. of CIKM*, 2022.
- [50] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. In *Procs. of NeurIPS*, 2018.
- [51] Zaixi Zhang, Qi Liu, Hao Wang, Chengqiang Lu, and Chee-Kong Lee. Motif-based graph self-supervised learning for molecular property prediction. *Procs. of NeurIPS*, 2021.
- [52] Tong Zhao, Tianwen Jiang, Neil Shah, and Meng Jiang. A synergistic approach for graph anomaly detection with pattern mining and feature learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.
- [53] Tong Zhao, Gang Liu, Daheng Wang, Wenhao Yu, and Meng Jiang. Learning from counterfactual links for link prediction. In *Procs. of ICML*, 2022.
- [54] Tong Zhao, Wei Jin, Yozen Liu, Yingheng Wang, Gang Liu, Stephan Günnemann, Neil Shah, and Meng Jiang. Graph data augmentation for graph machine learning: A survey. *IEEE DEBULL*, 2023.
- [55] Wenqing Zheng, Edward W Huang, Nikhil Rao, Sumeet Katariya, Zhangyang Wang, and Karthik Subbian. Cold brew: Distilling graph node representations with incomplete or missing neighborhoods. In *Procs. of ICLR*, 2021.
- [56] Qi Zhu, Carl Yang, Yidan Xu, Haonan Wang, Chao Zhang, and Jiawei Han. Transfer learning of graph neural networks with ego-graph information maximization. In *Procs. of NeurIPS*, 2021.
- [57] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. Deep Graph Contrastive Representation Learning. In *ICML Workshop on Graph Representation Learning and Beyond*, 2020.
- [58] Yanqiao Zhu, Yuanqi Du, Yinkai Wang, Yichen Xu, Jieyu Zhang, Qiang Liu, and Shu Wu. A survey on deep graph generation: Methods and applications. In *Procs. of LOG*, 2022.