

Spear and Shield: Adversarial Attacks and Defense Methods for Model-Based Link Prediction on Continuous-Time Dynamic Graphs

Dongjin Lee¹, Juho Lee², Kijung Shin^{1,2}

¹School of Electrical Engineering, Daejeon, South Korea

²Kim Jaechul Graduate School of AI, KAIST, Seoul, South Korea
{dongjin.lee, juholee, kijungs}@kaist.ac.kr

Abstract

Real-world graphs are dynamic, constantly evolving with new interactions, such as financial transactions in financial networks. Temporal Graph Neural Networks (TGNNs) have been developed to effectively capture the evolving patterns in dynamic graphs. While these models have demonstrated their superiority, being widely adopted in various important fields, their vulnerabilities against adversarial attacks remain largely unexplored. In this paper, we propose T-SPEAR, a simple and effective adversarial attack method for link prediction on continuous-time dynamic graphs, focusing on investigating the vulnerabilities of TGNNs. Specifically, before the training procedure of a victim model, which is a TGNN for link prediction, we inject edge perturbations to the data that are unnoticeable in terms of the four constraints we propose, and yet effective enough to cause malfunction of the victim model. Moreover, we propose a robust training approach T-SHIELD to mitigate the impact of adversarial attacks. By using edge filtering and enforcing temporal smoothness to node embeddings, we enhance the robustness of the victim model. Our experimental study shows that T-SPEAR significantly degrades the victim model's performance on link prediction tasks, and even more, our attacks are transferable to other TGNNs, which differ from the victim model assumed by the attacker. Moreover, we demonstrate that T-SHIELD effectively filters out adversarial edges and exhibits robustness against adversarial attacks, surpassing the link prediction performance of the naive TGNN by up to 11.2% under T-SPEAR. The code and datasets are available at <https://github.com/wooner49/T-spear-shield>.

1 Introduction

Graph Neural Networks (GNNs) have shown impressive capabilities across various tasks (Sun et al. 2022; Jin et al. 2021). However, their susceptibility to adversarial perturbations (Bojchevski and Günnemann 2019; Dai et al. 2018; Zügner, Akbarnejad, and Günnemann 2018; Xu et al. 2019) has raised concerns about their safe use in critical applications, including fraud detection (Hooi et al. 2016; Zeager et al. 2017) and malware detection (Hou et al. 2019). Despite substantial progress made in developing attack and defense techniques (Sun et al. 2022; Jin et al. 2021), the focus has primarily been on GNNs designed for static graphs.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

In the real world, graphs are dynamic, constantly changing over time. For example, in financial networks, transactions between accounts occur continuously over time, leading to ever-changing graphs. Dynamic graphs fall into two types: Discrete-Time Dynamic Graphs (DTDGs), denoted as a sequence of static graph snapshots, and Continuous-Time Dynamic Graphs (CTDGs), represented as a sequence of interactions over continuous time. Due to the capability of CTDGs to capture more intricate and fine-grained temporal patterns in dynamic graphs, CTDGs have recently garnered increasing attention, resulting in the development of numerous machine learning models on CTDGs for time-critical applications (Kazemi et al. 2020; Thomas et al. 2022).

To effectively capture the evolving patterns in CTDGs, Temporal Graph Neural Networks (TGNNs) (Rossi et al. 2020; Kumar, Zhang, and Leskovec 2019; Xu et al. 2020; Wang et al. 2021a; Zhou et al. 2022) have been developed. TGNNs jointly learn the temporal, structural, and contextual relationships present in CTDGs by encoding graph information into time-aware node embeddings. By incorporating temporal information, TGNNs have demonstrated superior performance over static GNNs in various tasks like link prediction (Cong et al. 2023; Wang et al. 2021b) and dynamic node classification (Rossi et al. 2020).

Nonetheless, in contrast to the actively researched adversarial attacks on static GNNs, the vulnerabilities of TGNNs to adversarial attacks remain underexplored, yet the significance of conducting such research is undeniable. For instance, in a financial attack (Fursov et al. 2021; Zeager et al. 2017) scenario, an adversary may inject adversarial transactions into the transaction graph, perturbing the timing and content of transactions to mislead the model running on the graph. This manipulation could lead to the model falsely predicting fraudulent transactions as legitimate, resulting in inaccurate risk predictions and potential financial losses.

In this context, we explore the weakness of TGNNs by introducing T-SPEAR, a simple and effective adversarial attack method for link prediction on CTDGs. From an adversarial perspective, designing effective attacks on CTDGs poses unique challenges compared to attacks on static graphs. First, in addition to the perturbation injected into graphs, attacks should carefully be timed. Second, attacks should stay stealthy and not deviate significantly from the usual evolving patterns of interactions in the original graph.

To address these challenges, we impose four constraints on perturbations to ensure the attack remains unnoticeable: in terms of the (C1) perturbation budget, (C2) distribution of time, (C3) endpoints of adversarial edges, and (C4) number of perturbations per node. While satisfying the above constraints, we generate edge perturbations by selecting edges that are unlikely to be formed when considering the evolution of dynamic graphs, and then inject them into the data before the training process of a victim model (i.e., a poisoning attack). Benchmarking on link prediction tasks, T-SPEAR outperforms baselines, all within the same perturbation budget. Moreover, T-SPEAR remains effective even when the victim model differs from the model assumed by the attacker, indicating that T-SPEAR is transferrable.

In terms of the defense against adversarial attacks, the techniques available for static GNNs (Mujkanovic et al. 2022) are not adequate, partly because they do not consider the dynamic nature of the graphs. For instance, in CTDGs, a normal edge may turn into an adversarial edge over time, but a defense algorithm for static GNNs cannot adapt to this change. To mitigate the impact of adversarial attacks on CTDGs, we propose T-SHIELD, a robust training method for TGNNs. By employing the edge filtering technique that considers graphs’ dynamics, T-SHIELD eliminates potential adversarial edges from the corrupted dynamic graphs without any prior knowledge of the attacks. In addition, we enforce temporal smoothness on node embeddings to prevent abrupt changes, further enhancing the models’ robustness.

In summary, our contributions are as follows:

- To the best of our knowledge, we are the first to formulate adversarial attacks on CTDGs under realistic constraints regarding unnoticeability.
- We propose T-SPEAR, a simple and effective poisoning attack method on CTDGs, and demonstrate its superiority over 5 baselines across 4 different datasets and 4 victim models on link prediction tasks.
- We propose T-SHIELD, a robust training method for TGNNs that shows robustness against adversarial attacks, surpassing the naïve TGNN by up to 11.2% on link prediction tasks.

2 Related Work

Temporal Graph Neural Networks (TGNNs). TGNNs are categorized into two types: discrete-time TGNNs and continuous-time TGNNs. Discrete-time TGNNs operate on sequences of static graph snapshots that are evenly sampled. Many discrete-time TGNNs integrate spatial and temporal information by combining GNNs with sequence models (Sankar et al. 2018; Seo et al. 2018; Pareja et al. 2020; Kazemi et al. 2020). Especially, DySAT (Sankar et al. 2018) leverages structural and temporal self-attention on DTDGs.

The continuous-time TGNNs to which our work relates directly learn the time-aware node embeddings on CTDGs, allowing them to capture finer temporal patterns compared to discrete-time TGNNs. Specifically, JODIE (Kumar, Zhang, and Leskovec 2019) employs two Recurrent Neural Network (RNN) modules to sequentially update the source and destination node embeddings as new edges arrive. TGAT (Xu et al. 2020) employs an attention-based

message-passing to aggregate messages from a historical neighborhood. To alleviate the expensive neighborhood aggregation of TGAT, TGN (Rossi et al. 2020) employs an RNN memory module to encode the history of each node. Our proposed method mainly focuses on targeting the above three representative continuous-time TGNNs.

In addition to these, various studies (Jin, Li, and Pan 2022; Wang et al. 2021a; Zhou et al. 2022; Cong et al. 2023) have been conducted on continuous-time TGNNs. Most research on continuous-time TGNNs has demonstrated their advancement using link prediction as a downstream task, so we mainly investigate link prediction as our target task to study the effect of adversarial attacks.

Adversarial Attacks on Graphs. Many studies have delved into adversarial attacks on static graphs (Sun et al. 2022; Jin et al. 2021). Moreover, graph adversarial attacks have been explored from diverse angles, such as perturbation types (Liu et al. 2019) and unnoticeability (Zügner, Akbarnejad, and Günnemann 2018; Dai et al. 2023).

Some studies consider adversarial attacks on DTDGs, focusing on degrading the link prediction performance. Specifically, TGA (Chen et al. 2021) is a white-box evasion attack that greedily selects edge perturbations across graph snapshots. TD-PGD (Sharma et al. 2023) is a Projected Gradient Descent (PGD) based attack under temporal dynamics-aware perturbation constraint. However, these methods are only applicable to DTDGs and cannot generate edge perturbations with continuous timestamps. Despite diverse studies, adversarial attacks on CTDGs are insufficiently studied.

Adversarial Defenses on Graphs. A bunch of work has proposed adversarial defense methods for static graphs (Mujkanovic et al. 2022). Especially, GCN-SVD (Entezari et al. 2020) substitutes the adjacency matrix with its low-rank approximation before plugging it into a regular GNN. ProGNN (Jin et al. 2020) alternately optimizes the parameters of the GNN and the graph structure. GNNGuard (Zhang and Zitnik 2020) filters edges in the message passing stage using cosine-similarity. However, these methods cannot be directly applied to CTDGs, and repeatedly applying them whenever new interactions occur is computationally prohibitive, and further fail to consider the dynamics in graphs.

In contrast, the literature on adversarial defense for dynamic graphs is limited. AdaNet (Li et al. 2022) proposes an adaptive neighbor selection technique for reliable message propagation on CTDGs. However, the method demonstrates its robustness against small amounts of contextual noise such as outdated links and node feature noise, rather than structural perturbations that are our main focus.

3 Problem Formulation

3.1 Notation and Preliminaries

From now on, unless specified otherwise, by a dynamic graph we refer to CTDG. A dynamic graph denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a sequence of temporal interactions, where \mathcal{V} is the set of nodes, $\mathcal{E} = \{e_1, e_2, \dots, e_{|\mathcal{E}|}\}$ is the set of interactions. Each interaction $e_i = (u_i, v_i, t_i)$ is a dyadic event between two nodes $u_i, v_i \in \mathcal{V}$ occurring at a specific

time $t_i \in \mathbb{R}^+$. For simplicity, we first assume these interactions are undirected, unattributed, and ordered chronologically (i.e., $t_i \leq t_{i+1}$). Later, we will discuss how our proposed method can be extended to attributed dynamic graphs.

TGNNs are commonly trained using temporal interactions as supervision (Thomas et al. 2022). Therefore, the capability of a model is evaluated by how accurately it may predict future interactions based on historical events. Specifically, given two nodes u and v at time t , TGNNs aim to learn their time-aware embeddings $\mathbf{h}_u(t)$ and $\mathbf{h}_v(t)$, where the presence of an interaction between them can be predicted with a downstream classifier, i.e., $\hat{y}_{uv,t} = \text{clf}(\mathbf{h}_u(t), \mathbf{h}_v(t))$. The loss function is typically formulated as:

$$\mathcal{L}_{\mathcal{E}} = - \sum_{(u,v,t) \in \mathcal{E}} \log \hat{y}_{uv,t} + \mathbb{E}_{n \sim P_n(\mathcal{V})} \log(1 - \hat{y}_{unt}), \quad (1)$$

where $P_n(\mathcal{V})$ is the negative sampling distribution. In this paper, we refer to the classifier $\text{clf}(\cdot)$ as the *edge scorer* and the output $\hat{y}_{uv,t}$ as the *edge score* of the edge (u, v, t) .

3.2 Dynamic Graph Adversarial Attack

Attacker’s Goal. Given a dynamic graph, an attacker makes slight yet intentional changes to the data to achieve a purpose, for instance, deteriorating the prediction performance of the models. In this work, we assume that the goal of the attacker is to degrade the dynamic link prediction performance of TGN models by injecting temporal adversarial edges into the original edge sequence. Formally, a corrupted dynamic graph is denoted as $\tilde{\mathcal{G}} = (\mathcal{V}, \mathcal{E} \cup \tilde{\mathcal{E}})$, where $\tilde{\mathcal{E}} = \{\tilde{e}_1, \tilde{e}_2, \dots, \tilde{e}_{\Delta}\}$ is the set of adversarial edges and Δ is the perturbation budget. Our proposed attack takes place before the victim model is trained, i.e., it is a poisoning attack.

Constraint for Unnoticeability. To ensure the effectiveness and stealthiness of adversarial attacks, we impose four constraints on dynamic graph attacks:

(C1) Perturbation budget: To control the extent of adversarial perturbations, we impose a budget constraint Δ on the attacks, limiting the number of adversarial edges to be no more than a budget, i.e., $|\tilde{\mathcal{E}}| \leq \Delta$. We control the budget with the perturbation rate p , i.e., $\Delta = \lfloor p \cdot |\mathcal{E}| \rfloor$.

(C2) Distribution of time: To enhance the stealthiness, we impose a distribution constraint on the timing of the adversarial edges. To this end, when adding an adversarial edge $\tilde{e} = (\tilde{u}, \tilde{v}, \tilde{t})$ at time \tilde{t} , we ensure that the time \tilde{t} follows the time distribution observed in the original edges, i.e., $\tilde{t} \sim P_t(\mathcal{E})$, where P_t is the time distribution¹. This constraint aims to align the occurrences of adversarial edges with the natural flow of interactions in dynamic graphs, making the adversarial edges less distinguishable from genuine edges.

(C3) Endpoints of adversarial edges: If an adversarial edge is composed of nodes that were activated a significant time ago, it becomes easily detectable because it deviates from the typical pattern of interactions. Thus, we enforce

a restriction that adversarial edges can only be formed using nodes from interactions that occurred within a specific time window W . Formally, when an adversarial edge \tilde{e} is added at time \tilde{t} , where $t_i \leq \tilde{t} \leq t_{i+1}$, we select two nodes \tilde{u} and \tilde{v} from the set of nodes comprising the preceding W edges, i.e., $\tilde{u}, \tilde{v} \in \mathcal{V}_{i,W} = \{u_{i-W+1}, v_{i-W+1}, \dots, u_i, v_i\}$. This constraint ensures that the adversarial edges are derived from recent interactions, making them less conspicuous.

(C4) Number of perturbations per node: A sudden increase in the degree of nodes can make the attack easily detectable (Bhatia et al. 2020). By limiting the number of perturbations per node to N , the attack becomes less noticeable. It is important to note that N is not a fixed value for the entire graph; it can vary over time, considering the dynamics of the original interactions in dynamic graphs.

Attacker’s Knowledge. We focus on limited-knowledge attacks, in which the attacker lacks knowledge about the victim model, including its architecture and trained weights. However, the attacker does possess the same level of knowledge about the data as the victim model, allowing them to observe all the dynamic graph data \mathcal{G} and to study the temporal interactions within the dynamic graph. Note that while our primary focus is on targeting TGN (Rossi et al. 2020), one of the representative TGNNs, the scope is not limited to this model. We investigate the transferability of the proposed attack, demonstrating its effectiveness for other victim models. These settings are realistic because the attacker cannot know the details of the victim model in real situations.

4 Proposed Attack Method: T-SPEAR

In this section, we introduce our proposed adversarial attack method T-SPEAR. Our attack exploits the edge scores ($\hat{y}_{uv,t}$) to create adversarial edges, leveraging the fact that a low edge score indicates the edge is unlikely to be formed, and thus likely to corrupt the natural evolution of the dynamic graph. The method aims to strategically generate and inject adversarial edges into the original edge sequence while adhering to the four constraints (C1-C4).

4.1 Surrogate Model

We face the challenge of limited knowledge about the victim model’s parameters and architecture. To address this issue, we employ a surrogate model, specifically the TGN model, which serves as a substitute for the victim model. This enables us to develop and optimize adversarial edges without directly accessing the victim model’s internals.

Node Memory. TGN uses node memory to summarize past interactions. When there is an interaction involving a node, the memory of that node is updated. For each interaction (u, v, t) , a message from v to u is computed as follows:

$$\mathbf{m}_u(t) = [\mathbf{s}_u(t^-) \parallel \mathbf{s}_v(t^-) \parallel \Phi(t - t_u^-)] e_{uv}. \quad (2)$$

Here, $\mathbf{s}_u(t^-)$ is the memory of node u just before time t , \parallel is the concatenation operator, and $\Phi(\Delta t) = \cos(\omega \Delta t + \phi)$ is the time encoder which encodes the time interval $\Delta t = t - t_u^-$ into a vector, where ω and ϕ are learnable vectors. Edge features are denoted as e_{uv} , but they can be ignored if not provided. The node memory is then updated as:

¹We use Kernel Density Estimation (KDE) to approximate the time distribution of the original edges. We employ the Gaussian kernel with a bandwidth of 0.1. The timestamps of the adversarial edges are drawn from the estimated temporal distribution.

$$s_u(t) = \text{mem}(s_u(t^-), m_u(t)), \quad (3)$$

where $\text{mem}(\cdot)$ is the GRU (Cho et al. 2014) memory updater. The memory of node v is also updated in the same manner.

Attention Aggregator. TGN employs the attention mechanism of Transformer (Vaswani et al. 2017) to aggregate information from temporal neighbors. The attention aggregation of node u is computed by the queries, keys, and values from its temporal neighbors $v \in \mathcal{N}_u(t)$, where $\mathcal{N}_u(T) = \{v : (u, v, t) \in \mathcal{E}, t \leq T\}$. The temporal embedding $h_u(t)$ of node u at time t can be computed by:

$$h_u(t) = \sum_{v \in \mathcal{N}_u(t)} \text{attn}(s_u(t), s_v(t), e_{uv}, v_u(t), v_v(t)), \quad (4)$$

where $\text{attn}(\cdot)$ is the temporal graph attention module (Xu et al. 2020), and $v_u(t)$ is the node features. The node features can also be disregarded if they are not provided. The attention aggregator enables TGN to integrate relevant information from its temporal neighbors, considering their importance in the context of the dynamic graph at time t .

Train the Surrogate Model. As discussed in Section 3.1, a surrogate TGN is trained in conjunction with a link classifier, which is responsible for predicting the occurrence of future interactions based on the time-aware node embeddings. After training, the classifier (i.e., edge scorer) estimates the likelihood of edge formation (i.e., edge score).

4.2 Naïve Adversarial Edge Selection

Given a well-trained surrogate model and its associated edge scorer, to create the adversarial edges $\tilde{\mathcal{E}} = \{\tilde{e}_1, \tilde{e}_2, \dots, \tilde{e}_\Delta\}$, we first determine the timestamps, i.e., $\tilde{t}_1, \dots, \tilde{t}_\Delta$, at which the adversarial edges are to be injected. Ensuring stealthiness, we sample $\tilde{t}_1, \dots, \tilde{t}_\Delta \sim P_t(\mathcal{E})$ to align them with the time distribution of the original edges (C2).

Next, for each time sampled \tilde{t} , we select two nodes \tilde{u} and \tilde{v} based on their edge score. To satisfy the constraint (C3), we choose two nodes from the set of nodes that comprise the previous W edges. When $t_i \leq \tilde{t} \leq t_{i+1}$ holds, we compute edge scores for all pairs of nodes in $\mathcal{V}_{i,W}$ using the node embeddings at time \tilde{t} . The node pair with the lowest edge score is chosen as the endpoints of the adversarial edge. It is formally described as follows:

$$(\tilde{u}, \tilde{v}) = \arg \min_{(u,v) \in \mathcal{E}(\mathcal{V}_{i,W})} \text{clf}(h_u(\tilde{t}), h_v(\tilde{t})). \quad (5)$$

Here, $\mathcal{E}(\mathcal{V}_{i,W}) = \{(u, v) | u, v \in \mathcal{V}_{i,W}, u \neq v\}$ denotes all pairs of nodes within $\mathcal{V}_{i,W}$. The lowest edge score indicates that the corresponding edge is the least likely to be formed at time \tilde{t} . Thus, the proposed adversarial edge selection rule may effectively perturb the temporal dynamics of the graph. We repeat the edge selection Δ times for all timestamps $\tilde{t}_1, \dots, \tilde{t}_\Delta$, thereby satisfying the constraint (C1).

4.3 Advanced Edge Selection on a Batch

Selecting adversarial edges one by one is effective but it may become impractical when dealing with large datasets. To overcome this issue, we propose an alternative where adversarial edges are selected in batches, improving efficiency.




Figure 1: Comparison of Low-K selection (left) and Hungarian selection (right) when K is 5. The Hungarian selection provides more inconspicuous and balanced perturbations.

The batch size $|B|$ is set to half of the window size $\lfloor |W|/2 \rfloor$ to adhere the constraint (C3), and we generate $K = \lfloor p \cdot |B| \rfloor$ adversarial edges per batch.

For each batch b , we first sample K timestamps from the time distribution $P_t(\mathcal{E})$, i.e., $\tilde{t}_{b,1}, \dots, \tilde{t}_{b,K} \sim P_t(\mathcal{E})$, where $t_{b,1} \leq \tilde{t}_{b,1} \leq \dots \leq \tilde{t}_{b,K} \leq t_{b,|B|}$. Here, $t_{b,1}$ and $t_{b,|B|}$ are the first and the last timestamps of the original edges in the batch b , respectively. Let \mathcal{V}_{b-1} be a set of nodes constituting the edges in the previous batch $b-1$. We then compute the edge scores for all node pairs within \mathcal{V}_{b-1} using the node embeddings at time $\tilde{t}_{b,1}$, which is the first timestamp of the perturbations. For an edge $(u, v, \tilde{t}_{b,1})$, its edge score is $\hat{y}_{uv\tilde{t}_{b,1}}$. Based on these edge scores, we select K edges as the adversarial edges.

Low-K Selection. The simplest way to select K edges is to pick the K edges with the lowest edge scores. However, we found that this method often leads to a substantial overlap of nodes. While it effectively corrupts information related to specific nodes, it may not sufficiently degrade the overall link prediction performance. Moreover, this method can result in a sudden increase in high-degree nodes, which violates constraint (C4), making the attack more detectable.

Hungarian Selection. To ensure the chosen edges are both impactful and inconspicuous, it is important to select adversarial edges with low edge scores and minimal overlapping endpoints. To this end, we reframe our challenge as the problem of assignment between the nodes in \mathcal{V}_{b-1} . Formally, let x_{uv} be an assignment indicator where $x_{uv} = 1$ if node u is matched with node v , otherwise $x_{uv} = 0$. We only allow one-to-one matching between distinct nodes. The optimal set of edges made within \mathcal{V}_{b-1} , denoted as E^* , can be obtained by solving our assignment problem formulated as:

$$E^* = \arg \min_{\{(u,v) | x_{uv}=1\}} \sum_{u \in \mathcal{V}_{b-1}} \sum_{v \in \mathcal{V}_{b-1}} \hat{y}_{uv\tilde{t}_{b,1}} x_{uv}, \quad (6)$$

$$\text{subject to } \sum_{v \in \mathcal{V}_{b-1}} x_{uv} = \sum_{v \in \mathcal{V}_{b-1}} x_{vu} = 1 \text{ and } x_{uu} = 0, \forall u \in \mathcal{V}_{b-1}.$$

Note that, if the sets of nodes available for the source and destination nodes differ (i.e., a bipartite dynamic graph), the problem changes to a rectangular assignment problem (Crouse 2016). To solve this problem, we employ the Hungarian algorithm (Jonker and Volgenant 1988), one of the most well-known solvers of an assignment problem. Solving this problem ensures that the selected edges have non-overlapping nodes, which helps improve the unnoticeability when considering (C4). Figure 1 illustrates the concept of the Hungarian selection and its impact on reducing

the occurrence of overlapping nodes. Furthermore, we apply the Hungarian algorithm N times, excluding previously selected edges, to create the candidate pool consisting of the selected edges. From this pool, we choose the K edges having the lowest edge score as the adversarial edges. Then, we randomly assign $\tilde{t}_{b,1}, \dots, \tilde{t}_{b,K}$ to the chosen edges. This approach allows us to limit the maximum number of adversarial edges per node to N , ensuring balanced contamination of all nodes in the graph while maintaining stealthiness (C4). In this paper, we maintain N as the minimum number of times that all perturbations can be selected, i.e., $N = \lceil \frac{K}{E_{\max}(\mathcal{V}_{b-1})} \rceil$, where $E_{\max}(\mathcal{V}_{b-1})$ is the maximum number of non-overlapping edges can be made within \mathcal{V}_{b-1} .

5 Proposed Defense Method: T-SHIELD

The goal of adversarial defense is to enhance the model’s robustness against perturbations and alleviate performance degradation. The defense for dynamic graphs is particularly challenging because the criteria for judging legitimate and adversarial edges change over time in CTDGs. To attain this goal, we identify and eliminate potential adversarial edges from the corrupted graphs considering the evolution of dynamic graphs without prior knowledge of adversarial attacks. Moreover, we enforce temporal smoothness on node embeddings to prevent sudden changes. In the following subsections, we elaborate on the proposed defense method called T-SHIELD.

5.1 Filtering Potential Adversarial Edges

T-SHIELD selectively excludes potential adversarial edges from the training process. For this, we leverage the edge scores generated by the link classifier (i.e., edge scorer), which is trained alongside the TGNN model.

For each edge $e = (u, v, t)$, we first compute its edge score $\hat{y}_{uv,t}$, and if the edge score is less than a threshold $\tau \in (0, 1)$, we classify it as a potential adversarial edge, i.e., $e \in \hat{\mathcal{E}}$, where $\hat{\mathcal{E}}$ is the set of potential adversarial edges. Subsequently, we can train the model on the cleaned dynamic graph $\mathcal{G}_c = (\mathcal{V}_c, \mathcal{E}_c)$, where \mathcal{V}_c is the set of remaining nodes and $\mathcal{E}_c = \mathcal{E} \cup \tilde{\mathcal{E}} - \hat{\mathcal{E}}$, obtained by removing the set of potential adversarial edges $\hat{\mathcal{E}}$ from the corrupted dynamic graph.

However, using a fixed threshold for edge filtering could lead to two significant issues. First, in the initial stages of training, the model and edge scorer might not be sufficiently trained, resulting in low confidence in the edge scores. Setting a high fixed threshold during this phase could remove numerous legitimate edges, impacting the model’s performance. Second, as training progresses and the model and edge scorer become well-trained, the edge scores become more reliable. In this later stage, using a low fixed threshold could increase the likelihood of failing to filter out adversarial edges effectively, leaving the model vulnerable to adversarial attacks. To address these concerns, we employ the cosine annealing scheduler to gradually increase the threshold from τ_s to τ_e where $0 \leq \tau_s \leq \tau_e \leq 1$. This approach provides a more adaptive and balanced filtering approach during the training process. The cleaned dynamic graph is then trained with the loss function $\mathcal{L}_{\mathcal{E}_c}$ (Eq. (1)).

5.2 Temporal Smoothness

Adversarial edges may change the node embeddings dramatically, preventing the model from learning the dynamics graphs. We ensure the temporal smoothness of node embeddings to mitigate the impact of adversarial edges that evade edge filtering. Maintaining temporal smoothness limits the difference between two consecutive embeddings to be relatively small, preventing the embeddings from deviating too dramatically from recent history. To impose temporal smoothness, we introduce a regularization term as follows:

$$\mathcal{L}_{tmp} = - \sum_{(u,v,t) \in \mathcal{E}_c} \sum_{i \in \{u,v\}} w(t - t_i^-) \text{sim}(\mathbf{h}_i(t), \mathbf{h}_i(t_i^-)), \quad (7)$$

where $w(\Delta t) = \exp(-\theta \Delta t)$ is the weight function that assigns higher weights to interactions within shorter periods and $\text{sim}(\cdot)$ denotes the cosine similarity. We set θ to 0.01.

Finally, by integrating edge filtering and temporal regularization, the loss function of T-SHIELD is the weighted sum of $\mathcal{L}_{\mathcal{E}_c}$ and \mathcal{L}_{tmp} , i.e., $\mathcal{L} = \mathcal{L}_{\mathcal{E}_c} + \lambda \mathcal{L}_{tmp}$, where λ is a hyperparameter to control the extent of temporal smoothness.

6 Complexity Analysis

The time complexities of T-SPEAR and T-SHIELD are presented in Lemmas 1 and 2, respectively. To focus on the time complexities of the proposed methods, we denote the time complexity of TGN for the inference of a single node’s embedding as C_{tgn} and the time complexity of the edge scorer as C_{clf} . Here, d denotes the dimension of node embeddings. The proofs are provided in Appendix B.

Lemma 1 (Time complexity of T-SPEAR). *The time complexity of T-SPEAR is $O((C_{tgn} + |W| \cdot C_{clf} + p \cdot |W|^2) \cdot |\mathcal{E}|)$.*

Lemma 2 (Time complexity of T-SHIELD). *The time complexity of T-SHIELD is $O((C_{tgn} + C_{clf} + d) \cdot |\mathcal{E}|)$.*

7 Experiments

7.1 Experimental Setup

Datasets. We conduct experiments on 4 real-world datasets: Wikipedia (Kumar, Zhang, and Leskovec 2019), MOOC (Feng, Tang, and Liu 2019), UCI (Panzarasa, Opsahl, and Carley 2009), and Bitcoin-OTC (Kumar et al. 2016). The details of them are provided in Appendix A.2. Across all datasets, we use the same 70%/15%/15% chronological splits for the train/validation/test sets as in (Rossi et al. 2020). Then, to make attacked graphs, we inject perturbations on all the split sets using the attack method.

Victim Models. We assess the proposed methods using 4 TGNNs: TGN (Rossi et al. 2020), JODIE (Kumar, Zhang, and Leskovec 2019), TGAT (Xu et al. 2020), and DySAT² (Sankar et al. 2018) (see Section 2 for their summaries). Our evaluation mainly focuses on TGN (see Section 4.1 for its details), while the other models are utilized to verify the transferability of our proposed attack.

²Technically, DySAT is a discrete-time TGNN.

Dataset	Victim	Clean	RANDOM	PA	JACCARD*	STRUCT-D	STRUCT-PR	T-SPEAR
Wikipedia	TGN	80.5 ± 0.5	66.0 ± 0.6	65.3 ± 1.1	-	66.4 ± 0.5	66.9 ± 0.9	60.0 ± 1.6
	JODIE	63.2 ± 1.2	36.3 ± 3.4	35.0 ± 2.3	-	33.0 ± 2.1	35.0 ± 2.5	<u>33.8 ± 2.4</u>
	TGAT	57.0 ± 0.6	41.4 ± 0.6	39.2 ± 1.0	-	41.1 ± 0.6	40.8 ± 0.5	<u>40.7 ± 0.5</u>
	DYSAT	66.6 ± 0.4	63.0 ± 0.5	62.4 ± 0.5	-	<u>62.1 ± 0.2</u>	62.8 ± 0.4	62.0 ± 0.4
MOOC	TGN	61.8 ± 2.1	55.4 ± 1.7	56.0 ± 1.6	-	50.3 ± 1.4	53.2 ± 1.2	48.2 ± 1.3
	JODIE	37.4 ± 2.0	23.8 ± 0.6	29.3 ± 3.9	-	26.2 ± 1.8	26.6 ± 1.6	22.8 ± 1.9
	TGAT	11.8 ± 0.1	10.2 ± 0.2	10.0 ± 0.1	-	10.7 ± 0.2	<u>9.6 ± 0.1</u>	9.3 ± 0.2
	DYSAT	18.4 ± 0.1	14.5 ± 0.1	14.5 ± 0.3	-	14.5 ± 0.2	<u>14.4 ± 0.2</u>	14.3 ± 0.2
UCI	TGN	44.2 ± 0.4	42.5 ± 0.3	44.2 ± 0.9	43.0 ± 0.5	42.5 ± 0.5	42.9 ± 0.8	41.6 ± 0.3
	JODIE	24.9 ± 0.4	21.9 ± 0.5	22.3 ± 0.4	21.9 ± 0.2	21.8 ± 0.6	21.4 ± 0.5	21.7 ± 0.2
	TGAT	16.3 ± 0.3	15.5 ± 0.3	<u>14.7 ± 0.1</u>	14.9 ± 0.2	14.7 ± 0.2	15.7 ± 0.4	14.6 ± 0.2
	DYSAT	71.3 ± 0.9	65.3 ± 0.7	<u>65.3 ± 0.3</u>	65.3 ± 0.5	65.4 ± 0.3	65.3 ± 0.6	64.7 ± 0.4
Bitcoin	TGN	48.3 ± 0.7	48.0 ± 1.3	48.0 ± 1.0	48.4 ± 0.9	46.8 ± 0.5	<u>46.2 ± 0.7</u>	45.2 ± 1.4
	JODIE	34.6 ± 0.7	31.4 ± 0.5	30.3 ± 0.6	32.1 ± 0.6	29.2 ± 0.8	27.9 ± 0.9	28.9 ± 0.7
	TGAT	21.8 ± 0.4	18.6 ± 0.3	17.3 ± 0.3	18.5 ± 0.4	17.1 ± 0.4	16.6 ± 0.6	<u>17.1 ± 0.2</u>
	DYSAT	84.7 ± 0.7	79.0 ± 0.8	<u>78.9 ± 0.2</u>	79.4 ± 0.6	79.0 ± 0.4	79.0 ± 0.2	78.8 ± 0.2
A.P.D.↓		-	-14.7%	-14.9%	-7.5%	-16.3%	<u>-16.4%</u>	-18.9%

*JACCARD cannot be applied to bipartite graphs.

Table 1: MRR at a perturbation rate $p = 0.3$. For each victim model, the best and second-best performances are highlighted in boldface and underlined, respectively. A.P.D. stands for average performance degradation.

Baselines. We compare T-SPEAR and T-SHIELD with various baseline approaches including 5 edge-perturbation-based adversarial attacks and 2 defense methods. A detailed description of these baselines is provided in Appendix A.3.

Evaluation Protocols. We adopt the Mean Reciprocal Rank (MRR) and hit rate (specifically, Hit@10) of the true edge among 100 randomly selected negative edges (without replacement) as a validation metric. Note that, all metrics are normalized by multiplying 100.

Because our settings involve perturbations in both the validation and test sets, we evaluate the metrics for all edges (including both ground truth and adversarial edges) in the validation set. For the test set, we only evaluate the metrics for the ground truth edges to be a fair comparison. Especially, since T-SHIELD and TGN-COSINE employ edge filtering, the evaluation process differs from others; for the validation set, we assess the metrics for the remaining edges after edge filtering, on the other hand, for the test set, we evaluate the metrics for the ground-truth edges before edge filtering. All the results are averaged over 5 runs. The results evaluated by Hit@10 are provided in Appendix C.

Implementation Details. Details on the implementation and hyperparameters of all victim models, attack models, and defense models are provided in Appendix A.4.

7.2 Attack Performance

We conduct a comprehensive comparison of our attack method against various baseline approaches. Table 1 summarizes the link prediction performance of all victim models under different attack methods with a perturbation rate of 0.3. Across all datasets and victim models, our proposed attack consistently outperforms the baselines (lower values indicate better attack performance). Overall, T-SPEAR

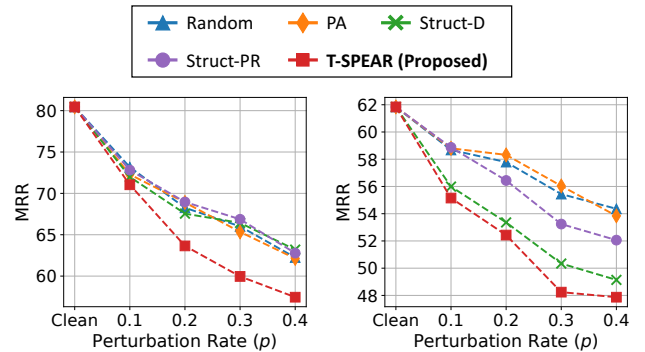


Figure 2: Link prediction performance of TGN on Wikipedia (left) and MOOC (right) as the perturbation rate increases.

shows the most substantial performance degradation, averaging -18.9%. Moreover, our attack significantly degrades the performance of other TGNs (i.e., JODIE, TGAT, and DySAT), which differ from the surrogate model (i.e., TGN) used by T-SPEAR. In Figure 2, we observe the drop in link prediction performance of TGN on Wikipedia and MOOC datasets as the number of adversarial edges increases. T-SPEAR consistently achieves a more significant reduction in performance compared to other baselines in any perturbation rates. In Appendix C, we assess the efficacy of our attack under the scenario where the attacker lacks knowledge about the embedding size of the victim model.

T-SPEAR demonstrates particular effectiveness when the victim models are TGN and JODIE. This is because these models incorporate a memory module to retain historical node interactions. As the perturbations generated by our attack are the least likely edges to be formed when considering the graphs' dynamics, these perturbations effectively

Model	Wikipedia (Clean: 80.5 ± 0.5)			MOOC (Clean: 61.8 ± 2.1)			UCI (Clean: 44.2 ± 0.4)			A.P.G.↑
	$p = 0.1$	$p = 0.2$	$p = 0.3$	$p = 0.1$	$p = 0.2$	$p = 0.3$	$p = 0.1$	$p = 0.2$	$p = 0.3$	
TGN	73.2 ± 1.1	68.3 ± 0.9	66.0 ± 0.6	58.7 ± 1.0	57.8 ± 0.9	55.4 ± 1.7	44.5 ± 0.3	43.5 ± 0.7	42.5 ± 0.3	-
TGN-SVD	65.3 ± 1.1	63.9 ± 1.2	60.6 ± 0.9	58.8 ± 0.3	58.0 ± 1.4	55.0 ± 1.1	44.5 ± 0.7	44.3 ± 0.5	44.7 ± 0.7	-2.1%
TGN-COSINE	75.8 ± 0.5	71.6 ± 0.7	68.8 ± 0.4	53.7 ± 2.7	52.8 ± 4.1	54.0 ± 3.2	43.2 ± 0.6	43.0 ± 0.6	43.8 ± 0.4	-0.9%
T-SHIELD-F	77.4 ± 0.5	77.4 ± 0.2	76.0 ± 0.3	59.2 ± 1.0	58.8 ± 1.3	58.0 ± 0.6	44.1 ± 0.2	44.2 ± 0.6	44.2 ± 0.7	+5.1%
T-SHIELD	77.3 ± 0.5	77.0 ± 0.3	76.5 ± 0.5	59.4 ± 1.2	58.4 ± 1.1	58.4 ± 0.6	43.9 ± 0.6	44.3 ± 0.7	44.6 ± 0.8	+5.3%

Table 2: MRR under RANDOM attack. A.P.G. stands for average performance gain over the naïve TGN.

Model	Wikipedia (Clean: 80.5 ± 0.5)			MOOC (Clean: 61.8 ± 2.1)			UCI (Clean: 44.2 ± 0.4)			A.P.G.↑
	$p = 0.1$	$p = 0.2$	$p = 0.3$	$p = 0.1$	$p = 0.2$	$p = 0.3$	$p = 0.1$	$p = 0.2$	$p = 0.3$	
TGN	71.0 ± 2.0	63.7 ± 2.6	60.0 ± 1.6	55.2 ± 0.6	52.4 ± 0.4	48.2 ± 1.3	43.6 ± 1.0	41.7 ± 0.5	41.6 ± 0.3	-
TGN-SVD	67.6 ± 0.9	61.2 ± 1.4	58.1 ± 1.8	55.7 ± 1.3	50.6 ± 2.6	50.1 ± 2.1	46.0 ± 0.5	45.5 ± 0.4	46.5 ± 0.4	+1.8%
TGN-COSINE	72.3 ± 0.8	62.9 ± 2.0	59.6 ± 2.1	53.5 ± 3.1	47.6 ± 1.7	45.2 ± 2.1	43.7 ± 0.4	44.1 ± 0.3	43.2 ± 1.3	-1.0%
T-SHIELD-F	77.3 ± 0.5	76.8 ± 0.3	75.3 ± 0.4	58.5 ± 0.5	55.7 ± 0.9	54.5 ± 0.9	44.4 ± 0.8	45.5 ± 0.3	44.7 ± 0.4	+11.0%
T-SHIELD	77.1 ± 0.2	76.9 ± 0.3	76.1 ± 0.3	58.8 ± 0.7	56.3 ± 1.2	53.9 ± 1.1	44.9 ± 0.3	44.3 ± 0.5	45.4 ± 0.8	+11.2%

Table 3: MRR under T-SPEAR attack. A.P.G. stands for average performance gain over the naïve TGN.

distort the node memory compared to baselines, which only consider basic graph statistics such as node degree, PageRank centrality, and the number of common neighbors. On the other hand, TGAT and DySAT operate without relying on node memory. They utilize a self-attention mechanism for adaptive message aggregation, which makes them less susceptible to perturbations over long periods. Nevertheless, since our attack injects perturbations considering the *current dynamics* of graphs, it can substantially reduce performance.

Additionally, we found that connecting two nodes with low node centrality (in terms of degree and PageRank centrality) is also an effective attack. This demonstrates that low-degree attacks (Hussain et al. 2021; Ma, Ding, and Mei 2020), which are effective in static graphs, also exhibit partial effectiveness in dynamic graphs.

7.3 Defense Performance

To demonstrate the generality of our proposed defense method, we evaluate its robustness under two distinct attacks: RANDOM and T-SPEAR. We first poison dynamic graphs using each attack, and then train the T-SHIELD and its baselines on the poisoned graphs to evaluate their link prediction performance. Tables 2 and 3 summarize the performance of the defense models against the RANDOM and T-SPEAR, respectively. T-SHIELD-F, which uses only edge filtering, is a variant of T-SHIELD to validate the effect of edge filtering itself. Notably, T-SHIELD and T-SHIELD-F consistently improve the link prediction performance across different attack methods, datasets, and perturbation rates. It verifies that our proposed defense method is a suitable defense method for real-world situations where the attacker’s details are unknown. The primary performance improvement comes from edge filtering, but enforcing temporal smoothness on node embeddings provides additional enhancement. On the other hand, baselines exhibit insufficient defense performance (even if carefully optimized), except for the UCI dataset. It indicates that cosine similarity is not appropriate

Attack	Model	Wikipedia	MOOC	UCI
RANDOM	TGN-COSINE	0.605	0.633	0.495
	T-SHIELD-F	0.848	0.697	0.531
	T-SHIELD	0.850	0.709	0.534
T-SPEAR	TGN-COSINE	0.474	0.556	0.518
	T-SHIELD-F	0.858	0.679	0.608
	T-SHIELD	0.864	0.674	0.620

Table 4: AUROC of adversarial edge classification at perturbation rate $p = 0.3$.

to filter out adversarial edges in bipartite graphs, and applying SVD to the entire training set can not capture the reliability of edges changing over time. To sum up, it is not straightforward to extend static defense methods to CTDGs.

Table 4 shows the AUROC of filtering-based defense models in classifying adversarial edges at $p = 0.3$, demonstrating that our proposed defense methods more accurately discriminate adversarial edges compared to the baselines.

8 Conclusion

In this paper, we propose T-SPEAR, an effective poisoning adversarial attack method for TGNN-based link prediction on CTDGs. While adhering to the four constraints for unnoticeability, we strategically generate adversarial edges, which are unlikely to be formed, and thus likely to corrupt the natural evolution of the dynamic graphs. Additionally, to mitigate the impact of adversarial attacks, we propose T-SHIELD, a robust training method for TGNNs. By employing the edge filtering technique and enforcing temporal robustness to node embeddings, we improve the robustness of the victim mode. Experimental results demonstrate that T-SPEAR significantly degrades victim model performance and is transferable across various TGNNs. Moreover, T-SHIELD yields significant performance improvements over the basic TGN model under different adversarial attacks.

Acknowledgements

This work was supported by Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-00157, Robust, Fair, Extensible Data-Centric Continual Learning) (No. 2019-0-00075, Artificial Intelligence Graduate School Program (KAIST)).

References

- Bhatia, S.; Hooi, B.; Yoon, M.; Shin, K.; and Faloutsos, C. 2020. Midas: Microcluster-based detector of anomalies in edge streams. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 3242–3249.
- Bojchevski, A.; and Günnemann, S. 2019. Adversarial attacks on node embeddings via graph poisoning. In *International Conference on Machine Learning*, 695–704. PMLR.
- Chen, J.; Zhang, J.; Chen, Z.; Du, M.; and Xuan, Q. 2021. Time-aware gradient attack on dynamic network link prediction. *IEEE Transactions on Knowledge and Data Engineering*.
- Cho, K.; Merriënboer, B.; Gulcehre, C.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.
- Cong, W.; Zhang, S.; Kang, J.; Yuan, B.; Wu, H.; Zhou, X.; Tong, H.; and Mahdavi, M. 2023. Do We Really Need Complicated Model Architectures For Temporal Networks? *arXiv:2302.11636*.
- Crouse, D. F. 2016. On implementing 2D rectangular assignment algorithms. *IEEE Transactions on Aerospace and Electronic Systems*, 52(4): 1679–1696.
- Dai, E.; Lin, M.; Zhang, X.; and Wang, S. 2023. Unnoticeable backdoor attacks on graph neural networks. In *Proceedings of the ACM Web Conference 2023*, 2263–2273.
- Dai, H.; Li, H.; Tian, T.; Huang, X.; Wang, L.; Zhu, J.; and Song, L. 2018. Adversarial attack on graph structured data. In *International Conference on Machine Learning*, 1115–1124. PMLR.
- Entezari, N.; Al-Sayouri, S. A.; Darvishzadeh, A.; and Papalexakis, E. E. 2020. All you need is low (rank) defending against adversarial attacks on graphs. In *Proceedings of the 13th International Conference on Web Search and Data Mining*, 169–177.
- Feng, W.; Tang, J.; and Liu, T. X. 2019. Understanding dropouts in MOOCs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 517–524.
- Fursov, I.; Morozov, M.; Kaploukhaya, N.; Kovtun, E.; Rivera-Castro, R.; Gusev, G.; Babaev, D.; Kireev, I.; Zaytsev, A.; and Burnaev, E. 2021. Adversarial attacks on deep models for financial transaction records. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2868–2878.
- Hooi, B.; Song, H. A.; Beutel, A.; Shah, N.; Shin, K.; and Faloutsos, C. 2016. Fraudar: Bounding graph fraud in the face of camouflage. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 895–904.
- Hou, S.; Fan, Y.; Zhang, Y.; Ye, Y.; Lei, J.; Wan, W.; Wang, J.; Xiong, Q.; and Shao, F. 2019. α cyber: Enhancing robustness of android malware detection system against adversarial attacks on heterogeneous graph based model. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, 609–618.
- Hussain, H.; Duricic, T.; Lex, E.; Helic, D.; Strohmaier, M.; and Kern, R. 2021. Structack: Structure-based adversarial attacks on graph neural networks. In *Proceedings of the 32nd ACM Conference on Hypertext and Social Media*, 111–120.
- Jin, M.; Li, Y.-F.; and Pan, S. 2022. Neural temporal walks: Motif-aware representation learning on continuous-time dynamic graphs. *Advances in Neural Information Processing Systems*, 35: 19874–19886.
- Jin, W.; Li, Y.; Xu, H.; Wang, Y.; Ji, S.; Aggarwal, C.; and Tang, J. 2021. Adversarial attacks and defenses on graphs. *ACM SIGKDD Explorations Newsletter*, 22(2): 19–34.
- Jin, W.; Ma, Y.; Liu, X.; Tang, X.; Wang, S.; and Tang, J. 2020. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 66–74.
- Jonker, R.; and Volgenant, T. 1988. A shortest augmenting path algorithm for dense and sparse linear assignment problems. In *DGOR/NSOR*, 622–622. Springer.
- Kazemi, S. M.; Goel, R.; Jain, K.; Kobayez, I.; Sethi, A.; Forsyth, P.; and Poupart, P. 2020. Representation learning for dynamic graphs: A survey. *The Journal of Machine Learning Research*, 21(1): 2648–2720.
- Kumar, S.; Spezzano, F.; Subrahmanian, V.; and Faloutsos, C. 2016. Edge weight prediction in weighted signed networks. In *2016 IEEE 16th International Conference on Data Mining*, 221–230. IEEE.
- Kumar, S.; Zhang, X.; and Leskovec, J. 2019. Predicting dynamic embedding trajectory in temporal interaction networks. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1269–1278.
- Li, H.; Li, C.; Feng, K.; Yuan, Y.; Wang, G.; and Zha, H. 2022. Robust knowledge adaptation for dynamic graph neural networks. *arXiv:2207.10839*.
- 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. *arXiv:1910.14147*.
- Ma, J.; Ding, S.; and Mei, Q. 2020. Towards more practical adversarial attacks on graph neural networks. *Advances in Neural Information Processing Systems*, 33: 4756–4766.
- Mujkanovic, F.; Geisler, S.; Günnemann, S.; and Bojchevski, A. 2022. Are Defenses for Graph Neural Networks Robust? *Advances in Neural Information Processing Systems*, 35: 8954–8968.
- Panzarasa, P.; Opsahl, T.; and Carley, K. M. 2009. Patterns and dynamics of users’ behavior and interaction: Network

- analysis of an online community. *Journal of the American Society for Information Science and Technology*, 60(5): 911–932.
- Pareja, A.; Domeniconi, G.; Chen, J.; Ma, T.; Suzumura, T.; Kanezashi, H.; Kaler, T.; Schardl, T.; and Leiserson, C. 2020. Evolvecn: Evolving graph convolutional networks for dynamic graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 5363–5370.
- Rossi, E.; Chamberlain, B.; Frasca, F.; Eynard, D.; Monti, F.; and Bronstein, M. 2020. Temporal graph networks for deep learning on dynamic graphs. *arXiv:2006.10637*.
- Sankar, A.; Wu, Y.; Gou, L.; Zhang, W.; and Yang, H. 2018. Dynamic graph representation learning via self-attention networks. *arXiv:1812.09430*.
- Seo, Y.; Defferrard, M.; Vandergheynst, P.; and Bresson, X. 2018. Structured sequence modeling with graph convolutional recurrent networks. In *International Conference on Neural Information Processing*, 362–373. Springer.
- Sharma, K.; Trivedi, R.; Sridhar, R.; and Kumar, S. 2023. Temporal Dynamics-Aware Adversarial Attacks on Discrete-Time Dynamic Graph Models. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2023–2035.
- Sun, L.; Dou, Y.; Yang, C.; Zhang, K.; Wang, J.; Philip, S. Y.; He, L.; and Li, B. 2022. Adversarial attack and defense on graph data: A survey. *IEEE Transactions on Knowledge and Data Engineering*.
- Thomas, J. M.; Moallem-Oureh, A.; Beddar-Wiesing, S.; and Holzhüter, C. 2022. Graph neural networks designed for different graph types: A survey. *arXiv:2204.03080*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Wang, X.; Lyu, D.; Li, M.; Xia, Y.; Yang, Q.; Wang, X.; Wang, X.; Cui, P.; Yang, Y.; Sun, B.; et al. 2021a. Apan: Asynchronous propagation attention network for real-time temporal graph embedding. In *Proceedings of the 2021 International Conference on Management of Data*, 2628–2638.
- Wang, Y.; Chang, Y.-Y.; Liu, Y.; Leskovec, J.; and Li, P. 2021b. Inductive representation learning in temporal networks via causal anonymous walks. *arXiv:2101.05974*.
- Xu, D.; Ruan, C.; Korpeoglu, E.; Kumar, S.; and Achan, K. 2020. Inductive representation learning on temporal graphs. *arXiv:2002.07962*.
- Xu, K.; Chen, H.; Liu, S.; Chen, P.-Y.; Weng, T. W.; Hong, M.; and Lin, X. 2019. Topology attack and defense for graph neural networks: An optimization perspective. In *International Joint Conference on Artificial Intelligence*.
- Zeager, M. F.; Sridhar, A.; Fogal, N.; Adams, S.; Brown, D. E.; and Beling, P. A. 2017. Adversarial learning in credit card fraud detection. In *2017 Systems and Information Engineering Design Symposium*, 112–116. IEEE.
- Zhang, X.; and Zitnik, M. 2020. Gnn-guard: Defending graph neural networks against adversarial attacks. *Advances in Neural Information Processing Systems*, 33: 9263–9275.
- Zhou, H.; Zheng, D.; Nisa, I.; Ioannidis, V.; Song, X.; and Karypis, G. 2022. Tgl: A general framework for temporal gnn training on billion-scale graphs. *arXiv:2203.14883*.
- Zügner, D.; Akbarnejad, A.; and Günnemann, S. 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.