

BANDITS FOR BLACK-BOX ATTACKS TO GRAPH NEURAL NETWORKS WITH STRUCTURE PERTURBATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Graph neural networks (GNNs) have achieved state-of-the-art performance in many graph-based tasks such as node classification and graph classification. However, many recent works have demonstrated that an attacker can mislead GNN models by slightly perturbing the graph structure. Existing attacks to GNNs are either under the less practical threat model where the attacker is assumed to access the GNN model parameters, or under the practical black-box threat model but consider perturbing node features that are shown to be not enough effective. In this paper, we aim to bridge this gap and consider black-box attacks to GNNs with structure perturbation. We propose to address this challenge through bandit techniques. Specifically, we formulate our attack as an online optimization with bandit feedback. This original problem is essentially NP-hard due to the fact that perturbing the graph structure is a binary optimization problem. We then propose a bandit convex optimization based attack which is proven to be sublinear to the query number T , i.e., $\mathcal{O}(\sqrt{NT}^{3/4})$ where N is the number of nodes in the graph. Finally, we evaluate our proposed attack algorithm by conducting experiments over multiple datasets and GNN models. The experimental results on various citation graphs and image graphs show that our attack is both effective and efficient.

1 INTRODUCTION

Graph neural networks (GNNs) have been emerging as the most prominent methodology for learning with graphs, such as social networks, chemical networks, superpixel graphs, etc. GNNs have also advanced many graph-related applications including but not limited to drug discovery (Shi et al., 2020), fake news detection on social media (Monti et al., 2019), traffic forecasting (Yu et al., 2018), and superpixel graph classification (Dwivedi et al., 2020). However, recent works have shown that GNNs are vulnerable to adversarial attacks (Dai et al., 2018; Zügner et al., 2018; Wu et al., 2019b; Zügner & Günnemann, 2019; Wang & Gong, 2019; Xu et al., 2019a; Sun et al., 2020; Ma et al., 2020a). That is, an attacker can easily fool a GNN model by slightly perturbing the graph structure (e.g., adding the fake edges into the graph or deleting the existing edges from the graph) or perturbing the node features. Most of the existing attacks to GNNs essentially rely on white-box or gray-box threat model (Dai et al., 2018; Zügner et al., 2018; Wu et al., 2019b; Zügner & Günnemann, 2019; Wang & Gong, 2019; Xu et al., 2019a; Sun et al., 2020). In this threat model, an attacker can not only obtain the predictions generated by the targeted GNN model, but also know the whole or partial GNNs’ inner parameters and network structure. Such a threat model enables the attacker to derive the true gradients that can be used to construct an (almost) optimal edge/feature perturbation via first-order optimizations approaches, e.g., projected gradient descent, or PGD for short.

In practice, however, an attacker often has limited knowledge about the GNN model. For instance, many models are deployed as an API due to the commercial value. In these practical scenarios, an attacker can only obtain the model predictions by querying the API, while without knowing the model’s other information. An attack based on such a realistic threat model is called a *black-box attack*, which significantly raises the bar for the attacker as he cannot obtain the gradient information. A recent work (Ma et al., 2020a) performs black-box attacks against GNNs. However, this work has two key drawbacks. First, it assumes that the attacker can only perturb the (continuous) node features. Existing works (e.g., (Zügner et al., 2018)) have shown that feature perturbation-based attacks to GNNs are significantly less effective than structure perturbation-based attacks. Second, the attack is implemented via a heuristic greedy algorithm, which has no theoretically guaranteed performance.

In this paper, we consider black-box attacks against GNNs with *structure perturbation*. Such a new attack setting is much more challenging, as finding the optimal structure perturbation is essentially an NP-hard problem (i.e., a binary optimization problem) and the attacker only obtains the predictions via querying the model. We take the first step to solve the structure perturbation-based black-box attacks. Specifically, we first reformulate our attack as a bandit optimization (i.e., online optimization with bandit feedback) problem, which characterizes the attacker’s query process on the black-box GNN model and captures the uncertainty during the query. Then, we handle the binary constraint brought by the discrete structure perturbation and integrate it into our bandit-based attack objective. Finally, we perform the theoretical analysis for our attack algorithm. Our contributions are as follows:

- We design the first structure perturbation-based black-box attack to GNNs via bandit optimization.
- We prove that our bandit-based attack algorithm theoretically yields a *sublinear* regret bound $\mathcal{O}(\sqrt{NT}^{3/4})$ within T queries for attacking a graph with N nodes.
- We conduct extensive experiments to evaluate our attack over multiple datasets and GNN models and demonstrate the effectiveness and efficiency of our attack algorithm.

2 PRELIMINARIES AND PROBLEM FORMULATION

2.1 GRAPH NEURAL NETWORKS

Let $G = (\mathcal{V}, \mathcal{E}, \mathbf{X})$ be a graph, where $u \in \mathcal{V}$ is a node, $(u, v) \in \mathcal{E}$ is an edge between u and v , and $\mathbf{X} = [\mathbf{x}_1; \mathbf{x}_2; \dots; \mathbf{x}_N] \in \mathbb{R}^{N \times d}$ is the node feature matrix. We denote $\mathbf{a}_v = [a_{v1}; a_{v2}; \dots; a_{vN}] \in \{0, 1\}^N$ as the adjacency vector of node v . $N = |\mathcal{V}|$ and $M = |\mathcal{E}|$ are the number of nodes and edges, respectively. We denote d_u and $\mathcal{N}(u)$ as u ’s node degree and the neighborhood set of u (including self-loop (u, u)). We consider GNNs for node classification in this paper¹. In this context, each node $u \in \mathcal{V}$ has a label y_u from a label set $\mathcal{Y} = \{1, 2, \dots, L_C\}$. Given a set of $\mathcal{V}_L \subset \mathcal{V}$ labeled nodes $\{(\mathbf{x}_u, y_u)\}_{u \in \mathcal{V}_L}$ as the training set, GNN for node classification is to take the graph G and labeled nodes as input and learn a node classifier that maps each unlabeled node $u \in \mathcal{V} \setminus \mathcal{V}_L$ to a class $y \in \mathcal{Y}$.

Generally speaking, GNN consists of two main steps: *neighborhood aggregation* and *node representation update*. Suppose a GNN has K layers. We denote v ’s representation in the k -th layer be $\mathbf{h}_v^{(k)}$, with $\mathbf{h}_v^{(0)} = \mathbf{x}_v$. In neighborhood aggregation, GNN obtains the representation $\mathbf{l}_v^{(k)}$ by aggregating representations of v ’s neighbors in the $(k-1)$ -th layer as follows:

$$\mathbf{l}_v^{(k)} = \text{AGG}(\{\mathbf{h}_u^{(k-1)} : u \in \mathcal{N}(v)\}). \quad (1)$$

In node representation update, GNN updates v ’s representation at the k -th layer via parameterizing the aggregated representation $\mathbf{l}_v^{(k)}$ with parameters matrix $\mathbf{W}^{(k)}$ as follows:

$$\mathbf{h}_v^{(k)} = \text{UPDATE}(\mathbf{l}_v^{(k)}, \mathbf{W}^{(k)}). \quad (2)$$

Different GNNs use different AGG and UPDATE functions. For instance, in Graph Convolutional Network (GCN) (Kipf & Welling, 2017), AGG is the element-wise mean pooling function and UPDATE is the ReLU activation function. More specifically, it has the following form:

$$\mathbf{h}_v^{(k)} = \text{ReLU}\left(\mathbf{W}^{(k)} \left(\sum_{u \in \mathcal{N}(v)} d_u^{-1/2} d_v^{-1/2} \mathbf{h}_u^{(k-1)} \right)\right).$$

A node v ’s final representation $\mathbf{h}_v^{(K)} \in \mathbb{R}^{|\mathcal{Y}|}$ can capture the structural information of all nodes within v ’s K -hop neighbors. Moreover, the final node representations of training nodes are used to train the node classifier. Specifically, let $\Theta = \{\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(K)}\}$ be the model parameters and v ’s output be $f_{\Theta}(\mathbf{a}_v) = \text{softmax}(\mathbf{h}_v^{(K)}) \in \mathbb{R}^{|\mathcal{Y}|}$, where $f_{\Theta}(\mathbf{a}_v)_y$ indicates the probability of node v being class y ². Then, Θ are learnt by minimizing the cross-entropy loss on the outputs of the training nodes \mathcal{V}_L , i.e.,

$$\Theta^* = \arg \min_{\Theta} - \sum_{v \in \mathcal{V}_L} \ln f_{\Theta}(\mathbf{a}_v)_y. \quad (3)$$

With the learnt Θ^* , we can predict the label for each unlabeled nodes $u \in \mathcal{V} \setminus \mathcal{V}_L$ as $\hat{y}_u = \arg \max_y f_{\Theta^*}(\mathbf{a}_u)_y$.

¹Our attack can be naturally generalized to GNNs for graph classification. We discuss this in Section 3.1.

²Note that the prediction also depends on v ’s node feature \mathbf{x}_v and the whole graph G . We omit \mathbf{x}_v and G for notation simplicity.

2.2 THREAT MODEL

Attacker’s knowledge. The considered black-box attack setting in this paper implies that an attacker does not know the internal configurations (i.e., the learned parameters) of the targeted GNN model. For a target node $v \in \mathcal{V}$, the only information the attacker knows is the predictions $f_{\Theta^*}(\mathbf{a}_v)$ via querying the GNN model f_{Θ^*} . Moreover, we also reasonably assume that the attacker knows her neighbors, i.e., the adjacency vector \mathbf{a}_v ³.

Attacker’s capability. We consider that the attacker can modify the connection status (e.g., inject new edges or remove existing edges) between the target node v and other nodes in the graph. In practice, it also incurs different costs for the attacker to manipulate different edges. The attacker’s budget of manipulating edges is often limited, and we denote by C the cost budget. We also constrain that the number of edges to be manipulated is bounded by B .

Attacker’s goal. Based on the attacker’s knowledge and capability, an attacker’s goal is to fool a targeted GNN, i.e., making her target node’s predicted label \tilde{y}_v different from the true label y_v , by manipulating her adjacency vector \mathbf{a}_v with cost budget C and the maximal perturbed number B .

2.3 PROBLEM FORMULATION

Given the target node v , node feature \mathbf{x}_v , label y_v , and adjacency vector \mathbf{a}_v , an attacker aims to modify the connection status related to the target node v such that the targeted GNN misclassifies v . Let $\mathbf{s}_v \in \{0, 1\}^N$ be the adversarial structure perturbation on v , where $s_{vu} = 1$ means the connection status between the nodes v and u is changed. Then, we define the perturbed adjacency vector for v as $\mathbf{a}_v \oplus \mathbf{s}_v$, where \oplus is the XOR operator between two binary vectors. Moreover, we denote $\mathbf{c}_v \in \mathbb{R}^N$ as the cost vector associated with v , i.e., c_{vu} is the cost to modify the connection status between v and u . In the focused black-box setting, we consider the untargeted attack. Let $L(\mathbf{a}_v)$ be the loss function for the targeted node v without attack. With the adversarial perturbation \mathbf{s}_v , we have the attack loss as $L(\mathbf{a}_v \oplus \mathbf{s}_v)$, which is defined as,

$$L(\mathbf{a}_v \oplus \mathbf{s}_v) = L(\mathbf{a}_v) - \max_{y \neq y_v} f_{\Theta^*}(\mathbf{a}_v \oplus \mathbf{s}_v)_y. \quad (4)$$

The first term is the original loss without attack. The second term is the maximum logit score of the label other than the original label y_v . In this paper, we use the popular CW attack loss function (Carlini & Wagner, 2017) with κ attack confidence. Specifically, it is defined as follows:

$$L(x) = \max\{f_{\Theta^*}(x)_y - \max_{\hat{y} \neq y} \{f_{\Theta^*}(x)_{\hat{y}}\}, -\kappa\}. \quad (5)$$

Finally, our problem of structure perturbation-based black-box attack to GNN can be formulated as

$$\min_{\mathbf{s}_v} L(\mathbf{a}_v \oplus \mathbf{s}_v), \quad \text{s.t.}, \quad \mathbf{1}^T \mathbf{s}_v \leq B, \quad \mathbf{c}_v^T \mathbf{s}_v \leq C, \quad \mathbf{s}_v \in \{0, 1\}^N, \quad (6)$$

where the first constraint means the number of edges to be perturbed is no more than B and the second constraint means the total costs of the perturbation are no more than C .

Lemma 1. *Our problem in Equation 6 is NP-hard.*

Proof. Our problem in Equation 6 is a combinatorial optimization problem, actually a type of knapsack problem, which is a classical NP-hard problem and so does our problem. \square

Lemma 1 implies that it is difficult to calculate the optimal perturbation vector \mathbf{s}_v^* within polynomial time under large graphs (i.e., \mathbf{s}_v has large dimension). To this end, we aim to design an approximation algorithm to derive sub-optimal solution. One algorithm is to relax the combinatorial binary constraint $\mathbf{s}_v \in \{0, 1\}^N$ into convex hull $\mathbf{s}_v \in [0, 1]^N$ and obtain a continuous optimization problem. Let $\hat{\mathbf{s}}_v$ be the solution of the continuous optimization problem. We can derive the sub-optimal solution for the original problem in Equation 6 by rounding $\hat{\mathbf{s}}_v$ into combinatorial space $\{0, 1\}^N$ using randomization sampling like Bernoulli sampling (Liu et al., 2015; Xu et al., 2019a). Then, we have the following lemma to characterize the relation between \mathbf{s}_v and $\hat{\mathbf{s}}_v$ in expectation:

Lemma 2. *When sampling \mathbf{s}_v element-wise in Bernoulli distribution using the probability from the relaxed vector $\hat{\mathbf{s}}_v \in [0, 1]^N$, then the expectation of \mathbf{s}_v is $\hat{\mathbf{s}}_v$, i.e., the following condition holds,*

$$\mathbb{E}[\mathbf{s}_v] = \hat{\mathbf{s}}_v. \quad (7)$$

³For graph classification, we assume attackers know the input graph.

Proof. This lemma holds due to the fact that a random variable \mathbf{X} subject to Bernoulli distribution on support $\{0, 1\}$ takes its probability as expectation, i.e., $\mathbb{E}[\mathbf{X}] = \mathbb{P}[\mathbf{X}]$. Applying this fact elements in \mathbf{s}_v , we can obtain Eq. (7). \square

Conventionally, to solve our relaxed continuous optimization problem, we can apply the projected gradient descent (PGD) approach by running gradient updates projected onto the feasible domain within several steps. However, PGD requires that the gradient information needs to be available. In our black-box attack setting, only the prediction information via querying the GNN is available instead of the exact gradient. Therefore, the attack problem becomes how to compute the approximate gradient such that the PGD method can still be applied. It is shown that zeroth-order optimization (short for ZOO) can be used to estimate the gradient (Chen et al., 2017; Ilyas et al., 2019; Liu et al., 2018). However, ZOO suffers from a low convergence rate and high query overhead due to only exploring the unknown gradient from the uniformly sampled edges. We aim to estimate the unknown gradient by controlling the exploration-exploitation tradeoff via bandit methods. Reinforcement learning (RL) can also control the exploration-exploitation tradeoff. However, RL-based attack, i.e., RL-S2V (Dai et al., 2018), is naturally heuristic. Our attack can address both issues in ZOO-based and RL-based attacks. Specifically, our attack has theoretical guarantees and better attack performance (See Section 6). Next, we reformulate our attack problem as a bandit optimization problem and then propose a solution to solve the bandit problem.

2.4 REFORMULATING OUR ATTACK AS A BANDIT PROBLEM

When the attacker selects a perturbation vector \mathbf{s}_v and uses the perturbed adjacent vector $\mathbf{a}_v \oplus \mathbf{s}_v$ to query the GNN, the GNN returns the prediction $f_{\Theta^*}(\mathbf{a}_v \oplus \mathbf{s}_v)$. Thus, the objective $L(\mathbf{a}_v \oplus \mathbf{s}_v)$ is revealed, which can be seen as a bandit feedback (i.e., reward) for the selected perturbation vector \mathbf{s}_v (i.e., an arm). Under the bandit feedback, the attacker wants to maximize the cumulative rewards. Note that since the attacker does not know the optimal arm \mathbf{s}_v^* in each round, it will incur a regret, i.e., the difference of the maximum reward under the optimal arm \mathbf{s}_v^* in hindsight and the reward of the attacker’s attack algorithm. Then, the attacker’s goal is to minimize the cumulative regrets. Let $\text{Reg}(T)$ be the cumulative regrets in T rounds, and \mathbf{s}_v^t be the perturbation vector selected at round t , then the cumulative regrets $\text{Reg}(T)$ can be calculated as,

$$\text{Reg}(T) = \mathbb{E}\left[\sum_{t=1}^T L(\mathbf{s}_v^t)\right] - TL((\mathbf{s}_v^t)^*). \quad (8)$$

In bandit optimization, it is important to design an arm selection algorithm with *sublinear* regret (i.e., $\text{Reg}(T) = o(T)$). This is because the selected arm \mathbf{s}_v at round T is asymptotically optimal when T is sufficiently large (i.e., $\lim_{T \rightarrow \infty} \frac{\text{Reg}(T)}{T} = 0$).

3 BANDIT-BASED BLACK-BOX ATTACKS TO GNN

In this section, we design an online attack to GNN based on bandits. We also prove that our attack algorithm has a *sublinear regret*.

First, we relax the binary perturbation vector $\mathbf{s}_v \in \{0, 1\}^N$ into a continuous convex hull $\hat{\mathbf{s}}_v \in [0, 1]^N$. In this case, we can define the arm set \mathcal{W} as follows:

$$\mathcal{W} = \{\hat{\mathbf{s}}_v \in [0, 1]^N \mid \mathbf{1}^T \hat{\mathbf{s}}_v \leq B, \mathbf{c}^T \hat{\mathbf{s}}_v \leq C\}, \quad (9)$$

where \mathcal{W} is convex. Note that our bandit instance for the black-box setting is different from the traditional multi-armed bandits (MAB), as our arm set \mathcal{W} contains infinite perturbation vectors. Therefore, the approaches like UCB (Auer & Ortner, 2010) and Thompson Sampling (Russo et al., 2017) in MAB fail to solve our problem.

Next, we need to address how to efficiently decide the next arm (i.e., $\hat{\mathbf{s}}_v^{t+1}$) at the end of round t such that the regret is minimized. We leverage online convex optimization (OCO) technique to derive the next arm $\hat{\mathbf{s}}_v^{t+1}$ at round $t+1$. We note that the loss function $L(\cdot)$ is often non-convex. We emphasize that gradient descent used in OCO is still useful as it is challenging to derive the closed-form solution for non-convex functions. OCO approach requires gradient information at the selected arm to conduct gradient descent. In our black-box attack setting, the attacker only receives the bandit feedback for the selected arm $\hat{\mathbf{s}}_v^t$. To estimate the gradient for the black-box attack at the selected arm $\hat{\mathbf{s}}_v^t$, the

Algorithm 1 Black-box attack to GNN for node classification via OCO with bandit**Input:** GNN f_{Θ}^* , target node v , budget B , cost C , PGD step η , $\delta > 0$, $\alpha \in [0, 1]$, query number T **Output:** Perturbed vector s_v

- 1: Initialize: $v^1 = \mathbf{0} \in \mathcal{W} = \{\hat{s}_v \in [0, 1]^N \mid \mathbf{1}^T \hat{s}_v \leq B, \mathbf{c}^T \hat{s}_v \leq C\}$
- 2: **for** $t = 1$ **to** T **do**
- 3: Attacker randomly chooses a unit vector \mathbf{u}^t from the unit sphere \mathcal{S} .
- 4: Attacker determines a perturbation $\hat{s}_v^t = v^t + \delta \mathbf{u}^t$.
- 5: Attacker converts \hat{s}_v^t to be binary s_v^t by setting top- B values in \hat{s}_v^t to be 1 and others to be 0.
- 6: Attacker queries the GNN model f_{Θ}^* with s_v^t to obtain the predictions and CW loss $L(s_v^t)$.
- 7: Attacker conducts PGD and updates: $v^{t+1} = \prod_{(1-\alpha)\mathcal{W}}(v^t - \eta \hat{\mathbf{g}})$, $\hat{\mathbf{g}} = \frac{N-1}{\delta} L(s_v^t) \mathbf{u}^t$.
- 8: **end for**
- 9: **return** s_v^T

attacker can only query the GNN to compute the gradient. We use *one point gradient estimation* (OPGE) (Granichin, 1989; Spall, 1997) technique to estimate the gradient. The idea of OPGE is to find a vector in unit sphere $\mathcal{S} = \{\mathbf{u} \in \mathbb{R}^{N-1} \mid \|\mathbf{u}\|_2 = 1\}$ such that its direction has small intersection angle with the gradient (i.e., \mathbf{u} is a good approximation of the gradient). To this end, we uniformly sample a unit vector in \mathcal{S} and derive an approximate gradient in the following lemma.

Lemma 3. *For a unit vector \mathbf{u} uniformly sampled in \mathcal{S} and a sufficient small $\delta > 0$, we have the following estimated gradient:*

$$\nabla \hat{L}(\hat{s}_v) = \mathbb{E}_{\mathbf{u} \in \mathcal{S}} \left[\frac{N-1}{\delta} L(\hat{s}_v + \delta \mathbf{u}) \mathbf{u} \right]. \quad (10)$$

After we obtain a continuous solution \hat{s}_v^t after T rounds, we round \hat{s}_v^t to a binary solution s_v by randomizing entry-wise components based on the probabilities in \hat{s}_v^t . The details of our attack algorithm are shown in Algorithm 1.

The inputs of our algorithm include the target GNN model f_{Θ}^* , target node v , budget B , cost C , PGD step η , a small $\delta > 0$, projection scale $\alpha \in [0, 1]$, and query number T . The output is a perturbed vector s_v for the target node v after T rounds. In line 1, we set $\mathbf{0}$ as the initial prior vector v^1 . In line 2–8, we calculate a sub-optimal perturbed vector to attack the targeted GNN based on OCO with bandit feedback. At round t , we randomly select a unit vector \mathbf{u}^t from the unit sphere \mathcal{S} as a stochastic gradient in line 3. In line 4, we derive a relaxed perturbed vector \hat{s}_v^t by updating the prior vector v^t according to the selected stochastic gradient \mathbf{u}^t . In line 5, we convert \hat{s}_v^t to binary s_v^t by setting its top- B values to be 1 and the remaining values to be 0. In line 6, we query the GNN model f_{Θ}^* with s_v^t and obtain a loss feedback L . In line 7, we conduct PGD to update the v^{t+1} for the next round. Finally, after T queries, we obtain the perturbed vector s_v^T for the target node v .

Extending our attack to GNNs for graph classification. Our proposed attack against node classification can be naturally extended to attack GNN models for graph classification without efforts of modifications. In a graph classification model, its input is an adjacent matrix of a graph and its output is the label of the graph. In node classification, we aim to perturb the adjacent vector of a target node, while in graph classification, we perturb the adjacent matrix. Let $\mathbf{A} \in \{0, 1\}^{N \times N}$ be the adjacent matrix of a graph with N nodes. Moreover, let $\mathbf{S} \in \{0, 1\}^{N \times N}$ be a perturbation matrix, where $S_{ij} = 1$ if the connection status between the edge (i, j) is modified, and $S_{ij} = 0$, otherwise. To perform our attack against graph classification models, we only need to flatten the matrix \mathbf{S} into a vector s and feed it as an input to our attack algorithm. After obtaining the perturbed s , we can reshape it to a perturbed adjacent matrix.

4 MAIN RESULTS

In this section, we analyze the regret bound of our attack algorithm where we assume the loss function is convex. We note that it is an interesting future work to generalize our analysis to non-convex loss function. Before providing the regret bound, we first present the following assumptions and lemmas. Due to the limited space, we defer the detailed proofs to the Appendix.

Assumption 1. *There exists a Lipschitz constant C_L such that the following inequality holds for any \mathbf{u} and \mathbf{v} ,*

$$|L(\mathbf{u}) - L(\mathbf{v})| \leq C_L \|\mathbf{u} - \mathbf{v}\|_2. \quad (11)$$

Lemma 4. For continuous $\hat{\mathbf{s}}_v^t$ and the rounded binary \mathbf{s}_v , the instant regret is bounded as:

$$\mathbb{E}[|L(\mathbf{s}_v) - L(\hat{\mathbf{s}}_v^t)|] \leq C_L(N-1)^{3/4} \sqrt{1 + \frac{\eta}{\delta}}. \quad (12)$$

Line 6 in Algorithm 1 can be seen as the stochastic gradient descent on $L(\hat{\mathbf{s}}_v + \delta \mathbf{u})$. Thus, we have the following lemma to characterize its regret bound.

Lemma 5 ((Flaxman et al., 2004)). Suppose that the arm set \mathcal{W} satisfies $\mathcal{W} \subseteq R\mathbb{B}$ for given radius $R > 0$, where \mathbb{B} is unit ball, i.e., $\mathbb{B} = \{\mathbf{u} \in \mathbb{R}^{N-1} : \|\mathbf{u}\|_2 \leq 1\}$. When loss function $L(\cdot)$ is convex, the cumulative regret for relaxed continuous variables are bound as,

$$\mathbb{E} \sum_{t=1}^T L(\hat{\mathbf{s}}_v^t) - TL(\hat{\mathbf{s}}_v^*) \leq C_L R \sqrt{T}, \quad (13)$$

where $\hat{\mathbf{s}}_v^t$ is continuous variable at round t and $\hat{\mathbf{s}}_v^*$ is optimal continuous solution of our relaxed optimization problem.

Note that Lemma 5 just captures the regret on the whole arm set \mathcal{W} . However, line 6 in Algorithm 1 updates the arm $\hat{\mathbf{s}}_v^t$ by projecting onto set $(1-\alpha)\mathcal{W}$ for $0 < \alpha < 1$. The incurred regret by conducting $(1-\alpha)$ -projection is captured by the following lemma.

Lemma 6. For time horizon T , the incurred regret due to $(1-\alpha)$ -projection is bounded as follows,

$$\min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T L(\mathbf{w}) - TL(\hat{\mathbf{s}}_v^*) \leq 2\alpha T. \quad (14)$$

Based on the above assumption and the lemmas, we have the following theorem on the regret bound:

Theorem 1. Under T rounds attack span, our proposed attack method incurs regret $\text{Reg}(T)$, which is upper bounded by $\mathcal{O}(\sqrt{NT}^{3/4})$ with T queries to the GNN.

Remark. Theorem 1 not only demonstrates the sublinear regret our attack achieves, but also presents that our attack is dimension-efficient, i.e., $\mathcal{O}(\sqrt{N})$. It implies that our attack can be scalable to large graphs. Note that gradient-free ZOO (Chen et al., 2017) has query complexity $\mathcal{O}(N)$. From the regret bound, we can see that our attack enjoys a better convergence rate $\mathcal{O}(1/T^{3/4})$ than ZOO (Liu et al., 2018), which has a $\mathcal{O}(1/T^{1/2})$ convergence rate. Note also that Ilyas et al. (2019) proposed a black-box attack with bandit to classifiers on image data, which can be adapted to solve our problem. However, their approach 1) does not provide theoretical results in terms of regret bound; and 2) is less efficient than ours due to requiring multiple gradient estimation in each iteration.

5 RELATED WORK

Recent attacks to GNNs mainly focus on white-box/grey box settings (Zügner et al., 2018; Dai et al., 2018; Zügner & Günnemann, 2019; Xu et al., 2019a; Sun et al., 2020; Chang et al., 2020; Zhang et al., 2020; Wang et al., 2020; Ma et al., 2020b). For instance, Zügner et al. (2018) proposed the first attack, called Nettack, against GCN for node classification by perturbing graph structure or/and node features. Specifically, Nettack learns a surrogate linear model of GCN by defining a graph structure-preserving perturbation that constrains the difference between the node degree distributions of the graph before and after an attack. Xu et al. (2019a) utilized the model gradient to generate perturbation on the topology of the graph. Chang et al. (2020) designed a grey-box attack against graph embedding methods. We study the black-box setting where gradient information is completely unknown to the attacker. A recent work (Ma et al., 2020a) performed black-box attacks against GNNs. However, this work focus on perturbing the continuous node features, which is less effective than discrete structure perturbation. In addition, the attack is implemented via a heuristic greedy algorithm, which has no theoretically guaranteed performance. Zang et al. (2020) studied the graph universal adversarial attacks where a set of anchor nodes is identified and their connection to the target node is flipped. However, how to select the optimal anchor nodes is not investigated.

Several black-box attacks (Ilyas et al., 2018; 2019; Cheng et al., 2019) for non-graph classification models have been proposed. However, these methods cannot solve our problem because their attack

problems are essentially continuous optimization problems. Note that Ilyas et al. (2019) also used the bandit to formulate the black-box attack problem. However, their work does not have a theoretical regret bound and is less efficient due to using multiple gradient estimations in each iteration.

6 EXPERIMENTS

6.1 EXPERIMENTAL SETUP

Dataset description and GNN models. In node classification experiments, we use three benchmark citation graphs, i.e., Cora, Citeseer, and Pubmed (Sen et al., 2008). In these graphs, each node represents a document and each edge indicates a citation between two documents. Each document treats the bag-of-words feature as the node feature vector, and has a label. We adopt the representative GCN (Kipf & Welling, 2017) and SGC (Wu et al., 2019a) for node classification, and evaluate our attack against the two models. In graph classification, we use two benchmark superpixel graphs, i.e., MNIST and CIFAR10 (Dwivedi et al., 2020), in computer vision. MNIST and CIFAR10 are classical image classification datasets. In our experiments, they are converted into graphs using the SLIC super-pixels (Achanta et al., 2012). Each node has the super-pixel coordinates as the feature and each super-pixel graph has a label. We adopt the representative GIN (Xu et al., 2019b) for graph classification, and evaluate our attack against GIN. Table 1 in Appendix summarizes the basic statistics of these graph datasets.

Training nodes/graphs and target nodes/graphs. We use the training nodes/graphs to train GNN models and use the target nodes/graphs to evaluate our attack against the trained models. Following existing works (Kipf & Welling, 2017; Zügner et al., 2018; Dwivedi et al., 2020), in citation graphs, we randomly sample 20 nodes from each class as the training nodes; sample 100 nodes that are correctly classified by each GNN model as the target nodes. In image graphs, we respectively use 55,000 graphs and 45,000 graphs in MNIST and CIFAR10 for training, and randomly sample 100 graphs correctly classified by the GIN model as the target graphs.

Compared methods. We compare our attack with two state-of-the-arts: *RL-based attack* (Dai et al., 2018) and *Zoo attack* (Chen et al., 2017). We also choose *random attack* for comparison, where we generate structure perturbations by randomly changing connection status between pairs of nodes.

Cost simulation. We specify a cost of modifying the connection state for each pair of nodes. Note that the costs could be application-dependent. W.l.o.g., we assume the costs for different node pairs are uniformly distributed among a certain interval (e.g., [1,5] in our experiments). Note that other costs (Wang & Gong, 2019) like equal cost can be seen as a special case of the uniform cost.

Parameter setting. We set the hyperparameters in our attack algorithm as follows: $\eta = 10^{-4}$, $\delta = 10^{-6}$, and $\alpha = 0.7$ for attacking node classification methods, and $\eta = 10^{-1}$, $\delta = 10^{-3}$, and $\alpha = 0.6$ for attacking graph classification methods. Consider the different graph sizes, we set the default maximal number of perturbed edges B as 5 and 15 on the three citation graphs, and on the two image graphs, respectively; and the default total costs C is bounded by 25 on the citation graphs and 75 on the image graphs, respectively. In addition, the total number of queries $T = 50$ by default. We also study the impact of these parameters in our experiments. We implement our algorithm in Python and conduct experiments using public source codes. All our experiments are run in a computer with Intel(R) Core(TM) i7-6700 CPU @3.4Hz processors with 4 cores, 32GB RAM, 1 TB disk space and 6G GPU. We run all experiments 30 times and report the average result.

Evaluation metric. We adopt the attack successful rate, i.e., a fraction of the total (i.e., 100) targeted nodes/graphs misclassified after our attack, as the metric to measure the effectiveness of our attack. We also use the number of queries (T) to measure the efficiency of our attack. Specifically, we count each PGD iteration in our Algorithm 1 as a query.

6.2 EXPERIMENTAL RESULTS

In this section, we provide a comprehensive evaluation of our black-box attack against GNNs for both node classification and graph classification. We aim to study our attack in terms of both effectiveness and efficiency. Our attack algorithm has three key factors: the number of perturbed edges B , total costs C , and number of queries T . We will study the impact of these parameters one by one. More experimental results are shown in Appendix.

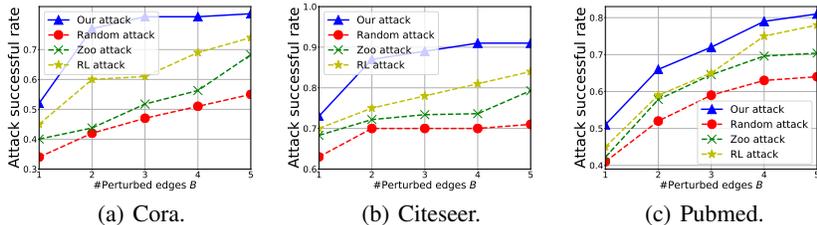


Figure 1: Attack successful rate vs. #perturbed edges B on GCN for node classification.

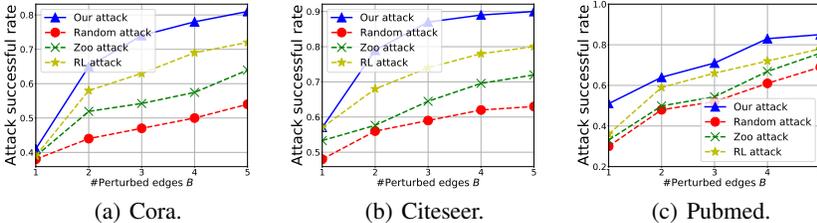


Figure 2: Attack successful rate vs. #perturbed edges B on SGC for node classification.

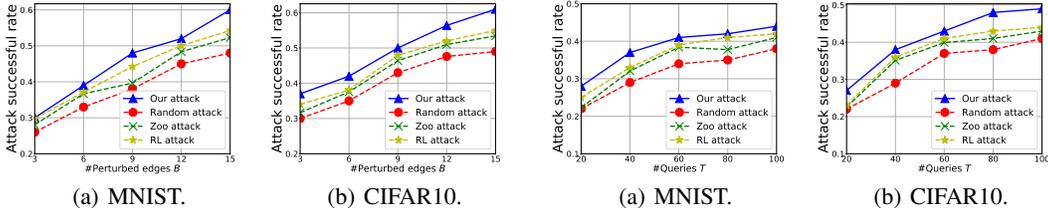


Figure 3: Attack successful rate vs. #perturbed edges B on GIN for graph classification.

Figure 4: Attack successful rate vs. #queries T on GIN for graph classification.

Impact of the number of perturbed edges. In this experiment, we separately fix the bounded total costs C to be 25 and 75, and set the number of queries T to be 50 on attacking both node classification models and graph classification models. The results of attacking these models are shown in Figure 1, Figure 2, and Figure 3, respectively.

We have the following observations: First, all attack approaches achieve a larger attack successful rate when the maximal number of perturbed edges B increases. This is because a larger B allows an adversary to have a better capability to perform the attack. Second, our attack significantly achieves a higher attack successful rate than the compared attacks in all GNNs models and graph datasets. For instance, when attacking GCN and SGC for node classification, our attack achieves the attack successful rate larger than 80% (and even 90%), while the second-best RL attack achieves the attack successful rate at most 80% across the three citation graphs. When attacking GIN for graph classification, our attack achieves 60% attack successful rates, while RL attack can achieve less than 53% attack successful rate in the two image graphs. All these results verify that the black-box bandit feedback in our algorithm is very useful to guide the selection of the optimal edge to be perturbed. In contrast, the heuristic RL attack is hard to select the optimal edge to be perturbed.

Impact of the number of queries. In this experiment, we separately fix the bounded total costs C to be 25 and 75, and the maximal number of perturbed edge B to be 2 and 5 on attacking node classification models and graph classification models, respectively. The results of attacking GIN for graph classification and attacking GCN and SGC for node classification are shown in Figure 4, Figure 5, and Figure 6, respectively.

We have the following observations: First, all attack approaches achieve a larger attack successful rate when the number of queries T increases. This is because a larger T allows an adversary to obtain more predictions via querying the GNN models and thus have a better capability to perform the attack. Second, our attack requires much fewer queries to achieve the same attack successful rate than the compared attacks. In particular, compared to our attack, RL attack requires 1.5-3x queries, when achieving comparable attack performance against node classification models and graph classification

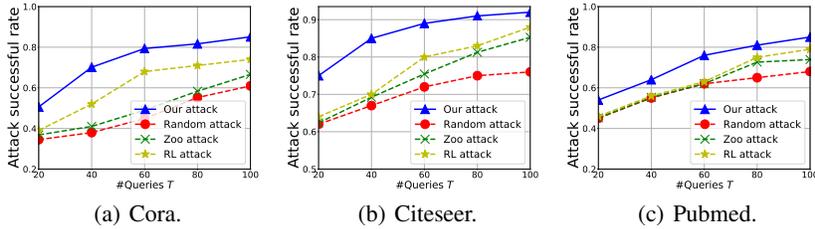


Figure 5: Attack successful rate vs. #queries T on GCN for node classification.

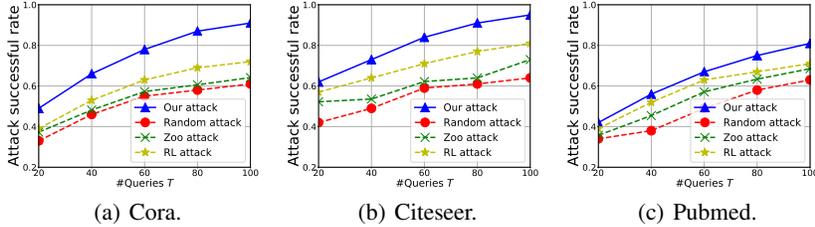


Figure 6: Attack successful rate vs. #queries T on SGC for node classification.

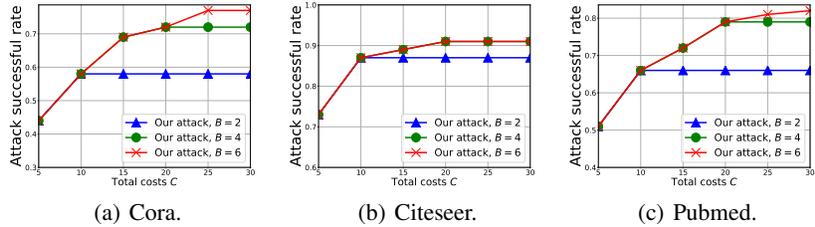


Figure 7: Attack successful rate vs. #total costs C on GCN for node classification.

models, respectively. Third, given the same number of queries, our attack achieves 10% – 20% higher successful rate than RL attack across all citation graphs and image graphs. Similarly, these results again verify that the black-box bandit feedback is really beneficial in designing our attack.

Impact of the total costs. We study the impact of the total costs C and fix the number of queries T to be 50. As our attack outperforms the compared attacks, we only study results on our attack on attacking GCN for simplicity. Note that we have similar observations on attacking SGC and GIN.

The result is shown in Figure 7. We observe that the attack successful rate becomes higher when the cost budget is larger. Moreover, note that the cost for perturbing each edge is within $[1, 5]$. If the cost budget is sufficient, i.e., $C \geq 10, 20, 30$ for $B = 2, 4, 6$, then the attack successful rate is stable, as our attack finds the same space for structure perturbation. In contrast, if the cost budget is insufficient, even the allowed number of perturbed edges is larger, our attack performance cannot be improved, e.g., $C = 20$ for $B = 4, 6$ on Citeseer. The result demonstrates that cost budget is a key factor to affect our attack performance.

7 CONCLUSION AND FUTURE WORK

In this paper, we study black-box attacks to GNNs via manipulating the graph structure. We first introduce the threat model of our black-box attack from three dimensions: attacker’s knowledge, attacker’s capability, and attacker’s goal. Then, we formulate our attack as a binary optimization problem, which is NP-hard. Next, we relax and reformulate our attack problem as a bandit optimization problem, and propose a bandit-based attack algorithm and rigorously prove that our attack yields a sublinear regret bound $\mathcal{O}(\sqrt{NT}^{3/4})$ within T queries for attacking a graph with N nodes. Finally, we evaluate our attack against GNN models for both node classification and graph classification. Our results demonstrate both the effectiveness and efficiency of our attack. Future work includes: 1) Proposing provable defenses against our attack; 2) Generalizing our regret bound to non-convex loss functions; and 3) Incorporating prior information in graphs to design better attacks.

REFERENCES

- Radhakrishna Achanta, Appu Shaji, Kevin Smith, Aurelien Lucchi, Pascal Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE TPAMI*, 2012.
- Peter Auer and Ronald Ortner. Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- Nicholas Carlini and David Wagner. Towards evaluating the robustness of neural networks. In *IEEE SP*, 2017.
- Heng Chang, Yu Rong, Tingyang Xu, Wenbing Huang, Honglei Zhang, Peng Cui, Wenwu Zhu, and Junzhou Huang. A restricted black-box adversarial framework towards attacking graph embedding models. In *AAAI*, 2020.
- Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order optimization based black-box attacks to deep neural networks without training substitute models. In *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security (AISec)*, pp. 15–26, 2017.
- Shuyu Cheng, Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Improving black-box adversarial attacks with a transfer-based prior. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 10934–10944, 2019.
- Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. Adversarial attack on graph structured data. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- Vijay Prakash Dwivedi, Chaitanya K Joshi, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *arXiv preprint arXiv:2003.00982*, 2020.
- Abraham D Flaxman, Adam Tauman Kalai, and H Brendan McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. *arXiv preprint cs/0408007*, 2004.
- ON Granichin. Stochastic approximation with input perturbation under dependent observation noises. *Vestn. Leningrad. Gos Univ*, (4):27–31, 1989.
- Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018.
- Andrew Ilyas, Logan Engstrom, and Aleksander Madry. Prior convictions: Black-box adversarial attacks with bandits and priors. In *International Conference on Learning Representations (ICLR)*, 2019. URL <https://openreview.net/forum?id=BkMiWhR5K7>.
- Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017. URL <https://openreview.net/pdf?id=SJU4ayYgl>.
- Antoine Lesage-Landry, Joshua A Taylor, and Duncan S Callaway. Online convex optimization with binary constraints. *arXiv preprint arXiv:2005.02274*, 2020.
- Sijia Liu, Swarnendu Kar, Makan Fardad, and Pramod K Varshney. Sparsity-aware sensor collaboration for linear coherent estimation. *IEEE Transactions on Signal Processing*, 63(10):2582–2596, 2015.
- Sijia Liu, Xingguo Li, Pin-Yu Chen, Jarvis Haupt, and Lisa Amini. Zeroth-order stochastic projected gradient descent for nonconvex optimization. In *IEEE GlobalSIP*, 2018.
- Jiaqi Ma, Shuangrui Ding, and Qiaozhu Mei. Black-box adversarial attacks on graph neural networks with limited node access. *arXiv preprint arXiv:2006.05057*, 2020a.
- Jiaqi Ma, Shuangrui Ding, and Qiaozhu Mei. Towards more practical adversarial attacks on graph neural networks. *Advances in neural information processing systems*, 2020b.

- Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannion, and Michael M Bronstein. Fake news detection on social media using geometric deep learning. *arXiv preprint arXiv:1902.06673*, 2019.
- Daniel Russo, Benjamin Van Roy, Abbas Kazerouni, Ian Osband, and Zheng Wen. A tutorial on thompson sampling. *arXiv preprint arXiv:1707.02038*, 2017.
- Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. Collective classification in network data. *AI magazine*, 29(3):93–93, 2008.
- Chence Shi, Minkai Xu, Zhaocheng Zhu, Weinan Zhang, Ming Zhang, and Jian Tang. Graphaf: a flow-based autoregressive model for molecular graph generation. *arXiv preprint arXiv:2001.09382*, 2020.
- James C Spall. A one-measurement form of simultaneous perturbation stochastic approximation. *Automatica*, 33(1):109–112, 1997.
- Yiwei Sun, Suhang Wang, Xianfeng Tang, Tsung-Yu Hsieh, and Vasant Honavar. Adversarial attacks on graph neural networks via node injections: A hierarchical reinforcement learning approach. In *The Web Conference*, 2020.
- Binghui Wang and Neil Zhenqiang Gong. Attacking graph-based classification via manipulating the graph structure. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 2023–2040, 2019.
- Binghui Wang, Tianxiang Zhou, Minhua Lin, Pan Zhou, Ang Li, Meng Pang, Cai Fu, Hai Li, and Yiran Chen. Efficient evasion attacks to graph neural networks via influence function. *arXiv preprint arXiv:2009.00203*, 2020.
- Felix Wu, Tianyi Zhang, Amauri Holanda de Souza Jr, Christopher Fifty, Tao Yu, and Kilian Q Weinberger. Simplifying graph convolutional networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2019a.
- Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. Adversarial examples on graph data: Deep insights into attack and defense. In *IJCAI*, 2019b.
- Kaidi Xu, Hongge Chen, Sijia Liu, Pin-Yu Chen, Tsui-Wei Weng, Mingyi Hong, and Xue Lin. Topology attack and defense for graph neural networks: An optimization perspective. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 3961–3967, 2019a.
- Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019b. URL <https://openreview.net/forum?id=ryGs6iA5Km>.
- Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In *IJCAI*, 2018.
- Xiao Zang, Yi Xie, Jie Chen, and Bo Yuan. Graph universal adversarial attacks: A few bad actors ruin graph learning models. *arXiv preprint arXiv:2002.04784*, 2020.
- Zaixi Zhang, Jinyuan Jia, Binghui Wang, and Neil Zhenqiang Gong. Backdoor attacks to graph neural networks. *arXiv*, 2020.
- Daniel Zügner and Stephan Günnemann. Adversarial attacks on graph neural networks via meta learning. *ICLR*, 2019.
- Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2847–2856, 2018.

Table 1: Dataset statistics

Datasets	#Graphs	#Ave. Nodes	#Ave. Edges	#Classes	Task
Cora	1	2,708	5,429	7	(1)
Citeseer	1	3,327	4,732	6	(1)
Pubmed	1	19,717	44,338	3	(1)
MNIST	70K	70.57	564.53	10	(2)
CIFAR10	60K	117.63	941.07	10	(2)

(1) Node classification, (2) Graph classification.

A PROOF OF LEMMA 3

Proof. We first consider the case where $N - 1 = 1$, i.e., scalar setting. Thus, we have $|u| = 1$ and u takes value in $\{-1, 1\}$. When uniformly sampling its value, we have,

$$\begin{aligned} \nabla \hat{L}(\hat{\mathbf{s}}_v) &= \mathbb{E}_{u \in \{-1, 1\}} \left[\frac{1}{\delta} L(\hat{\mathbf{s}}_v + \delta u) u \right] \\ &= \frac{1}{2} \frac{L(s + \delta)}{\delta} - \frac{1}{2} \frac{L(s - \delta)}{\delta} = L'(s). \end{aligned} \quad (15)$$

For $N - 1 > 1$ setting, the proof is similar. \square

B PROOF OF LEMMA 4

Proof. We start the proof from Lipschitz continuity assumed for $L(\cdot)$. Based on Eq. (11), we have

$$\begin{aligned} \mathbb{E}[|L(\mathbf{s}_v) - L(\hat{\mathbf{s}}_v^t)|] &\stackrel{(a)}{\leq} \mathbb{E}[C_L \|\mathbf{s}_v - \hat{\mathbf{s}}_v^t\|_2] \\ &\stackrel{(b)}{=} C_L \mathbb{E} \left[\sqrt{\sum_{i=1}^{N-1} (\mathbf{s}_v(i) - \hat{\mathbf{s}}_v^t(i))^2} \right] \\ &\stackrel{(c)}{\leq} C_L \sqrt{\sum_{i=1}^{N-1} \mathbb{E}[(\mathbf{s}_v(i) - \hat{\mathbf{s}}_v^t(i))^2]} \\ &\stackrel{(d)}{=} C_L \sqrt{\sum_{i=1}^{N-1} \mathbb{E}[(\mathbf{s}_v(i) - \mathbb{E}[\mathbf{s}_v(i)])^2]} \\ &\stackrel{(e)}{=} C_L \sqrt{\sum_{i=1}^{N-1} \mathbb{E}[\mathbf{s}_v(i)^2] - \mathbb{E}[\mathbf{s}_v(i)]^2} \\ &\stackrel{(g)}{\leq} C_L \sqrt{\sum_{i=1}^{N-1} \mathbb{E}[\mathbf{s}_v(i)^2]} \stackrel{(f)}{=} C_L \sqrt{\|\hat{\mathbf{s}}_v^t\|_1} \\ &\stackrel{(h)}{\leq} C_L \sqrt{\sqrt{N-1} \|\hat{\mathbf{s}}_v^t\|_2} \stackrel{(i)}{\leq} C_L \sqrt{\sqrt{N-1} \|\mathbf{v}_{t-1} - \eta \hat{\mathbf{g}}_{t-1}\|_2} \\ &\stackrel{(j)}{\leq} C_L \sqrt{\sqrt{N-1} \|\mathbf{v}_{t-1}\|_2 + \eta \sqrt{N-1} \|\hat{\mathbf{g}}_{t-1}\|_2} \\ &\stackrel{(k)}{\leq} C_L (N-1)^{3/4} \sqrt{1 + \frac{\eta}{\delta}}, \end{aligned}$$

where (a) is due to Lipschitz continuity of attack loss function $L(\cdot)$, (b) is due to the definition of 2-norm, (c) holds due to applying Jensen's inequality, (d) is due to the random rounding between $\hat{\mathbf{s}}_v^t$ and \mathbf{s}_v , (e) is due to definition of variance related to random variable $\mathbf{s}_v(i)$. In (g), we drop the negative term that does not affect the bound. In (f), we use the fact that each component $\mathbf{s}_v(i) \in \{0, 1\}$ is bounded by one such that $\mathbf{s}_v(i)^2 = \mathbf{s}_v(i)$. Besides, we apply the definition of 1-norm. In (h), we use the inequality $\|\hat{\mathbf{s}}_v^t\|_1 \leq \sqrt{N-1} \|\hat{\mathbf{s}}_v^t\|_2$ for 1-norm and 2-norm. In (i), we substitute the $\hat{\mathbf{s}}_v^t$ by the projected gradient descent equation and apply the non-extensive property of the projection onto convex set. In (j), we apply the triangle inequality. In (k), we expand the 2-norm using the definition of \mathbf{v}_{t-1} and $\hat{\mathbf{g}}_{t-1}$. \square

Similar claim can be found in bOGD (Lesage-Landry et al., 2020). However, the differences of bOGD (Lesage-Landry et al., 2020) and our result in Lemma 4 are two-fold: 1) we consider the bandit setting of OCO, which requires to estimate gradient while bOGD performs based on the derived gradient when observing the loss function; 2) we only assume that the loss function is Lipschitz continuous related to 2-norm while bOGD additionally assume the loss is Lipschitz continuous related to 1-norm and 2-norm.

C PROOF OF THEOREM 1

Proof. Suppose the arm set \mathcal{W} satisfies $R_1\mathbb{B} \subseteq \mathcal{W} \subseteq R_2\mathbb{B}$. We begin the proof from the regret definition in Eq. (8).

$$\begin{aligned}
\text{Reg}(T) &= \mathbb{E}\left[\sum_{t=1}^T L(\mathbf{s}_v)\right] - TL(\mathbf{s}_v^*) \\
&\stackrel{(a)}{\leq} \mathbb{E}\left[\sum_{t=1}^T L(\mathbf{s}_v)\right] - TL(\hat{\mathbf{s}}_v^*) \\
&\stackrel{(b)}{\leq} \mathbb{E}\left[\sum_{t=1}^T L(\mathbf{s}_v)\right] - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T L(\mathbf{w}) + \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T L(\mathbf{w}) - TL(\hat{\mathbf{s}}_v^*) \\
&\stackrel{(c)}{\leq} \mathbb{E}\left[\sum_{t=1}^T L(\mathbf{s}_v)\right] - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T L(\mathbf{w}) + 2\alpha T \\
&\stackrel{(d)}{\leq} \mathbb{E}\left[\sum_{t=1}^T [L(\mathbf{s}_v) - L(\hat{\mathbf{s}}_v^t)]\right] + \mathbb{E}\left[\sum_{t=1}^T L(\hat{\mathbf{s}}_v^t)\right] - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T L(\mathbf{w}) + 2\alpha T \\
&\stackrel{(e)}{\leq} \mathbb{E}\left[\sum_{t=1}^T L(\hat{\mathbf{s}}_v^t)\right] - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T L(\mathbf{w}) + 2\alpha T + C_L T(N-1)^{3/4} \sqrt{1 + \frac{\eta}{\delta}} \\
&\stackrel{(f)}{\leq} \mathbb{E}\left[\sum_{t=1}^T \tilde{L}(\mathbf{v}^t)\right] - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T \tilde{L}(\mathbf{w}) + \mathbb{E}\left[\sum_{t=1}^T [L(\hat{\mathbf{s}}_v^t) - \tilde{L}(\mathbf{v}^t)]\right] \\
&\quad + \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T [\tilde{L}(\mathbf{w}) - L(\mathbf{w})] + 2\alpha T + C_L T(N-1)^{3/4} \sqrt{1 + \frac{\eta}{\delta}}.
\end{aligned}$$

In (a), we consider that $L(\mathbf{s}_v^*) \leq L(\hat{\mathbf{s}}_v^*)$ holds for the relaxed variable $\hat{\mathbf{s}}_v^*$. In (b), we consider the $(1-\alpha)$ -projection in our algorithm (i.e., line 6). In (c), we apply with the result of Lemma 6. In (d), we consider the relaxed variable $\hat{\mathbf{s}}_v^t$. In (e), we apply with the result of Lemma 4. In (f), we adapt the terms to make convenience of using the result of Lemma 5.

We define a smooth function $\tilde{L}(\cdot)$ to approximate the original loss function $L(\cdot)$ as follows,

$$\tilde{L}(\mathbf{v}) = \mathbb{E}_{\mathbf{u} \in \mathbb{B}} [L(\mathbf{v} + \delta \mathbf{u})]. \quad (16)$$

When δ is small, $\tilde{L}(\mathbf{v}) \approx L(\mathbf{v})$. Thus, $\tilde{L}(\mathbf{v})$ and $L(\mathbf{v})$ share similar gradient at any \mathbf{v} . According to the Lipschitz continuity in Eq. (11), we have,

$$|\tilde{L}(\mathbf{w}) - L(\mathbf{w})| \leq C_L \delta, \quad (17)$$

for all $\mathbf{w} \in \mathcal{W}$. Due to $\hat{\mathbf{s}}_v^t = \mathbf{v}^t + \delta \mathbf{u}^t$, we obtain,

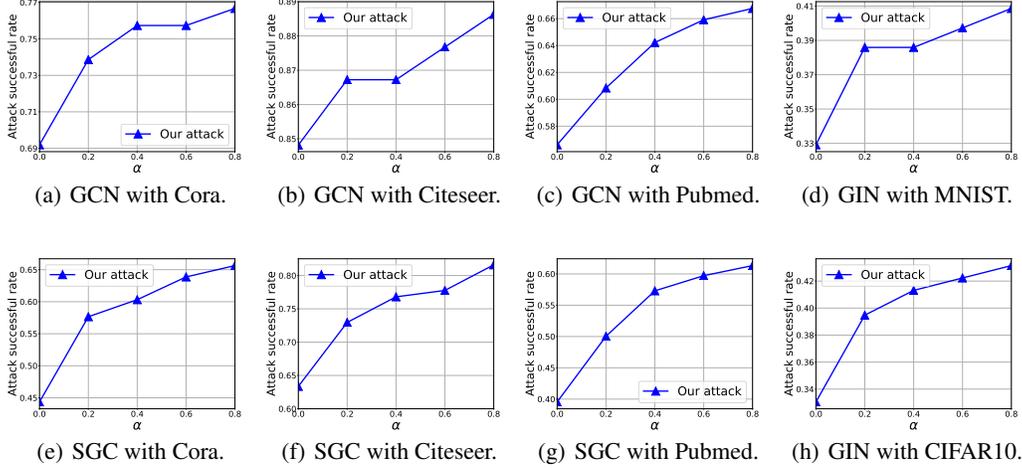
$$|L(\hat{\mathbf{s}}_v^t) - L(\mathbf{v}^t)| = |L(\mathbf{v}^t + \delta \mathbf{u}^t) - L(\mathbf{v}^t)| \leq C_L \delta, \quad (18)$$

where we further have by applying triangle inequality,

$$|L(\hat{\mathbf{s}}_v^t) - \tilde{L}(\mathbf{v}^t)| \leq |L(\hat{\mathbf{s}}_v^t) - L(\mathbf{v}^t)| + |L(\mathbf{v}^t) - \tilde{L}(\mathbf{v}^t)| \leq 2C_L \delta, \quad (19)$$

for all $t \in [T]$. To continue bounding $\text{Reg}(T)$ in Eq. (C).(f), we have,

$$\begin{aligned}
\text{Reg}(T) &\leq \sum_{t=1}^T \tilde{L}(\mathbf{v}^t) - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T \tilde{L}(\mathbf{w}) + 3C_L \delta T + 2\alpha T \\
&\quad + C_L T(N-1)^{3/4} \sqrt{1 + \frac{\eta}{\delta}}.
\end{aligned} \quad (20)$$

Figure 8: Sensitivity of hyperparameter α .

At last, we bound the regret $\sum_{t=1}^T \tilde{L}(\mathbf{v}^t) - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T \tilde{L}(\mathbf{w})$, which corresponds to the stochastic gradient descent in line 6 of Algorithm 1 where the gradient is $\frac{N-1}{\delta} L(\hat{\mathbf{s}}_v^t) \mathbf{u}^t$ and the feasible domain is $(1-\alpha)\mathcal{W}$. Note that we have $(1-\alpha)\mathcal{W} \subseteq \mathcal{W} \subseteq R_2\mathbb{B}$ for any $0 < \alpha < 1$. It is easy to see that $\|\frac{N-1}{\delta} L(\hat{\mathbf{s}}_v^t) \mathbf{u}^t\|_2 \leq \frac{N-1}{\delta}$. According to Lemma 5, we have

$$\sum_{t=1}^T \tilde{L}(\mathbf{v}^t) - \min_{\mathbf{w} \in (1-\alpha)\mathcal{W}} \sum_{t=1}^T \tilde{L}(\mathbf{w}) \leq \frac{N-1}{\delta} R_2 \sqrt{T}. \quad (21)$$

In summary, we can bound $\text{Reg}(T)$ as,

$$\begin{aligned} \text{Reg}(T) &\leq \frac{N-1}{\delta} R_2 \sqrt{T} + 3C_L \delta T + 2\alpha T + C_L T (N-1)^{3/4} \sqrt{1 + \frac{\eta}{\delta}} \\ &\stackrel{(a)}{\leq} \frac{N-1}{\delta} R_2 \sqrt{T} + 3C_L \delta T + 2\alpha T + C_L T (N-1)^{3/4} \sqrt{\frac{R_2}{(N-1)\sqrt{T}}} \\ &\stackrel{(b)}{\leq} \frac{N-1}{\delta} R_2 \sqrt{T} + 3C_L \delta T + 2\alpha T + C_L \sqrt{R_2} (N-1)^{1/2} T^{3/4} \\ &\stackrel{(c)}{\leq} \frac{N-1}{\delta} R_2 \sqrt{T} + \frac{3(C_L R_1 + 1)}{R_1} \delta T + C_L \sqrt{R_2} (N-1)^{1/2} T^{3/4} \\ &\stackrel{(d)}{\leq} 2\sqrt{\frac{3R_2(C_L R_1 + 1)(N-1)}{R_1}} T^{3/4} + C_L \sqrt{R_2} (N-1)^{1/2} T^{3/4} \\ &\stackrel{(e)}{=} \Lambda \sqrt{N-1} T^{3/4} = \mathcal{O}(\sqrt{N} T^{3/4}). \end{aligned} \quad (22)$$

In (a), we set the step $\eta = \frac{R_2}{\sqrt{T(N-1)/\delta}} - \delta$. We further simplify the last term, which leads to (b).

In (c), we set $\alpha = \frac{\delta}{R_1}$, which can ensure $\hat{\mathbf{s}}_v^t \in \mathcal{W}$ (Flaxman et al., 2004). In (d), we set parameter $\delta = T^{-1/4} \sqrt{\frac{R_1 R_2 (N-1)}{3(C_L R_1 + 1)}}$ to minimize the r.h.s. of (c). In (e), we define the leading constant as

$\Lambda = 2\sqrt{\frac{3R_2(C_L R_1 + 1)}{R_1}} + C_L \sqrt{R_2}$. Based on above deduction, we can prove this theorem. \square

D MORE EXPERIMENTAL RESULTS

In the following sections, we conduct experiments to evaluate the sensitivities of hyperparameters, i.e., projection scale α and update step δ of the prior vector. We set the default queries as 50 for both node classification and graph classification, i.e., $T = 50$. The default number of the perturbed edges is set to 2 and 6 for node classification and graph classification, respectively. Other parameter settings are consistently configured as the main experiments like learning rate η . α is ranged in

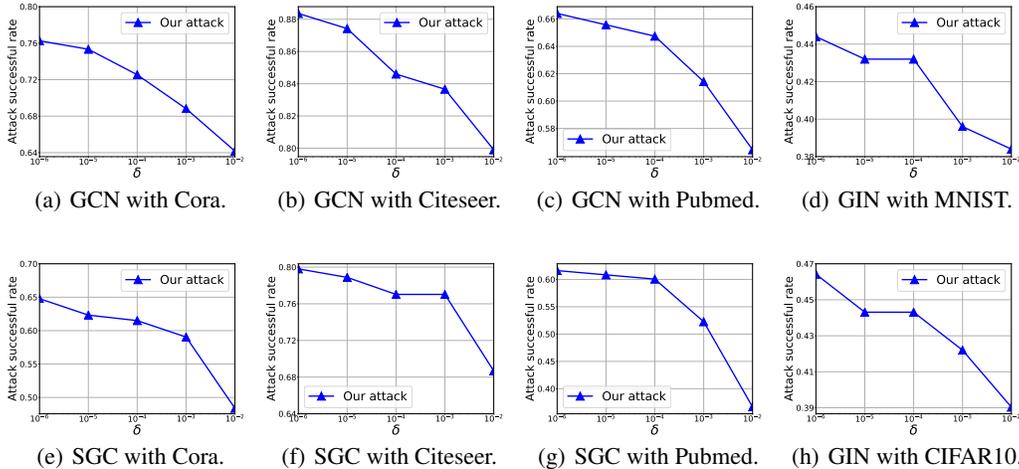


Figure 9: Sensitivity of hyperparameter δ .

$\{0, 0.2, 0.4, 0.6, 0.8\}$ and δ is ranged in $\{10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\}$. The results are shown in Figure 8 and Figure 9. Overall, the derivation of the attack successful rate is no more than 10% under different hyperparameters, which demonstrates that our attack is stable and robust to sensitivity.

Sensitivity evaluation of hyperparameter α . In Figure 8, we evaluate the sensitivity performance of hyperparameter α over different GNNs (i.e., GCN, SGC and GIN) and different datasets (i.e., Cora, Citeseer, Pubmed, MNIST and CIFAR10). We can observe that the attack successful rate increases as α increases. This can be explained that the projected domain in PGD becomes larger when α goes larger. Thus, there is a high probability that the optimal perturbation is located in the larger projected domain.

Sensitivity evaluation of hyperparameter δ . In Figure 9, we evaluate the sensitivity performance of hyperparameter δ over different GNNs (i.e., GCN, SGC and GIN) and different dataset (i.e., Cora, Citeseer, Pubmed, MNIST and CIFAR10). We can observe that the attack successful rate decreases as δ increases. This can be explained that the estimated gradient $\hat{\mathbf{g}} = \frac{N-1}{\delta} L(\hat{\mathbf{s}}_v^t) \mathbf{u}^t$ becomes more inaccurate when δ goes larger. Consequently, it is impossible to derive the optimal perturbation with an inaccurate gradient.