Are Gradients on Graph Structure Reliable in Gray-box Attacks?

Zihan Liu
[1] Zhejiang University
Hangzhou, Zhejiang, China
[2] AI Lab, Westlake University
Hangzhou, Zhejiang, China
liuzihan@westlake.edu.cn

Yun Luo
Westlake University
Hangzhou, Zhejiang, China

Lirong Wu
AI Lab, Westlake University
Hangzhou, Zhejiang, China

Siyuan Li
AI Lab, Westlake University
Hangzhou, Zhejiang, China

Zicheng Liu
AI Lab, Westlake University
Hangzhou, Zhejiang, China

Stan Z. Li*
AI Lab, Westlake University
Hangzhou, Zhejiang, China
stan.zq.li@westlake.edu.cn

ABSTRACT

Graph edge perturbations are dedicated to damaging the prediction of graph neural networks by modifying the graph structure. Previous gray-box attackers employ gradients from the surrogate model to locate the vulnerable edges to perturb the graph structure. However, unreliability exists in gradients on graph structures, which is rarely studied by previous works. In this paper, we discuss and analyze the errors caused by the unreliability of the structural gradients. These errors arise from rough gradient usage due to the discreteness of the graph structure and from the unreliability in the meta-gradient on the graph structure. In order to address these problems, we propose a novel attack model with methods to reduce the errors inside the structural gradients. We propose edge discrete sampling to select the edge perturbations associated with hierarchical candidate selection to ensure computational efficiency. In addition, semantic invariance and momentum gradient ensemble are proposed to address the gradient fluctuation on semantic-augmented graphs and the instability of the surrogate model. Experiments are conducted in untargeted gray-box poisoning scenarios and demonstrate the improvement in the performance of our approach.

CCS CONCEPTS

• **Computing methodologies** → *Supervised learning by classification.*

KEYWORDS

graph adversarial attack, untargeted gray-box attack, poisoning attack

1 INTRODUCTION

Graph Neural Networks (GNNs) demonstrate excellent performance on various applications with structural, relational data [25], such as traffic [19], recommendation systems [5], and social networks [18]. As the prospects for the applications of GNNs expand rapidly, their reliability and robustness are beginning to be of interest. Several works have presented experimental evidence that GNNs are vulnerable to adversarial attacks [3, 26, 27]. They design undetectable perturbations which successfully mislead GNNs' prediction of targeted nodes or degrade the performance on a global scope. Subsequently, many works have been carried out around the graph adversarial attack and defense [15, 22].

Gray-box attacks allow attackers to access the training labels from the victim model. The attacker aims to damage the prediction of the victim model by injecting indistinguishable perturbations into the graph. The attacker should search for vulnerable edges and attack them by modifying the graph. In the field of adversarial attack, gradients are widely-used for attacking attributes that are spatially continuous [22]. However, for graphs, the sparseness and discreteness of the graph structure make it challenging to perturb the graph structure in the way of Fast Gradient Sign Method (FGSM) [4] or Projected Gradient Descent (PGD) [12]. To solve this problem, Zugner et al. [27] firstly introduce the meta-gradient on the graph structure to determine the perturbation. The attacker chooses one edge at a time to perturb based on the saliency of the gradient and iterates this step until the entire attack budget is consumed. Subsequent works focus on improving the attack strategy after deriving the gradient and the surrogate model [9, 10]. However, few works focus on whether the saliency of the gradient on the graph structure is reliable.

Meta-gradients are demonstrated to be noisy in attribution problems [17]. The gradients on the graph structure originate from the aggregation of node features, which means noises are equally propagated into the structural gradients. Moreover, edge flipping changes the value of edges across a considerable step size (i.e.,

adding edges are from 0 to 1 and deleting edges are from 1 to 0). Existing methods select the edge to be flipped based on the saliency of the gradient. It is worth noting that, during the edge flipping, the structural gradient varies since the aggregation of node features is influenced by the edge value. We consider edge perturbation to be a problem of modeling continuous gradient distributions rather than a discrete problem. However, the structural gradients are imprecisely assumed to be constant, ignoring the variance in the edge-flip interval.

This paper points out the gradient errors that negatively affect the untargeted gray-box attacks on graph structure. The discreteness of edges leads previous works to consider edge perturbation as a discrete problem about gradients. However, the gradient at the current edge state commonly used in previous works is inaccurate for describing the gradient over the edge-flip interval. We propose to transform the edge perturbation from a discrete problem to an approximation of a continuous problem and propose *edge discrete sampling* to reduce this error. Edge discrete sampling calculates the gradient of the transition process between the edge-flip interval in batches. It reduces the error from the discrete approximation to the continuous gradient distribution, which performs a more accurate structural gradient estimation. Since the space of edge perturbations is about the square of the number of nodes, the computational cost is unacceptable if the edge discrete sampling is performed for the whole space. *Hierarchical candidate selection* is then proposed to reduce the computational complexity. It retains a bag of candidate edges that are more likely to be effective than processing the whole space. In this step, we exploit the saliency of the gradient on the graph structure while the error still exists. The random initialization of parameters leads to variance in surrogate model training, which affects the structural gradients via back-propagation. Besides, the error occurs in gradient fluctuations on semantically identical graph augmentations. These two errors are to be considered since gray-box attacks focus on attack transferability, so the gradient information from the surrogate model is expected to be more general and representative. To address these errors, we first propose *momentum gradient ensemble* to mitigate the instability of the structural gradients provided by the surrogate model at each attack iteration. Then, we propose surrogate *semantic invariance* based on graph augmentation in a limited semantic range. These two methods of reducing the error on the gradient of the graph structure allow hierarchical candidate selection to provide better quality edge candidates. Candidate selection ensures that the computational cost of edge discrete sampling does not grow exponentially, allowing the attacker to perform poisoning attacks.

The contributions of this paper are summarized as follows:

- We analyze the errors in structural gradients caused by model instability and the discreteness of graph structure in untargeted gray-box attacks.
- We propose edge discrete sampling to approximate the continuous distributions of gradient over the edge-flip interval. Hierarchical candidate selection is proposed to ensure that the computational complexity does not explode.

---

†Corresponding author: Stan Z. Li.

- We propose semantic invariance and momentum gradient ensemble to address the gradient fluctuation on semantic graph augmentation and the instability of the surrogate model.
- We demonstrate the improvement of our approach and prove each module's effectiveness in the ablation study.

2 PRELIMINARIES

Before presenting the methodology and demonstrating the experiments, we first introduce the notations and backgrounds in this section. The notations used in the following sections can be referenced in Subsection 2.1. Subsection 2.2 introduces how to obtain the gradient on the graph structures by attacking loss in a generic attack strategy with edge perturbations.

2.1 Notations

A graph G is represented as $G = (V, E, X)$, where $V = \{v_1, v_2, ..., v_n\}$ is the node set, $E \subseteq V \times V$ is the edge set, and X is the attribution set. In node classification tasks, a vertex v_i is a sample with its features $x_i \in \mathbb{R}^d$ and label $y_i \in \mathbb{R}^c$. A total of N samples in the graph are distributed across c classes. The edges are represented as a binary adjacent matrix $A = \{0, 1\}^{N \times N}$, where $A_{ij} = 1$ if $(i, j) \subseteq E$. We can also represent a graph with an adjacent matrix and nodes' attributes $G = (A, X)$. For a gray-box attack, the attacker constructs a surrogate model, denoted by f_θ, to simulate the process of the victim model being attacked. The prediction of f_θ is denoted by the probability distribution P_{v_i}. We use $\mathcal{L}$ to represent the loss functions. The perturbed graph is indicated as $G' = (A', X')$ to distinguish the perturbed graph from the original graph. The attack is limited by a budget Δ.

2.2 Edge Perturbations

In the case of untargeted edge perturbation, the attacker is restricted to perturbing only by modifying the adjacency matrix (i.e., flipping edges). The l_0 norm of the changes in the perturbed adjacency matrix with respect to the original one is bounded by the attacker's budget Δ. For an undirected graph, Δ is set as:

$$\|A - A'\|_0 \leq 2\Delta, \tag{1}$$

where the budget Δ is generally equal to or less than 5% of the number of edges in the original graph.

Gradient-based attack models now become mainstream methods of edge perturbations on the graph structure [9, 23, 27]. In contrast to the gradient-based attacks widely used in computer vision [24], the discrete graph structure restricts the gradient from being added directly to the adjacency matrix. The gradient on the adjacent matrix (i.e. graph structure) [27] A^{grad} can be derived by the following equations:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} \, \mathcal{L}_{ce}(f_\theta(G), Y), \tag{2}$$

$$A^{grad} = \nabla_A \mathcal{L}_{atk}(f_{\theta^*}(G)), \tag{3}$$

where $\mathcal{L}_{atk}$ is the attack loss function and f_{θ^*} is the properly trained surrogate model. To facilitate understanding, we elaborate in a more intuitive way of deriving A^{grad}. First, a GNN surrogate model f_{θ^*} is trained until it fits the training samples. Subsequently, the attack loss is backpropagated through the surrogate model generating gradients on the input adjacency matrix. The attack loss is

expressed as $\mathcal{L}_{atk} = -\mathcal{L}_{CE}$, which is a negative cross-entropy loss. For the edge between nodes v_i and v_j, if $A_{i,j} = 1$ and $A_{i,j}^{grad} < 0$, or if $A_{i,j} = 0$ and $A_{i,j}^{grad} > 0$, then flipping edge $E_{i,j}$ is considered as a perturbation that has the potential to negatively affect the victim model. Among these edges, the one with the most significant gradient value is considered the most desirable perturbation for the current graph. The process of perturbing the graph using the gradient information can be represented as:

$$G'_t = \phi(\nabla_A \mathcal{L}_{atk}(f_{\theta^*}(G'_{t-1})), A'_{t-1}), \tag{4}$$

where ϕ denotes the strategy for choosing the edge to be attacked. The factors that influence the perturbation include the attack loss $\mathcal{L}_{atk}$ as well as the strategy ϕ and the surrogate model f_{θ^*}.

3 METHODOLOGY

This section introduces the errors of the gradient on the graph structure and the methods to solve these errors. Section 3.1 first analyzes the error caused by interpreting edge perturbations as a discrete problem and proposes the solution *edge discrete sampling*. To rationalize the computational cost, we propose *hierarchical candidate selection*. It selects a bag of edge candidates based on the meta-gradient on the graph structure so that edge discrete sampling requires only a small number of edges to be processed in batches. When we rethink the meta-gradient on the graph structure, we find room for improving the reliability of the structural gradient. We analyze and discuss this part in Sections 3.2 and 3.3 and give the corresponding solutions. Section 3.4 describes the overall attack flow.

3.1 Error Caused by Edges' Discreteness

As indicated by Eq.2&3, the gradient on A_{ij} is the partial derivative of the attack loss $\mathcal{L}_{atk}$ to the adjacent matrix. In the existing approaches, the attackers treat the gradients on graph structure as a discrete problem of choosing the perturbations directly based on the saliency of the gradient [9, 27]. This means that the previous approach assumes that the gradient value is maintained at its value on the original state of the edge during the edge flipping (i.e., the state of an edge turns from 0 to 1 or from 1 to 0). We give an example of the gradient approximation error introduced by this approach in Fig.1. In contrast, we consider the gradient used to determine the edge perturbation as a continuous problem with continuous distribution approximation.

In Fig.1, the blue curve indicates the gradient distribution with edge value between 0 and 1, and the red line indicates the approximation of this distribution by the existing approaches. The light yellow shading indicates the error of the previous method compared to the ground truth continuous gradient distribution. Assuming that the interval for deriving the continuous distribution shown in the blue curve is δ, the time complexity required to compute the gradient distribution for each edge is $O(N^2/\delta)$. The problem is that the computational cost of calculating the gradient distribution for all edges is unacceptable. In order to reduce the computational complexity, we propose hierarchical candidate selection. As mentioned in Section 2.2, the gradients on the adjacency matrix A^{grad} can be derived from a trained surrogate model by back-propagation of the attack loss. The first candidate selection step
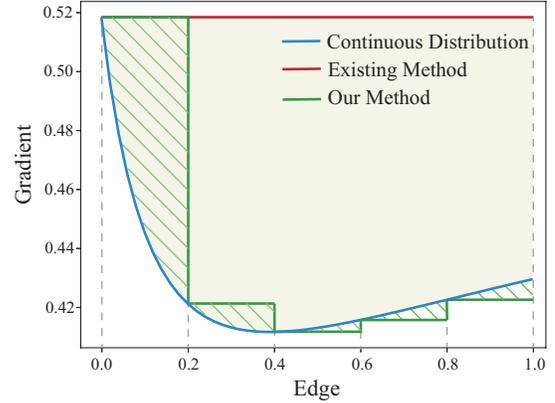
Figure 1: Illustration of the error generated by our method and the existing method in estimating gradient distribution on edge. The blue curve shows the continuous distribution. The thick lines in green and red are the estimation of our method and the existing method, while their estimation error to the continuous distribution is presented in thin green slash and light grey background.

removes edges (set the saliency of the gradient to 0) where the sign of the gradient and the change of the edge does not match [27], after which the gradient matrix is processed to be $\tilde{A}^{grad}$. For the second step, we sort the rest of candidates according to the saliency of the gradient values in descending and retain the top C edges as candidates:

$$S = \{e_{(u_1,v_1)}, ..., e_{(u_c,v_c)}\} = \underset{e=(u,v)}{top_C}(\tilde{A}^{grad}), \tag{5}$$

where function top_C extracts C edge candidates of high saliency from the gradient matrix. For these candidates, we reduce the time complexity to $O(C/\delta)$, where $C << N^2$. Up to this point, the time complexity of the algorithm is still enormous because δ is tiny. Therefore, in order to reduce the error while being able to ensure computational efficiency, we propose a discrete sampling method to approximate the continuous gradient distribution. The expression of the this approximation is:

$$
\begin{aligned}
g_{uv}^{int} &= \int_0^1 [\nabla_A \mathcal{L}_{atk}(f_{\theta^*}(A, X))]_{u,v} \, dA_{uv} \\
&\approx \lambda \sum_{s=1}^{1/\lambda} [\nabla_A \mathcal{L}_{atk}(f_{\theta^*}(A_{uv} = s\lambda, X))]_{u,v},
\end{aligned}
\tag{6}
$$

where g_{uv}^{int} represents the integral gradient as the edge flips from 0 to 1 and $f_{\theta^*}(A_{uv} = s\lambda, X)$ is the result of modifying A_{uv} to a transitional value $s\lambda$ without retraining. The solid thick green line in Fig.1 indicates the approximation by our algorithm, and the area in the thin green slash indicates the error caused by ours. Compared with the error of the previous method indicated by the blue shading, our method substantially reduces the error generated in the edge-flipping process. Based on the above algorithm, the time complexity of our method decreases to $O(C/\lambda)$. Considering that a perturbation can add or delete an edge, we adopt the one with
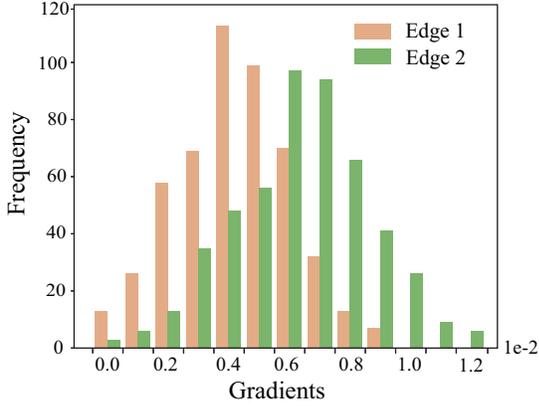
Figure 2: Possibility distributions of the gradient on two example edges. The x-axis is the interval of the gradient values, and the y-axis is the frequency of the gradient in the interval over 500 replicate experiments.

the highest saliency of integrated gradient as a perturbation:

$$e' = \underset{e=(u,v)}{top_1} \left(\{ (1 - 2A_{uv}) \, g_{uv}^{int} \mid (u,v) \in S \} \right), \tag{7}$$

where e' denotes the selected edge to be perturbed, $(1 - 2A_{uv})$ is to invert the value of g_{uv}^{int} for those candidates to delete an edge.

To further increase the computational efficiency, we introduce batch processing of candidates. A batch contains batch size bs candidates. Since the gradients on the adjacent matrix come from the aggregation of node features, for a general 2-layer graph neural network, the state of one edge will influence the gradient on the other when the same node joints two edges. In other words, simultaneously changing the states of both edges causes a small amount of error. However, for a batch of candidates selected from N^2 space, the probability that the candidates happen to be connected is minimal. Therefore we adopt batch processing which reduces the time complexity to $O(C/(\lambda * bs))$.

The method mentioned in Section 3.1 has *a high dependence on the gradients on the adjacency matrix* A^{grad}. Hierarchical candidate selection selects C from the N^2 space, which requires strong reliability and stability of the gradients. In Sections 3.2 and 3.3, we discuss the errors present in A^{grad} as well as their solutions.

3.2 Error Caused by Uncertainty of Model Optimization

When the samples are not dense enough to describe the manifold of the data, the model is prone to fall into local optima. The local optima of a model based on backpropagation optimization is related to the initialization of the mapping function of the neural network (i.e., the initialization of the learnable weights). For the surrogate model, it tends to perform differently after retraining with different parameter initializations (i.e., in Equation 2 different θ leads to different θ^*). We give an example to verify the uncertainty of structural gradients A^{grad}, shown in Fig.2.

Fig.2 shows the gradients on two edges of the consistent graph after retraining the surrogate model with different weight initializations. We can see that the gradient expectation of Edge 2 is around 0.007, which is higher than 0.004 of Edge 1. We conduct experiments showing that attacking Edge 2 is more effective than attacking Edge 1. However, due to the randomness of weight initialization, the gradient of Edge 1 is possible to be higher than that of Edge 2, thus misleading the attacker to make a wrong judgment.

To enhance the reliability of A^{grad} at each attack iteration, we expect each structural gradient obtained from the surrogate model to be as close as possible to its expectation in such distribution shown in Fig.2. To this end, there is an easy way to solve this problem using an ensemble algorithm. We sample several weight initializations and integrate the gradients from retrained surrogate models to approximate the expectations of the structural gradients, reducing the probability of the gradient being sampled to outliers. The expression of the gradient ensemble is:

$$A^{grad} = \frac{1}{k} \sum_{i=1}^{k} A_{\theta_i^*}^{grad}, \tag{8}$$

where θ_i^* represents the parameters of the surrogate model after training under initialization θ_i and k denotes the number of integrations. This is an ensemble-based method that randomly initializes the model parameters with the identical uniform distribution between $[0, 1]$. However, retraining the surrogate model is unacceptable, for it causes k times the computational cost to be spent.

Considering that the attack on graph edges is a perturbation-by-perturbation iterative process, we propose a momentum gradient ensemble as a more efficient solution. The variation between perturbed graphs G_t' and G_{t-1}' is limited. The structural gradient A_t^{grad} at iteration t can reduce the instability of the gradient brought by each retraining of the surrogate model by fusing the structural gradient at previous iterations. The derivation of the structural gradient at attack iteration t is redefined as:

$$A_t^{grad} = \nabla_A \mathcal{L}_{atk}(f_{\theta^*}(G_t')) + p A_{t-1}^{grad}, \tag{9}$$

where p is the coefficient of the momentum term. Compared with the ensemble method in Equation 8, our momentum-based method consumes no additional computational cost. It makes full use of the structural gradients from previous iterations and avoids retraining the surrogate model multiple times in a single iteration.

3.3 Error Caused by Model's Unrobustness

The gradients of continuous data features are demonstrated to be noisy [17]. Similar to data features, the graph structure is explicitly involved in the forward process in GNNs. Taking a 2-layer GCN as an example, the expression of the forward process of graph neural network and the structural gradient is shown below.

$$f_\theta(A, X) = softmax(\hat{A}\tau(\hat{A}XW^{(0)})W^{(1)}) \tag{10}$$

$$\mathcal{L}_{atk} = -\mathcal{L}_{CE} = \log P(y_i | f_\theta(A, X)), \tag{11}$$

$$\nabla_{A_{ij}} \mathcal{L}_{atk} = \frac{\partial \log P(y_i | f_\theta)(A, X)}{\partial A_{ij}}, \tag{12}$$

where $\hat{A}$ is the normalized adjacent matrix, τ is the activation function, W is the weight matrix, $\mathcal{L}_{atk}$ equals to $-\mathcal{L}_{ce}$ and $P(y_i' | f_\theta)$ represents the prediction at label class y_i. It can be seen from Eq.10,11&12
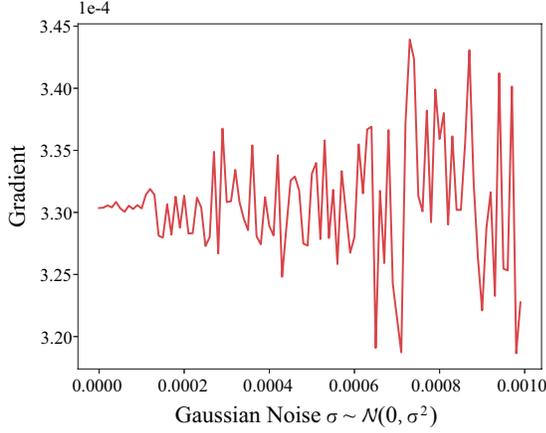
Figure 3: Illustration of the noisy gradients on graph structures. The curve represents the gradient change on an example edge under Gaussian noise on node attributes with variances σ.

that the derivation of the structural gradient involves the features of the nodes (which is related to the message passing mechanism in GNNs), resulting in the noise on the features being able to be propagated to the structural gradient. The graph structure explicitly participates in the forward process of the model, so both input features and model parameters are influencing factors of the structural gradient. Therefore, similar to the sample features, the structural gradient is noisy.

Fig.3 is an example which shows the gradients on an example edge on Citeseer [14]. It is a citation network in which 10^{-3} can be considered as a non-disturbing semantic perturbation for node attributes. We inject Gaussian noise on the attributes of nodes v_i and v_j and their 1-hop neighbors, where the horizontal axis of Fig.3 is the expectation of Gaussian noise and the vertical axis is the value of the gradient. As can be seen in Fig.3, Gaussian noise with a standard deviation of 10^{-3} can make the gradient noisy. Considering the need for transferability for gray-box attacks, the instability of the surrogate model in the semantic range will affect the attacker's estimation of the retrained victim model.

It is worth noting that the noise is added to the node attributes rather than the graph structure. We consider adding noise to the graph structure would harm the graph homophily. For multi-class tasks, the majority of the noises are added between inter-class edges while the minority are added between intra-class edges. This leads to the fact that adding Gaussian noise to the graph structure means *decreasing the homophily ratio of the graph*. Therefore, adding Gaussian noise on the node attributes is relatively reasonable than the graph structure.

In order to maintain the consistency of gradients in semantic graph augmentations, we propose a semantic invariance strategy based on graph augmentation. Computing the expectation of gradient over a high-dimensional augmentation space is intractable, so

we approximate this expectation by averaging sampled augmentations in the local space of the original graph. The semantic invariance is expressed mathematically as:

$$A_{si}^{grad} = \frac{1}{n} \sum_{i=1}^{n} \nabla_A \mathcal{L}_{atk}(f_{\theta^*}(A, X + \mathcal{N}(0, \sigma^2))), \qquad (13)$$

where n is the number of samples, and $\mathcal{N}(0, \sigma^2)$ represents a matrix of Gaussian noise with variable standard deviation σ. This method preserves the invariance of gradients on semantically consistent augmented graphs. The variance of Gaussian noise is a hyperparameter specific to the dataset. We empirically select the hyperparameter with the optimal performance on the validation set for testing based on grid search. For example, on Citeseer, the variance of the Gaussian noise is set to 5e-4. We verify that Gaussian noise at this variance has little effect on the classification accuracy of GNN, and that is how we define the noise as semantically non-disturbed.

3.4 The Overall Attack Model

This section describes the implementation of the above three error reduction methods in the attack model. Algorithm 1 is used to illustrate the whole flow of our proposed attack.

Algorithm 1 Attack Pipeline

Input: Original graph $G = (A, X)$, label of training data Y, attack budget Δ, momentum coefficient p;

Output: Graph with edge perturbations $G' = (A', X)$;

1: $t = 0; A' = A$;
2: **while** $\|A - A'\|_0 < 2\Delta$ **do**
3:     Train surrogate model f_{θ^*} with $G' = (A', X)$ and Y;
4:     Derive the semantic invariance structrual gradient A_{si}^{grad};
5:     stack the structural gradient with that at the previous iteration in the way of momentum $A_t^{grad} = A_{si}^{grad} + pA_{t-1}^{grad}$ ($A_0^{grad} = A_{si}^{grad}$ at $t = 0$);
6:     Filter the candidates whose gradient symbols and edge states mismatch, and take the absolute value to get the saliency $\tilde{A}_t^{grad}$;
7:     Retain top C edges in $\tilde{A}_t^{grad}$ as candidate set S_t and $S_{t \cdot copy}$
8:     **while** S_t is not none **do**
9:         Sample a batch b of edge candidates from S_t;
10:         Derive the integral gradient g_{uv}^{int} of each edge e_{uv} in batch b at the edge flip interval [0,1];
11:         Remove the batch b from S_t;
12:     **end while**
13:     Extract the edge e_{mn} with the most significant g_{mn}^{int} in edge candidates $S_{t \cdot copy}$;
14:     Update the perturbed graph $A'[m, n] = 1 - A'[m, n]$ and $A'[n, m] = 1 - A'[n, m]$;
15:     $t = t + 1$;
16: **end while**

The whole attack process is decomposed into Δ iterations, with one edge perturbed in each iteration. At the beginning of each iteration, a surrogate model is trained using the perturbed graph G' (Equation 2) to simulate the victim model under poisoning attacks.

With the trained surrogate model, we derive the semantic invariance structural gradient A_{si}^{grad} following Equation 13. Then we add the momentum term onto the structural gradient to minimize the error arising from the model's local optima (Equation 9). Line 6-7 in the Algorithm describe the hierarchical candidate selection, which is to choose candidate set S_t from the N^2 space. Line 8-11 in Algorithm perform edge discrete sampling (Equation 6) to approximate the continuous gradient distribution for each edge in S_t. Afterward, we select the edge with the most significant integral gradient g_{mn}^{int} as the perturbed edge at the t-th iteration. Finally, The perturbed graph structure is updated, and the $t + 1$ iteration is started.

This algorithm ensures that the surrogate model needs to be trained only once in each iteration. Therefore, it achieves high computational efficiency while reducing the errors in structural gradients.

4 EXPERIMENTS

We present experiments to demonstrate the effectiveness of our proposed attack model, named AtkSE [1] (Attacking by Shrinking Errors). The experimental settings are detailed in Section 4.1. In the following two sections, we compare our approach with other gray-box untargeted poisoning attack methods and provide ablation studies to verify our proposed improvements' validity. In Section 4.4, we provide the gradients' distribution on the edges' values between 0 and 1 to prove the necessity for our proposed approximation method for continuous gradient distribution.

4.1 Experimental Settings

4.1.1 Datasets. In this paper, we use the citation networks Citeseer [14], Cora [13] and Cora-ML [13] as well as the social network Polblogs [1] as the datasets. Consistent with the experimental setup of baselines, we randomly divide the dataset into 10% of labeled nodes and 90% of unlabeled nodes. The labels of the unlabeled nodes are not provided to the attacker or the victim model, and they are only used when evaluating the performance of the victim model.

4.1.2 Victim Models. The widely-used victim model is GCN [8] used in baseline papers. This paper extends GraphSage [6] as a more advanced GNN victim model. It is worth noting that there are some gray-box attacks, such as [9] and [10], where the network architecture of the victim model is considered to be known. The attack scenario in this paper considers that the victim model's architecture is unknown to study the poisoning attack's transferability better. Therefore, the GCN victim model differs from the surrogate model in linearity and number of neurons.

We uniformly use a 5% perturbation rate as the attack budget for imperceptibility needs. We repeat the experiments ten times for each experimental scenario and present the mean and variance of each set of experiments in the results. To ensure the fairness of the experiments, we test the perturbed graphs generated by each method with uniform and independent test files.

4.1.3 Baselines. Random, DICE [20], EpoAtk [9], Meta-Train & Meta-Self [27] are used as baselines in the experiments.

[1]https://github.com/Zihan-Liu-00/AtkSE

Random removes or add edges randomly.
DICE randomly removes edges between nodes from the same class or adds edges between nodes from different classes.
EpoAtk is originally a white-box attack model, transferred to the gray-box attack in our experiments. It proposes an exploration strategy where the attacker chooses edges from a set of candidates.
Meta-Self and Meta-Train consider the adjacent matrix as hyperparameters to optimize via meta-learning. Two attack models differ in the node subset used to calculate attack loss.

4.1.4 Hyperparameters. In the implementation of our attack model, the interval of edge discrete sampling λ is set to 0.2, the number of edge candidates C is set to 64, the batch bs is set to 16, the momentum coefficient is set to 0.8, and n in semantic invariance is 5.

4.2 Performance of AtkSE

Table 1 shows the comparison of our approach with baselines on various datasets and victim models. Among the eight experiments, our proposed AtkSE outperforms baselines in seven of them. Meta-Self is the most effective baseline in other gradient-based baselines, followed by Meta-Train and EpoAtk. The randomness-based methods DICE and Random are the worst. Our method ranks second behind the best effect in Polblogs/GraphSage. When the victim models are GraphSages, our method improves the attack success rate over the second-best method by 1.6%, 0.9%, and 0.1% on datasets Cora, Cora-ML, and Citeseer, respectively. On Polblogs, our method ranks second, below the first place by 1.5%. This result may be due to the independence of victim models' representation learning and the difficulty of transferring attacks. When the victim models are GCNs, our method is better than other methods across the board. The attack performance of our method improves 2.1%, 1.8%, 0.9%, and 4.3% on Cora, Cora-ML, Citeseer, and Polblogs, compared to the second-place method Meta-Self.

Averaged over experiments, the attack effect of our method is 1.275% higher than that of Meta-Self. Our proposed AtkSE outperforms the second-ranked model by more than 1% in four experiments and 2% in two experiments. Experiments prove that our attack model, AtkSE, is more effective than other methods. It proves that errors exist in the previous methods and that reducing these errors can improve the effectiveness of the attack model.

4.3 Ablation Study

To verify the effectiveness of each error reduction module, we conduct the ablation study. We ablate the three error reduction modules in Sections 3.1 (**H**ierarchical candidate selection and edge discrete sampling), 3.2 (**M**omentum gradient ensemble), and 3.3 (**S**emantic invariance), respectively. The ablated attack models are denoted as AtkSE-H, AtkSE-M, and AtkSE-S. Table 2 shows the results of the ablation study.

In the ablation experiments, the worse the attack of the ablated model is, the more critical the ablated module is. AtkSE-M has the best accuracy on Cora, Citeseer/GraphSage, and Polblogs/GraphSage. It ranks second on Cora-ML/GraphSage and Citeseer/GCN and worse than the other scenarios. AtkSE-H has the highest accuracy when the dataset is Cora-ML, Citeseer and Polblogs, and the victim model is GCN. AtkSE-H ranks second on four experiments

Table 1: Experimental results comparing our method with other methods. The group 'Clean' denotes the performance for unperturbed graph. The victim models are GraphSages and GCNs. The results are shown in the classification accuracy (%) under perturbation rate 5% on Cora, Cora-ML, Citeseer and Polblogs. The best results from experiments are bold.

	Cora		Cora-ML		Citeseer		Polblogs	
Victim	GraphSage	GCN	GraphSage	GCN	GraphSage	GCN	GraphSage	GCN
Clean	80.8±0.4	81.7±0.3	81.9±0.6	84.0±0.4	69.8±0.5	69.9±0.4	90.5±4.4	95.4±0.4
Random	80.7±0.6	81.2±0.3	80.6±0.8	82.8±0.4	68.7±0.7	67.8±0.4	78.1±1.1	84.0±0.9
DICE	80.0±0.4	80.9±0.5	80.7±0.9	82.5±0.4	69.1±0.8	69.0±0.4	76.1±1.9	79.2±0.7
EpoAtk	78.9±0.8	77.0±0.6	80.2±0.7	81.3±0.4	68.8±0.5	66.3±0.4	69.4±1.8	82.6±0.3
Meta-Train	76.1±1.0	76.4±0.2	77.7±1.8	79.0±0.3	67.9±0.9	65.5±0.4	71.2±4.0	83.5±0.7
Meta-Self	74.9±0.8	75.8±0.4	76.3±1.6	76.2±0.3	60.8±0.6	60.4±0.4	**64.1±2.6**	75.8±0.5
AtkSE(ours)	**73.3±0.6**	**73.7±0.4**	**75.4±1.0**	**74.0±1.0**	**60.7±0.4**	**59.5±0.5**	65.6±0.8	**71.5±0.3**

Table 2: Ablation study of our proposed error reduction methods, where AtkSE-H ablate the hierarchical candidate selection as well as discrete sampling in Sec 3.1, AtkSE-M ablates the gradient ensemble in Sec 3.2, and AtkSE-S ablates the semantic invariance in Sec 3.3. The highest accuracy in the ablation is highlighted by underlining, representing the module that is ablated as more important in that experiment.

	Cora		Cora-ML		Citeseer		Polblogs	
Victim	GraphSage	GCN	GraphSage	GCN	GraphSage	GCN	GraphSage	GCN
AtkSE	73.3±0.6	73.7±0.4	75.4±1.0	74.0±1.0	60.7±0.4	59.5±0.5	65.6±0.8	71.5±0.3
AtkSE-H	73.9±0.6	74.6±0.3	75.2±1.1	<u>75.3±0.8</u>	61.0±0.6	<u>60.4±0.5</u>	67.6±0.9	<u>72.4±0.3</u>
AtkSE-M	<u>74.0±0.7</u>	<u>75.2±0.2</u>	75.5±1.0	74.7±0.7	<u>61.6±0.5</u>	60.3±0.5	<u>68.0±1.3</u>	71.8±0.1
AtkSE-S	73.8±0.6	74.5±0.4	<u>75.6±1.1</u>	75.1±0.7	60.6±0.8	60.0±0.4	66.1±2.3	72.0±0.3

and worse on Cora-ML/GraphSage. AtkSE-S ranks first on Cora-ML/GraphSage, while it ranks lower on five experiments. AtkSE-M has the highest accuracy on four experiments, while AtkSE-H and AtkSE-S have three and merely one, respectively.

Overall, the momentum gradient ensemble is the module that enhances the attack model the most. The results indicate that the modules are ranked from highest to lowest importance as the momentum gradient ensemble, the hierarchical candidate selection and edge discrete sampling, and the semantic invariance. By comparing ablated models with AtkSE, the experiments prove the effectiveness of each module in AtkSE.

4.4 Gradient between the Edge-flip Interval

In this paper, graph edge perturbation is considered a problem of modeling a continuous distribution of gradients on edges. A possible challenge for our approach is whether the problem is worthy of being solved as a continuous problem over the edge-flip interval. To answer this question, we give examples of continuous gradient distributions of edge candidates in Fig.4. We can see that the gradients of the edges are continuous and smooth on the interval. We use the blue slash to indicate cases where the estimate is above the actual distribution and the red slash to indicate cases where the estimate is below the actual distribution. Referring to the error analysis in Section 3.1, in Fig.4, the previous methods adopt a gradient value higher than the integral of the continuous distribution on the (a,b,c) plot and lower than the integral on the (d) plot. Our approach
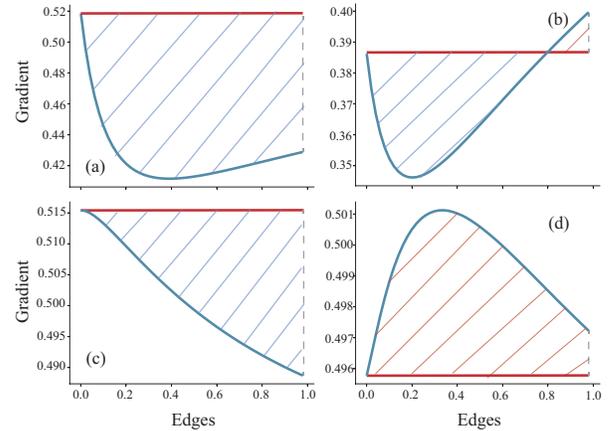
Figure 4: Illustrations of several variations of the structural gradient between the edge-flip interval.

considers the variation of gradients and transforms the edge perturbation from a discrete problem to a continuous gradients modeling problem. The error observed in Fig.4 proves the necessity for modeling continuous structural gradients. It also demonstrates why our approach improves the effectiveness of the attack.

5 RELATED WORKS

5.1 Graph Adversarial Attack

Graph adversarial attacks aim to disrupt the performance of graph neural networks using imperceptible attacks. There are three ways to attack graph-structured data: modifying the node features, modifying the graph structure, and injecting nodes [2, 7]. Among the researches in this field, more studies focus on modifying the graph structure [3, 20], and node injection [16], due to the specificity of graph data compared to general data. Among the general imperceptibility measures, attacks on graph structure are constrained by the L0 norm [27], and node injections are constrained by the number of manual nodes and their degrees [16]. Depending on the attacker's knowledge, attacks are classified as white-box attacks, gray-box attacks, and black-box attacks. White-box attacks [21, 23] open all the information of victim models. Gray-box [9, 10] attacks open labels of training set. Black-box attacks [11, 22] allow the attackers to query the predictions of the victim models. If the victim model is retrained, the attack is referred to as a poisoning attack, otherwise, it is referred to as an evasion attack. This paper studies the gray-box poisoning attacks. We aim to transfer the attack from the surrogate model to an unknown victim model, also referred to as the study of attack transferability.

5.2 Graph Edge Perturbation

This paper investigates attackers which globally perturb the graph structure. Mainstream approaches are based on the gradient derived from the loss function by backpropagation on the graph structure (or the adjacent matrix). Metattack [27] is the first gradient-based edge perturbation work on graph networks. It is a global poison attack model with a gray-box setting and the basis of other gradient-based edge perturbation methods. The attack strategy of Metattack is to search for the edge with the most significant gradient by the gradient on the adjacency matrix for modification, modifying one edge per iteration until an upper limit of the budget is reached. Another work [23] makes use of PGD [12] to search for an optimal perturbation matrix $S \in \{0,1\}^n$. EpoAtk [9] tries to add exploration to the attack model. EpoAtk has a small probability of not directly perturbing the edge with the most significant gradient, but instead of generating a perturbation from many candidate edges based on the loss values. These three works ignore whether the gradients they use are reliable. They only use the gradient derived from the current state of the graph and do not take into account the variation of the gradients on edges over the perturbation interval from 0 to 1. A method of integrated gradient on edges [21] is proposed to solve this problem. It requires traversing the adjacency matrix to score all edges and multiple gradient calculations to score each edge. It needs to traverse the adjacency matrix to score all edges, which has an exploded computational complexity $O(N^2)$. Moreover, it requires a massive amount of computing resources to perform a single time gradient calculation on all edges, limiting it to be used only in evasion attacks of small graphs, not any of the other scenarios. In addition, these efforts do not consider the gradient instability during the model optimization process. The above problems are addressed in our proposed method.

6 CONCLUSION

This paper proposes that the gradient on the graph structure in graph adversarial attack is subject to errors. This paper aims to illustrate these errors and propose corresponding modules to reduce them and implement them into our proposed attack model AtkSE. This paper first analyzes the error caused by treating graph structure attacks as a discrete problem with respect to gradients. We propose edge discrete sampling to reduce this error by approximating the continuous distribution of gradients over the edge flipping interval [0,1]. To ensure the computational cost of this step, we propose hierarchical candidate sampling based on the gradient saliency of the graph structure to select a small bag of edges as the candidates to be attacked rather than the entire N^2 space. Subsequently, we discuss the errors present in the gradient of the graph structure, including GNNs' unrobustness on the semantic space of the graph and the instability of the GNNs' representations due to the randomness of the parameter initialization. We propose semantic invariance and momentum gradient ensemble to solve these two errors, respectively. We integrate the above error reduction modules and propose the attack model AtkSE. In the experiments, we validate the effectiveness of our proposed method by comparing it with state-of-the-art baselines and showing the ablation study in untargeted poisoning gray-box attacks. The results demonstrate that our approach improves the attacker's performance, proving the reliability of our discussion on the error analysis and the effectiveness of our approach.

REFERENCES

[1] Lada A Adamic and Natalie Glance. 2005. The political blogosphere and the 2004 US election: divided they blog. In *Proceedings of the 3rd international workshop on Link discovery*. 36–43.

[2] Liang Chen, Jintang Li, Jiaying Peng, Tao Xie, Zengxu Cao, Kun Xu, Xiangnan He, and Zibin Zheng. 2020. A survey of adversarial learning on graphs. *arXiv preprint arXiv:2003.05730* (2020).

[3] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial attack on graph structured data. In *International conference on machine learning*. PMLR, 1115–1124.

[4] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. 2015. EXPLAINING AND HARNESSING ADVERSARIAL EXAMPLES. *stat* 1050 (2015), 20.

[5] Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. 2020. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering* (2020).

[6] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 1025–1035.

[7] Wei Jin, Yaxing Li, Han Xu, Yiqi Wang, Shuiwang Ji, Charu Aggarwal, and Jiliang Tang. 2021. Adversarial Attacks and Defenses on Graphs. *SIGKDD Explor. Newsl.* (2021), 19–34.

[8] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*.

[9] Xixun Lin, Chuan Zhou, Hong Yang, Jia Wu, Haibo Wang, Yanan Cao, and Bin Wang. 2020. Exploratory Adversarial Attacks on Graph Neural Networks. In *2020 IEEE International Conference on Data Mining (ICDM)*. IEEE, 1136–1141.

[10] Zihan Liu, Yun Luo, Zelin Zang, and Stan Z Li. 2022. Surrogate Representation Learning with Isometric Mapping for Gray-box Graph Adversarial Attacks. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*. 591–598.

[11] Yao Ma, Suhang Wang, Tyler Derr, Lingfei Wu, and Jiliang Tang. 2019. Attacking graph convolutional networks via rewiring. *arXiv preprint arXiv:1906.03750* (2019).

[12] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.

[13] Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. 2000. Automating the construction of internet portals with machine learning. *Information Retrieval* 3, 2 (2000), 127–163.

[14] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective classification in network data. *AI magazine* 29, 3 (2008), 93–93.

[15] Lichao Sun, Yingtong Dou, Carl Yang, Ji Wang, Philip S. Yu, Lifang He, and Bo Li. 2018. Adversarial Attack and Defense on Graph Data: A Survey. *arXiv preprint arXiv:1812.10528* (2018).

[16] Yiwei Sun, Suhang Wang, Xianfeng Tang, Tsung-Yu Hsieh, and Vasant Honavar. 2020. Adversarial attacks on graph neural networks via node injections: A hierarchical reinforcement learning approach. In *Proceedings of The Web Conference 2020*. 673–683.

[17] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*. PMLR, 3319–3328.

[18] Hao Wang, Tong Xu, Qi Liu, Defu Lian, Enhong Chen, Dongfang Du, Han Wu, and Wen Su. 2019. MCNE: an end-to-end framework for learning multiple conditional network representations of social network. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1064–1072.

[19] Xiaoyang Wang, Yao Ma, Yiqi Wang, Wei Jin, Xin Wang, Jiliang Tang, Caiyan Jia, and Jian Yu. 2020. Traffic flow prediction via spatial temporal graph neural network. In *Proceedings of The Web Conference 2020*. 1082–1092.

[20] Marcin Waniek, Tomasz P Michalak, Michael J Wooldridge, and Talal Rahwan. 2018. Hiding individuals and communities in a social network. *Nature Human Behaviour* 2, 2 (2018), 139–147.

[21] Huijun Wu, Chen Wang, Yuriy Tyshetskiy, Andrew Docherty, Kai Lu, and Liming Zhu. 2019. Adversarial examples on graph data: Deep insights into attack and defense. In *International Joint Conference on Artificial Intelligence*.

[22] Han Xu, Yao Ma, Hao-Chen Liu, Debayan Deb, Hui Liu, Ji-Liang Tang, and Anil K Jain. 2020. Adversarial attacks and defenses in images, graphs and text: A review. *International Journal of Automation and Computing* 17, 2 (2020), 151–178.

[23] Kaidi Xu, Hongge Chen, Sijia Liu, Pin Yu Chen, Tsui Wei Weng, Mingyi Hong, and Xue Lin. 2019. Topology attack and defense for graph neural networks: An optimization perspective. In *28th International Joint Conference on Artificial Intelligence, IJCAI 2019*. 3961–3967.

[24] Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. 2019. Adversarial examples: Attacks and defenses for deep learning. *IEEE transactions on neural networks and learning systems* 30, 9 (2019), 2805–2824.

[25] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. Graph neural networks: A review of methods and applications. *AI Open* 1 (2020), 57–81.

[26] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial attacks on neural networks for graph data. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2847–2856.

[27] Daniel Zügner and Stephan Günnemann. 2019. Adversarial Attacks on Graph Neural Networks via Meta Learning. In *International Conference on Learning Representations*.