

Practicable Black-Box Evasion Attacks on Link Prediction in Dynamic Graphs —a Graph Sequential Embedding Method

Jiate Li^{1,2}, Meng Pang^{1*}, Binghui Wang^{2*}

¹Nanchang University, Nanchang, China

²Illinois Institute of Technology, Chicago, USA

jetrichardlee@gmail.com, pangmeng1992@gmail.com, bwang70@iit.edu

Abstract

Link prediction in dynamic graphs (LPDG) has been widely applied to real-world applications such as website recommendation, traffic flow prediction, organizational studies, etc. These models are usually kept local and secure, with only the interactive interface restrictively available to the public. Thus, the problem of the black-box evasion attack on the LPDG model, where model interactions and data perturbations are restricted, seems to be essential and meaningful in practice. In this paper, we propose the first practicable black-box evasion attack method that achieves effective attacks against the target LPDG model, within a limited amount of interactions and perturbations. To perform effective attacks under limited perturbations, we develop a graph sequential embedding model to find the desired state embedding of the dynamic graph sequences, under a deep reinforcement learning framework. To overcome the scarcity of interactions, we design a multi-environment training pipeline and train our agent for multiple instances, by sharing an aggregate interaction buffer. Finally, we evaluate our attack against three advanced LPDG models on three real-world graph datasets of different scales and compare its performance with related methods under the interaction and perturbation constraints. Experimental results show that our attack is both effective and practicable.

Code — <https://github.com/JetRichardLee/GSE-METP>

Introduction

Dynamic graph sequences are a crucial structure for capturing the temporal dynamics of real-world systems. They represent the timing information within a dynamic graph, such as the historical evolution of interactions in a system. Link Prediction in Dynamic Graphs (LPDG), illustrated in Figure 1, leverages a dynamic graph sequence as input to predict future links for the next timestamp. Accurate predictions are highly valuable for a wide range of applications, including online recommendations, traffic management, and modeling disease contagion. In recent years, numerous methods for LPDG have been developed, including DynGEM (Goyal et al. 2018), MTCP (Zhou et al. 2018), EvolveGCN (Pareja et al. 2019), DyGCN (Manessi, Rozza, and Manzo 2020),

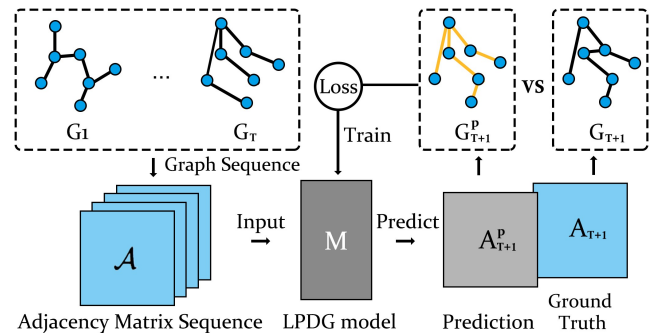


Figure 1: Illustration of LPDG: a dynamic graph sequence is taken as an input, and the LPDG model is trained to predict the future graph in the next time slice.

and MetaDyGNN (Yang et al. 2022a). These methods differ in their prediction settings. For example, MTCP and DyGCN focus on predicting the entire future graph at the next timestamp, whereas EvolveGCN and DynGEM provide predictions for the next set of changing edges.

Like adversarial attacks on static graphs, recent work (Fan et al. 2021) shows that LPDG for dynamic graphs is also vulnerable to *black-box* evasion attack, a practical attack where the attacker has no access to parameters or structure of the target model. However, the existing attack (called SAC) (Fan et al. 2021) based on reinforcement learning (RL) requires millions of interactions with the target model. Despite that, due to a coarse embedding method on the graph sequence, SAC is unable to attack large-scale graphs with large state spaces. These limitations inhibit SAC from being a practicable attack method in real-world scenarios.

We propose the *first practicable* black-box evasion attack on LPDG methods. Given a target model and instance, we generate the perturbed instance with only a few interactions with the model, while achieving promising attack results. Our attack method mainly consists of two novel designs:

- 1) **A graph sequential embedding (GSE).** GSE has a static degree embedding for every graph in the sequence, which enables it to find the desired state representations of the embedding sequence using two Long Short-term Memory (LSTM) (Hochreiter and Schmidhuber 1997) networks. Later, these representations can be applied to per-

*Corresponding authors

form the black-box attack based on a deep deterministic policy gradient (DDPG) (Silver et al. 2014) framework.

- 2) **A multi-environment training pipeline (METP)**. Considering that our model does not have sufficient resources for RL due to the limit on the interactions per instance, we hypothesize that there are similarities between instances from the same dataset and the same target model, and that training experience can be shared between them. Based on this hypothesis, we design a multi-environment training pipeline that sets multiple target instances as environments and trains a single model to attack them all.

Via testing on three recent LPDG models, Dynamic Graph Convolutional Network (DyGCN) (Manessi, Rozza, and Manzo 2020), Attention Based Spatial-Temporal Graph Convolutional Network (ASTGCN) (Guo et al. 2019) and Hyperbolic Temporal Graph Network (HTGN) (Yang et al. 2022b), over three real-world datasets, our method proves to be both reliable and effective under the practicable constraints and achieves state-of-the-art performance. Ablation experiments demonstrate the rationality of our graph sequential embedding design and multi-environment design.

Our main contributions are summarized as follows:

- We propose the first practicable black-box evasion attack against LPDG, which learns effective attacks by only a few amount of interactions with the target model.
- We develop a GSE method that assists to find reliable state embeddings for the RL framework to learn effective attacks. Furthermore, we design a METP to overcome the constraint of target model interactions by sharing experience between multiple instances.
- Evaluation results demonstrate the superiority of our practicable black-box attacks. Ablation experiments prove the rationality and effectiveness of our designs.

Related Work

Adversarial Attacks on Static Graphs

Existing attacks on graph learning for static graphs are classified as poisoning attacks and evasion attacks. Poisoning attack (Dai et al. 2018; Zügner and Günnemann 2019; Xu et al. 2019; Takahashi 2019; Liu et al. 2019; Sun et al. 2020; Zhang et al. 2021; Wang, Pang, and Dong 2023; Yang et al. 2024) is performed in the training phase (and testing phase). In training-time poisoning attacks, given a graph learning algorithm and a graph, an attacker carefully perturbs the graph (e.g., inject new edges to or remove the existing edges from the graph, perturb the node features) in the training phase, such that the learnt model misclassifies as many nodes, links, or graphs as possible on the clean testing graph(s). A special type of poisoning attack, called backdoor attack (Zhang et al. 2021), also perturbs testing graphs in the testing phase. Further, Yang et al. (2024) generalize the backdoor attack on graphs in the federated learning setting (Wang et al. 2022).

Evasion attack (Dai et al. 2018; Xu et al. 2019; Wang and Gong 2019; Wu et al. 2019; Ma et al. 2019; Ma, Ding, and Mei 2020; Li et al. 2021; Mu et al. 2021; Wang, Li, and Zhou 2022; Wang, Pang, and Dong 2023; Wang et al. 2024) is performed in the inference phase. In these attacks, given

a graph learning model and clean graph(s), an attacker carefully perturbs the graph structure or/and node features such that as many nodes or graphs as possible are misclassified on the perturbed graph(s) by the given model.

Deep Deterministic Policy Gradient

RL approaches with the actor-critic structure have been proposed and received great attention. In Haarnoja et al. (2018), the soft actor-critic (SAC) is described, with an additional reward on the entropy of the policy distribution. In Schulman et al. (2017), it proposes the proximal policy optimization (PPO), based on the trust region policy optimization (TRPO) (Schulman et al. 2015), with the clipped surrogate objective added to the loss function as refinement. In (Silver et al. 2014), the deterministic policy gradient (DPG), whose policy gives a deterministic action instead of a possibility distribution, is proposed. And in (Lillicrap et al. 2015), deep deterministic policy gradient (DDPG) is proposed with the extension of deep reinforcement learning to DPG.

Practicable Black-box Evasion Attack

In this section, we first define the practicable black-box evasion attack problem, and then introduce our two main designs, the graph sequential embedding (GSE) and the multi-environment training pipeline (METP). GSE learns the hidden sequential features of the dynamic graph sequence, and gives the embedding states, which are used by the policy network and the Q network of a DDPG agent. The agent’s policy network and Q network are trained under the METP and propagate gradients back to train the GSE models.

Problem Definition

LPDG Model Simulating a real-world attacking scenario, the LPDG model is private on its structure and parameters, but public on prediction interactions in limited chances. We only have the interface for inputting a dynamic graph sequence and receiving the prediction.

Clean input data A finite sequence of undirected graphs $\mathcal{G} = \{G_1, G_2, \dots, G_T\}$ is given as the input data, where $G_t = (V_t, E_t), \forall t \in [1 : T]$ denotes a snapshot graph at the time slice t . V_t is the set of nodes and E_t is the set of edges. In our attack setting, we could perform limited perturbations on the edge data $\mathcal{E} = \{E_t, \forall t \in [1 : T]\}$.

Output and metric The target LPDG model M is a black-box function. It predicts the edges of the graph at the next time slice $G_{T+1}^p = M(\mathcal{G}) = (V_{T+1}^p, E_{T+1}^p)$ and evaluates the performance by comparing E_{T+1}^p with the ground truth $E_{T+1} \in G_{T+1}$. We use the F1 score as a metric to evaluate the prediction and expect the attack to fool the model by having as low F1 as possible.

$$\begin{aligned}
 Precision &= \frac{|E_{T+1} \cap E_{T+1}^p|}{|E_{T+1}^p|} \\
 Recall &= \frac{|E_{T+1} \cap E_{T+1}^p|}{|E_{T+1} \cap E_{T+1}^p| + |E_{T+1} \cap \bar{C}E_{T+1}^p|} \\
 F_1 &= \frac{2 * Precision * Recall}{Precision + Recall}
 \end{aligned} \tag{1}$$

Attack Under Reinforcement Learning Since we have no knowledge of the LPDG model but only chances to interact with it, it requires us to learn from experience generated by limited exploration. Therefore, we apply a reinforcement learning method as our basic framework to perform the attack. We define our attack problem as follows.

State The perturbed dynamic graph sequence $\hat{\mathcal{G}} = \{\hat{G}_1, \hat{G}_2, \dots, \hat{G}_T\}$ describes the ground information during the attack. Since we only care about the $\hat{E}_{T+1}^p \in \hat{G}_{T+1}^p$ in the link prediction result, we assume that the nodes in the graph sequence are fixed as $\hat{V}_t = V, t \in [1, T+1]$ and basically, use an adjacency matrix sequence $\hat{\mathcal{A}} = \{\hat{A}_1, \hat{A}_2, \dots, \hat{A}_T\}$ to represent the state of the dynamic graph sequence. This raw representation has a large state space and lacks the essential feature extraction, making the policy and the Q difficult to converge. Therefore, in this paper we will introduce a new way to represent the state.

Action The action we take in the RL framework is the perturbation we apply to the dynamic graph sequence. Each time we perform an action a , we add one edge $(u, v), u, v \in V$ and delete one edge $(u', v'), u', v' \in V$ for all the $\hat{A}_t \in \hat{\mathcal{A}}$ in the sequence and obtain a new matrix sequence $\hat{\mathcal{A}}'$:

$$\hat{A}'_t = \hat{A}_t([u][v] = 1, [u'][v'] = 0), t \in [1, T] \quad (2)$$

Environment and reward In our black-box attack on LPDG, the environment is the LPDG model M and the original dynamic graph sequence \mathcal{G} . At the attack step k , we perform an attack action a_k to the edge sequence \hat{A}_k and get the next attacked sequence \hat{A}_{k+1} . We feed it into the model M and get the output $\hat{E}_{T+1, k+1}^p$. The F1 score metric f_{k+1} is computed by comparing it with the original $E_{T+1} \in G_{T+1}$. Depending on how much the prediction performance decreases from the previous prediction metric f_k , we give the reward for the action a_k and the state $S_k = \hat{A}_k$:

$$r(S_k, a_k) = f^{k+1} - f^k \quad (3)$$

Practicable constraints There are two constraints for the agent to satisfy the practicability requirement. First, for each attack instance M and \mathcal{G} , it is only able to perturb a few amount of edges. We set the δ as the ratio limit and n as the amount limit. For an attack attempt, we could only have K perturbations in total, which satisfies the constraint:

$$K = \min(\delta|E_{max}|, n) \quad (4)$$

$|E_{max}| = |V|^2/2$ is the maximum number of edges on the graph and n is a relatively small constant. In the second constraint, the agent only has limited chances to interact with the target model M . To learn attacks for one instance, we could only query M within I times during the training.

Graph Sequential Embedding

As mentioned before, the state space of the raw matrix sequence $\hat{\mathcal{A}}$ is too large, and makes the deep reinforcement learning difficult to converge. Therefore, we develop a graph sequential embedding (GSE) method, to embed it into a

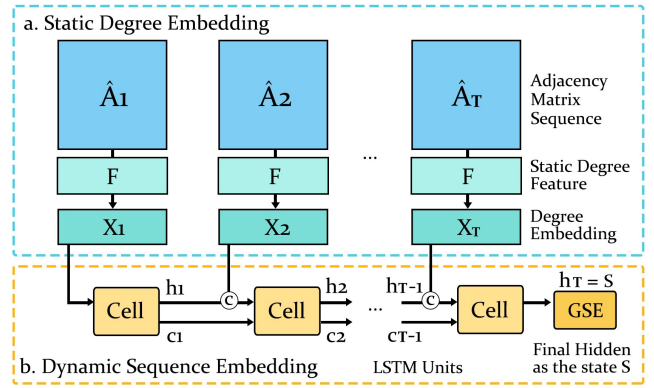


Figure 2: Illustration of the proposed GSE method.

smaller state space. Our GSE method consists of two modules, a static degree embedding module and a dynamic sequential embedding module.

Static Degree Feature In the static degree embedding, we first compute a degree feature to embed each graph matrix in the sequence first. In step a, we multiply the average adjacency matrix $\tilde{A} = \text{average}(\mathcal{A})$ of the original dynamic graph sequence, with an all-one tensor in the shape of $|V| \times 1$, which gives the average degree of each node in the original sequence. We multiply this again by the average adjacency matrix and get the amount of nodes connected within two steps. This doesn't work exactly for the three steps and the four steps, but we use the same method to approximate them. We describe this as:

$$d = \{d_0, \tilde{A}d_0, \tilde{A}^2d_0, \tilde{A}^3d_0\}, d_0 = [1]_{|V| \times 1} \quad (5)$$

In the step b, we normalize each of the last three separately with the batchnorm function bn . We then additionally inject a $|V| \times 4$ random noise feature by concatenation to enhance the feature's expressiveness (Sato, Yamada, and Kashima 2021). Our final degree feature F is represented by:

$$F = [[\text{Random}]_{|V| \times 4}, d_0, bn(d_1), bn(d_2), bn(d_3)] \quad (6)$$

This feature is calculated at the start of the attack for each instance, and it remains static throughout the learning.

Dynamic Sequence Embedding After we compute the static degree feature F , we multiply it with each matrix $\hat{A}_t, t \in [1, T]$ to get the degree embedding sequence $\mathcal{X} = \{X_1, X_2, \dots, X_T\} \in R^{T \times |V| \times |S|}$:

$$\mathcal{X} = \hat{\mathcal{A}}F \quad (7)$$

This process is described as step a in Figure 2. Then in step b, we use an LSTM module, to dynamically find the efficient embedding state S for the graph sequence through training. For simplicity, we denote this process as SE :

$$\begin{aligned} h_0 &= X_1, \quad c_0 = 0, \\ L_i(-) &= W_{i,-}X_t + b_{i,-}, \quad L_h(-) = W_{h,-}h_{t-1} + b_{h,-}, \\ i_t &= \sigma(L_i(i) + L_h(i)), \quad f_t = \sigma(L_i(f) + L_h(f)), \\ g_t &= \tanh(L_i(g) + L_h(g)), \quad o_t = \sigma(L_i(o) + L_h(o)), \\ c_t &= f_t c_{t-1} + i_t g_t, \quad h_t = o_t \tanh(c_t), \quad \forall t \in [1, T] \\ S &= SE(\mathcal{X}) = h_T \end{aligned} \quad (8)$$

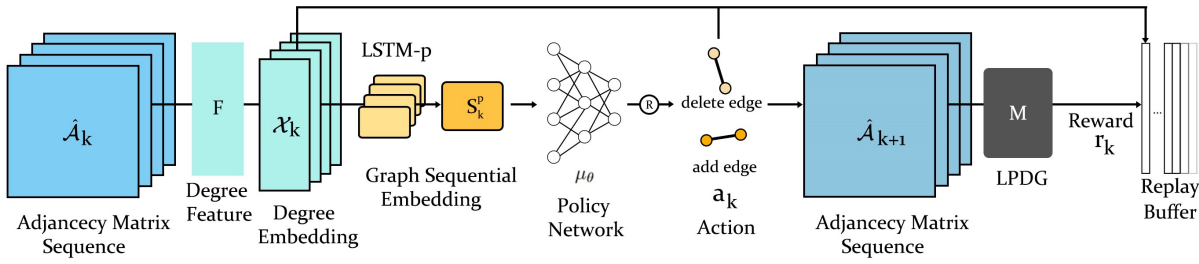


Figure 3: An overview of the agent interaction.

In our attack design, we actually use two different GSE models to train embedding states for the actor and the critic. They share the same degree embedding \mathcal{X} , but separate LSTM modules and result states. We denote the LSTM modules as LSTM-p and LSTM-q, the embedding processes as SE_p and SE_q , the embedding states as S^p and S^q for the policy network and the Q network respectively.

Multi-Environment Training

Our multi-environment training is designed based on a DDPG method. In traditional reinforcement learning, an agent is trained in one or more but identical environments. However, in a practicable attack, we have a limit on the amount of interactions for each attack instance. Hence, we train one agent under several separate instances instead. We attack them alternately and learn our agent from the collective experience they share. These instances are under the same dataset and the same target model, and we believe that they have similarities in the state space and help to explore more efficiently.

There are three main tasks in our multi-environment training pipeline, (1) perform attack interactions with each instance against the target model; (2) train the Q network and the Q GSE module; (3) train the policy network and the policy GSE module. In each attack step, we perform the three tasks sequentially.

Attack Interaction The Figure 3 describes how our agent performs the attack on each instance. In each attack step $k \in [1, K]$, we have the adjacency matrix sequence \hat{A}_k generated from the $k - 1$ step, and the static degree feature F computed in advance. In step a, we apply our graph sequential embedding method, through the degree feature F and the LSTM-p module, and obtain the degree embedding \mathcal{X}_k and the graph sequential embedding S_k^p . In step b, we design our policy network μ_θ based on an MLP network, denoted MLP_θ , to select our add and delete action:

$$\mu_\theta(S_k^p) = \text{Sigmod}(MLP_\theta(S_k^p)) * |V| \quad (9)$$

However, there are two situations where we will use a random action instead. First, the interaction records in the shared buffer are smaller than the batch size, so we need to fill it with random attacks first and allow our agent to start training in the following sections. Otherwise, the buffer would be filled with similar actions generated by the untrained policy, leading the agent to a meaningless training

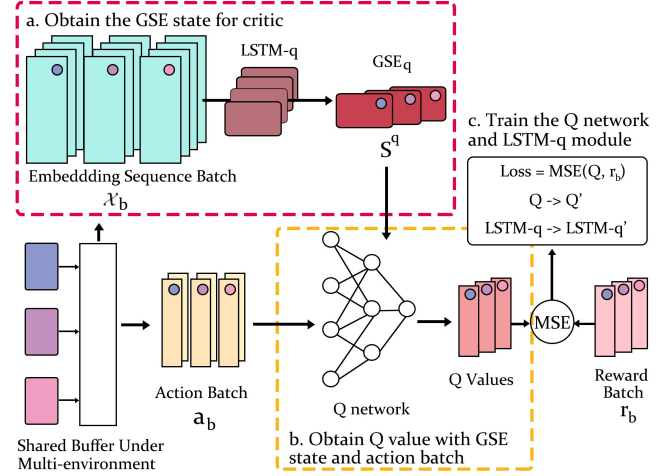


Figure 4: Training pipeline for the Q network and Q GSE.

result. The second case is that we hope to explore the state space during the attack, so when we go for a certain attack steps we will apply a random action instead. Finally our action a_k is selected from:

$$a_k = \begin{cases} \mu_\theta(S_k^p), & \text{attack} \\ a_{rand}, & \text{random} \end{cases} \quad (10)$$

After obtaining the action a_k , we apply it to the previous sequence \hat{A}_k , generate a new graph sequence \hat{A}_{k+1} , perform the prediction on the target LPDG model M and obtain the reward r_k in step c. We write the embedding sequence \mathcal{X}_k , the action a_k and the reward r_k to the shared replay buffer in step d.

Q Training The Figure 4 describes the process of training the Q network and the Q GSE model. The shared buffer has received experience from several instances, and we sample them collectively. The sample consists of a batch of embedding sequences \mathcal{X}_b , actions a_b and rewards $r_b = r(\mathcal{X}_b, a_b)$ from the shared buffer. In step a, we feed them to the LSTM-q to get their sequential embedding states S^q for the critic:

$$S^q = SE_q(\mathcal{X}_b) \quad (11)$$

In step b, we use an MLP model, denoted as MLP_q , as our Q network:

$$Q(S^q, a_b) = MLP_q([a_b, S^q]) \quad (12)$$

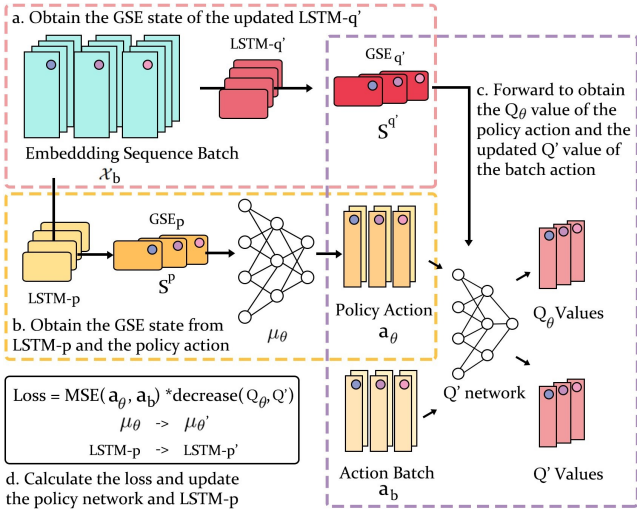


Figure 5: Training pipeline for the policy network and policy GSE.

In the deep reinforcement learning, the value function $Q(S, a)$ is usually defined as:

$$Q(S, a) = r(S, a) + \gamma Q(S', \mu_\theta(S')) \quad (13)$$

S' is the next state after S takes the action a . For the LPDG attack, however, the order of actions is not critical. If a set of actions $[a_1, a_2, \dots, a_K]$ leads the state from S_1 to S_{K+1} and the order is changed arbitrarily, the final state will still be S_{K+1} . Therefore, we concentrate on learning the Q-value for the reward itself:

$$Q(S, a) = r(S, a) \quad (14)$$

and in step c we use the MSE loss to define the Q loss for the Q network and the Q GSE to train:

$$loss_Q = MSE(Q(S^q, a_b), r_b) \quad (15)$$

Finally, we back propagate the Q loss to update the Q network and the Q GSE model with the learning rate τ :

$$Q' \leftarrow \tau Q + (1 - \tau)Q' \quad (16)$$

$$LSTM'_q \leftarrow \tau LSTM_q + (1 - \tau)LSTM'_q \quad (17)$$

Policy Training The Figure 5 describes the process of training the policy network and the policy GSE under the multi-environment. Similar to the Q training, we have the same batch of interaction samples from the shared buffer: embedding sequences \mathcal{X}_b and actions a_b . In step a, we forward embedding sequences to the updated LSTM-q' module to obtain the GSE states $S^{q'}$ of Q:

$$S^{q'} = SE_{q'}(\mathcal{X}_b) \quad (18)$$

In step b, we forward embedding sequences to the LSTM-p to obtain the GSE states S^p . Then we feed them to the policy network μ_θ to get the policy actions a_θ :

$$S^p = SE_p(\mathcal{X}_b), \quad a_\theta = \mu_\theta(S^p) \quad (19)$$

Name	Node	Average Edges	Graph type
Hagggle	274	12584	Human contact
Facebook	1000	97779	Social circle
AS	6474	141845	Traffic flows

Table 1: A brief description of datasets.

We define our policy loss as the scalar product of the differences between the actions and the decreases in the Q values. This means that if an action in a_θ is worse than that in the a_b , we hope to learn towards the batch one, otherwise we hope to stay the same. We update the policy network and the policy GSE with the learning rate τ as follows:

$$d(S^{q'}, a_b) = \max(0, Q'(S^{q'}, a_b) - Q_\theta(S^{q'}, a_\theta)) \quad (20)$$

$$loss_{\mu_\theta}(S^{q'}, a_b) = MSE(a_b, a_\theta) \cdot d(S^{q'}, a_b) \quad (21)$$

$$\mu'_\theta \leftarrow \tau \mu_\theta + (1 - \tau)\mu'_\theta \quad (22)$$

$$LSTM'_p \leftarrow \tau LSTM_p + (1 - \tau)LSTM'_p \quad (23)$$

Experiments

In this section, we test our attack method on three real-world datasets against three LPDG models, compared with four baseline black-box evasion methods. The experiments consist of two aspects, the performance evaluation and the interaction scale impact test for ablation study.

For each setting, it consists of the attack method \mathcal{C} , the dataset \mathcal{D} , the LPDG model M , the perturbation constraint $K = \min(\delta|E_{max}|, n)$ and the interaction constraint I . We first take 10 different instances from the dataset \mathcal{D} , each consisting of 11 temporal graphs. The first 10 graphs are the input to M and the last 1 is the ground truth for the prediction. We use these instances to train the target model M , and apply the attack method \mathcal{C} to perform a black-box attack. Each attack consists of several attempts. In each attempt, the agent gives at most K perturbations to the sequence, and the lowest prediction metric achieved during the perturbations is taken as the result of that attempt. The total number of interactions with the target model, taken by all attempts, could not exceed I . Once interactions are exhausted, we take the best result of all attempts as the performance of the setting.

Experiments Setting

Datasets We use three real-world datasets with varying scales, and their properties are shown in Table 1.

Hagggle This is a social network available at KONECT and published in (Kunegis 2013), representing the connection between users measured by wireless devices.

Facebook This is a subgraph of the ‘‘Social circles: Facebook’’ social networks from SNAP (Leskovec and Sosi c 2016). We randomly delete edges to generate the dynamic graph sequences.

Autonomous systems ‘‘AS-733’’ is a large traffic flow networks available on SNAP (Leskovec and Sosi c 2016). We randomly delete edges to generate the dynamic graph sequences as well.

Compared Attack

Random attack In this attack method, the agent randomly chooses two nodes to add an edge and two nodes to delete an edge as the action.

SAC attack This attack is introduced by (Fan et al. 2021), which claims to be the first black-box evasion attack on LPDG problem. However, without any constraint on interactions, we conclude that this method is impracticable, and prove this in our experiments.

SAC-METP attack This attack is an ablation version of our method. It applies the multi-environment training pipeline, but takes the same graph embedding method as (Fan et al. 2021).

GSE attack This attack is an ablation version of our method. It applies the graph sequential embedding method, but the models for different instances are trained and applied separately.

GSE-METP attack This attack is the complete version of our method. It applies both the graph sequential embedding method and the multi-environment training pipeline.

LPDG Model

DyGCN Dynamic Graph Convolutional Network (DyGCN) (Manessi, Rozza, and Manzo 2020), is an extension of GCN-based methods. It generalizes the embedding propagation scheme of GCN to the dynamic setting in an efficient manner, and propagates the change along the graph to update node embedding.

ASTGCN Attention Based Spatial-Temporal Graph Convolutional Networks (ASTGCN) (Guo et al. 2019), has several independent components. Each of them consists of two parts, the spatial-temporal attention mechanism and the spatial-temporal convolution. In our adaption for experiments, we use one component to consist a test ASTGCN.

HTGN Hyperbolic Temporal Graph Network (HTGN) (Yang et al. 2022b) is a temporal graph embedding method, which learns topological dependencies and implicitly hierarchical organization of each graph sequence individually and gives link predictions on it.

Attack Settings For the performance evaluation, we set the default attack rate limit to $\delta = 0.02$, the default attack amount limit to $n = 1000$, and the default interaction limit to $I = 5K$ to make the attack attempts complete in tests. Then, in the interaction impact test, we show our attack results in full on different $I \in [K, 10K]$.

Experiments Results

Effectiveness Evaluation Table 2 shows the performance of the four attack methods under the default setting against three LPDG models. As shown in the table, our GSE and GSE-METP methods perform best in these tests. Under constraints, the SAC method has good results on the small dataset, while in most cases it behaves close to or worse than the random method. This suggests that it is not a practicable

Method	Haggle	Facebook	AS
Edge Ratio(δ)	2%	0.2%	4.8e-5
Edge Amount(n)	751	1000	1000
Interaction(I)	3755	5000	5000
DyGCN (Original)	0.9930	0.9745	0.8990
Random	0.8979	0.9681	0.8862
SAC	0.8149	0.9665	0.8899
SAC-METP	0.8094	0.9651	0.8890
GSE(Ours)	0.8043	0.9629	0.8910
GSE-METP(Ours)	0.8118	0.9653	0.8573
ASTGCN (Original)	0.9852	0.9862	0.8303
Random	0.9752	0.9298	0.8276
SAC	0.9825	0.9853	0.8264
SAC-METP	0.9750	0.9859	0.8263
GSE(Ours)	0.9748	0.9407	0.8226
GSE-METP(Ours)	0.9702	0.9011	0.8020
HTGN (Original)	0.9753	0.9375	0.8665
Random	0.8083	0.8652	0.8377
SAC	0.9603	0.9145	0.8526
SAC-METP	0.9676	0.9141	0.8527
GSE(Ours)	0.8311	0.8930	0.7944
GSE-METP(Ours)	0.7417	0.8494	0.7649

Table 2: Performance evaluations on the default setting.

method. Our GSE-METP method behaves effectively on all experiment settings. On the largest dataset AS, GSE-METP has significantly better performance than others. This validates that GSE-METP is the first practicable black-box evasion attack against LPDG methods.

Ablation Study Figure 6 shows our method is more effective as the interaction limit grows compared with other methods. First, on the small Haggle, SAC and SAC-METP are effective when the interactions are enough for training, but fail to converge to good results on the large Facebook and AS. Instead, our GSE and GSE-METP converge to better attacks. This ablation study proves the effectiveness of our GSE design. Second, SAC-METP and GSE-METP converge faster than SAC and GSE, and also result in better performances. This ablation study shows the efficiency of our multi-environment training pipeline design.

Why SAC fails? We further tracked the states and actions during SAC’s attack and ours to explore the reason that SAC fails. We found that SAC’s represented state has a relatively small variance during the attack, meaning it nearly does not change, while ours changed apparently. For instance, on a sample from Haggle on DyGCN, we observed SAC’s add actions converge to the $\{0,0\}$ action, and the variance of its delete actions also decreases from $\{8032, 4890\}$ to $\{39, 84\}$. In contrast, GSE-METP has average actions on $\{97, 126\}$ for adding and $\{152, 130\}$ for deleting and exhibits greater variance $\{4655, 4708\}$ for adding and $\{4388, 7405\}$ for deleting. This suggests GSE-METP adapts its action nodes more responsively in different states, aligning with the ideal agent behavior. As for final performance, SAC only reduces accuracy to 0.854, while GSE-METP to 0.724.

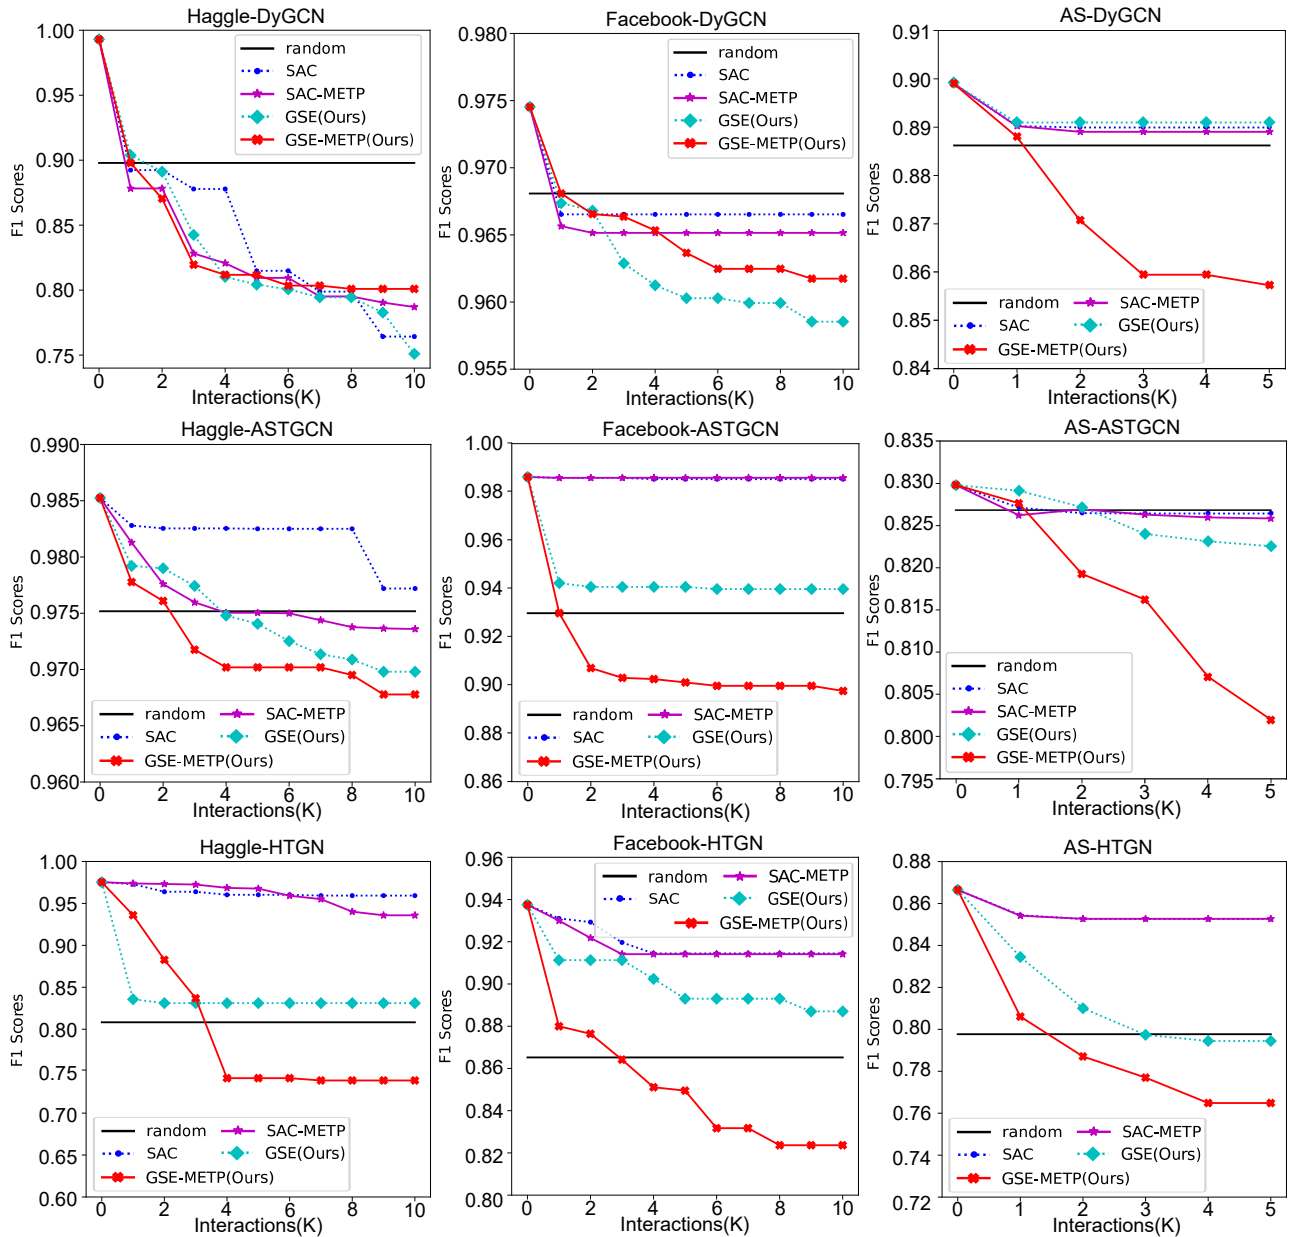


Figure 6: Interaction impacts on DyGCN, ASTGCN and HTGN.

Based on the observation above, we made an inference for SAC’s poor performance: SAC relies on a ranking of node degrees as the state, which is relatively static during learning as each action pair changes the degrees of at most four nodes. Consequently, the generated state representations of dynamic graphs stay nearly constant during the attack process. When the replay buffer is filled with repetitive states, each with a high negative reward, the policy network tends to produce extreme actions to avoid further negative rewards. This results in SAC repeatedly selecting delete-add actions on the same node pair— $\{0,0\}$ or $\{V, V\}$, which is apparently an undesired behavior.

Conclusion

We propose the first practicable black-box evasion attack against the link prediction in dynamic graph models. We design a graph sequential embedding method and a multi-environment training pipeline, and combine them with a deep reinforcement learning method, DDPG, to perform effective attacks under interaction and perturbation constraints. Experiments on three advanced LPDG methods demonstrate the effectiveness of our attack. Crucial future work is to design provably robust LPDG against the proposed evasion attacks, inspired by existing certified defense on static graphs (Wang et al. 2021; Xia et al. 2024).

Acknowledgments

Pang is supported in part by Natural Science Foundation of China (62466036), Natural Science Foundation of Jiangxi Province (20232BAB212025), High-level and Urgently Needed Overseas Talent Programs of Jiangxi Province (20232BCJ25024) and Key Research and Development Program of Jiangxi Province (20243BBG71035). Wang is supported in part by the Cisco Research Award and the National Science Foundation under grant Nos. ECCS-2216926, CCF-2331302, CNS-2241713 and CNS-2339686.

References

- Dai, H.; Li, H.; Tian, T.; and more. 2018. Adversarial attack on graph structured data. In *ICML*.
- Fan, H.; Wang, B.; Zhou, P.; and et al. 2021. Reinforcement learning-based black-box evasion attacks to link prediction in dynamic graphs. In *IEEE HPCC*.
- Goyal, P.; Kamra, N.; He, X.; and Liu, Y. 2018. DynGEM: Deep Embedding Method for Dynamic Graphs. *CoRR*.
- Guo, S.; Lin, Y.; Feng, N.; Song, C.; and Wan, H. 2019. Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting. *AAAI*, 33(01): 922–929.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *ICML*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Computation*, 9(8): 1735–1780.
- Kunegis, J. 2013. KONECT: The Koblenz Network Collection. In *WWW Companion*.
- Leskovec, J.; and Sosič, R. 2016. SNAP: A General-Purpose Network Analysis and Graph-Mining Library. *ACM TIST*.
- Li, J.; Xie, T.; Chen, L.; Xie, F.; He, X.; and Zheng, Z. 2021. Adversarial attack on large scale graph. *IEEE TKDE*.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2015. Continuous control with deep reinforcement learning. *arXiv*.
- Liu, X.; Si, S.; Zhu, X.; Li, Y.; and Hsieh, C.-J. 2019. A unified framework for data poisoning attack to graph-based semi-supervised learning. In *NeurIPS*.
- Ma, J.; Ding, S.; and Mei, Q. 2020. Towards More Practical Adversarial Attacks on Graph Neural Networks. In *NeurIPS*.
- Ma, Y.; Wang, S.; Derr, T.; Wu, L.; and Tang, J. 2019. Attacking graph convolutional networks via rewiring. *arXiv*.
- Manessi, F.; Rozza, A.; and Manzo, M. 2020. Dynamic graph convolutional networks. *Pattern Recognition*.
- Mu, J.; Wang, B.; Li, Q.; Sun, K.; Xu, M.; and Liu, Z. 2021. A Hard Label Black-box Adversarial Attack to Graph Neural Networks. In *CCS*.
- Pareja, A.; Domeniconi, G.; Chen, J.; Ma, T.; Suzumura, T.; Kanezashi, H.; Kaler, T.; and Leiserson, C. E. 2019. EvolveGCN: Evolving Graph Convolutional Networks for Dynamic Graphs. *CoRR*, abs/1902.10191.
- Sato, R.; Yamada, M.; and Kashima, H. 2021. Random features strengthen graph neural networks. In *SIAM SDM*.
- Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *ICML*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv*.
- Silver, D.; Lever, G.; Heess, N.; Degris, T.; Wierstra, D.; and Riedmiller, M. 2014. Deterministic Policy Gradient Algorithms. In *ICML*.
- Sun, Y.; Wang, S.; Tang, X.; Hsieh, T.-Y.; and Honavar, V. 2020. Adversarial Attacks on Graph Neural Networks via Node Injections: A Hierarchical Reinforcement Learning Approach. In *The Web Conference*.
- Takahashi, T. 2019. Indirect Adversarial Attacks via Poisoning Neighbors for Graph Convolutional Networks. In *BD*.
- Wang, B.; and Gong, N. 2019. Attacking Graph-based Classification via Manipulating the Graph Structure. In *CCS*.
- Wang, B.; Jia, J.; Cao, X.; and Gong, N. Z. 2021. Certified robustness of graph neural networks against adversarial structural perturbation. In *KDD*.
- Wang, B.; Li, A.; Pang, M.; Li, H.; and Chen, Y. 2022. Graphfl: A federated learning framework for semi-supervised node classification on graphs. In *ICDM*.
- Wang, B.; Li, Y.; and Zhou, P. 2022. Bandits for Black-box Attacks to Graph Neural Networks with Structure Perturbation. In *CVPR*.
- Wang, B.; Lin, M.; Zhou, T.; Zhou, P.; Li, A.; Pang, M.; Li, H.; and Chen, Y. 2024. Efficient, direct, and restricted black-box graph evasion attacks to any-layer graph neural networks via influence function. In *WSDM*.
- Wang, B.; Pang, M.; and Dong, Y. 2023. Turning strengths into weaknesses: A certified robustness inspired attack framework against graph neural networks. In *CVPR*.
- Wu, H.; Wang, C.; Tyshetskiy, Y.; Docherty, A.; Lu, K.; and Zhu, L. 2019. Adversarial examples on graph data: Deep insights into attack and defense. In *IJCAI*.
- Xia, Z.; Yang, H.; Wang, B.; Jia, J.; et al. 2024. GNN-Cert: Deterministic Certification of Graph Neural Networks against Adversarial Perturbations. In *ICLR*.
- Xu, K.; Chen, H.; Liu, S.; and more. 2019. Topology attack and defense for graph neural networks: An optimization perspective. In *IJCAI*.
- Yang, C.; Wang, C.; Lu, Y.; Gong, X.; Shi, C.; Wang, W.; and Zhang, X. 2022a. Few-Shot Link Prediction in Dynamic Networks. In *WSDM*.
- Yang, M.; Zhou, M.; Xiong, H.; and King, I. 2022b. Hyperbolic Temporal Network Embedding. *IEEE TKDE*, 1–14.
- Yang, Y.; Li, Q.; Jia, J.; Hong, Y.; and Wang, B. 2024. Distributed backdoor attacks on federated graph learning and certified defenses. In *CCS*.
- Zhang, Z.; Jia, J.; Wang, B.; and Gong, N. 2021. Backdoor attacks to graph neural networks. *SACMAT*.
- Zhou, L.; Yang, Y.; Ren, X.; Wu, F.; and Zhuang, Y. 2018. Dynamic Network Embedding by Modeling Triadic Closure Process. In *AAAI*.
- Zügner, D.; and Günnemann, S. 2019. Adversarial attacks on graph neural networks via meta learning. In *ICLR*.