

A Restricted Black-Box Adversarial Framework Towards Attacking Graph Embedding Models

Heng Chang,^{*1} Yu Rong,² Tingyang Xu,² Wenbing Huang,² Honglei Zhang,²
Peng Cui,³ Wenwu Zhu,^{†1,3} Junzhou Huang²

¹Tsinghua-Berkeley Shenzhen Institute, Tsinghua University, China

²Tencent AI Lab, China

³Department of Computer Science and Technology, Tsinghua University, China

changh17@mails.tsinghua.edu.cn, yu.rong@hotmail.com, Tingyangxu@tencent.com, hwenbing@126.com, zhanghonglei@gatech.edu, {cuip, wwzhu}@tsinghua.edu.cn, jzhuang@uta.edu

Abstract

With the great success of graph embedding model on both academic and industry area, the robustness of graph embedding against adversarial attack inevitably becomes a central problem in graph learning domain. Regardless of the fruitful progress, most of the current works perform the attack in a white-box fashion: they need to access the model predictions and labels to construct their adversarial loss. However, the inaccessibility of model predictions in real systems makes the white-box attack impractical to real graph learning system. This paper promotes current frameworks in a more general and flexible sense – we demand to attack various kinds of graph embedding model with black-box driven. To this end, we begin by investigating the theoretical connections between graph signal processing and graph embedding models in a principled way and formulate the graph embedding model as a general graph signal process with corresponding graph filter. As such, a generalized adversarial attacker: *GF-Attack* is constructed by the graph filter and feature matrix. Instead of accessing any knowledge of the target classifiers used in graph embedding, *GF-Attack* performs the attack only on the graph filter in a black-box attack fashion. To validate the generalization of *GF-Attack*, we construct the attacker on four popular graph embedding models. Extensive experimental results validate the effectiveness of our attacker on several benchmark datasets. Particularly by using our attack, even small graph perturbations like one-edge flip is able to consistently make a strong attack in performance to different graph embedding models.

Introduction

Graph embedding models (Scarselli et al. 2009; Cui et al. 2018), which elaborate the expressive power of deep learning on graph-structure data, have achieved promising success in various domains, such as predicting properties over molecules (Duvenaud et al. 2015), biology analysis (Hamilton, Ying, and Leskovec 2017), financial surveillance (Paranjape, Benson, and Leskovec 2017) and structural role classification (Tu et al. 2018). Given the increasing popularity and success of these methods, a bunch of

recent works have posed the risk of graph embedding models against adversarial attacks, just like what the researchers are anxious for convolutional neural networks (Akhtar and Mian 2018). A strand of research works (Dai et al. 2018; Zügner, Akbarnejad, and Günnemann 2018; Bojchevski and Günnemann 2019) have already shown that various kinds of graph embedding methods, including Graph Convolutional Networks, DeepWalk, etc., are vulnerable to adversarial attacks. Undoubtedly, the potential attacking risk is rising for modern graph learning systems. For instance, by sophisticated constructed social bots and following connections, it’s possible to fool the recommendation system equipped with graph embedding models to give wrong recommendations.

Regarding the amount of information from both target model and data required for the generation of adversarial examples, all graph adversarial attackers fall into three categories (arranged in an ascending order of difficulties):

- **White-box Attack (WBA)**: the attacker can access any information, namely, the training input (e.g., adjacency matrix and feature matrix), the label, the model parameters, the predictions, etc.
- **Practical White-box Attack (PWA)**: the attacker can any information except the model parameters.
- **Restrict Black-box Attack (RBA)**: the attacker can only access the training input and limited knowledge of the model. The access of parameters, labels and predictions is prohibited.

Despite the fruitful results (Sun et al. 2018; Zügner, Akbarnejad, and Günnemann 2018; Zügner and Günnemann 2019) which absorb ingredients from exiting adversarial methods on convolutional neural networks, obtained in attacking graph embeddings under both WBA and PWA setting, however, the target model parameter as well as the labels and predictions are seldom accessible in real-life applications. In the other words, the WBA and PWA attackers are almost impossible to perform a threatening attack to real systems. Meanwhile, current RBA attackers are either reinforcement learning based (Dai et al. 2018), which has low computational efficiency and is limited to edge deletion, or derived merely only from the structure information without considering the feature information (Bojchevski and Günnemann

*Work done while interning at Tencent AI Lab.

†Wenwu Zhu is the corresponding author.

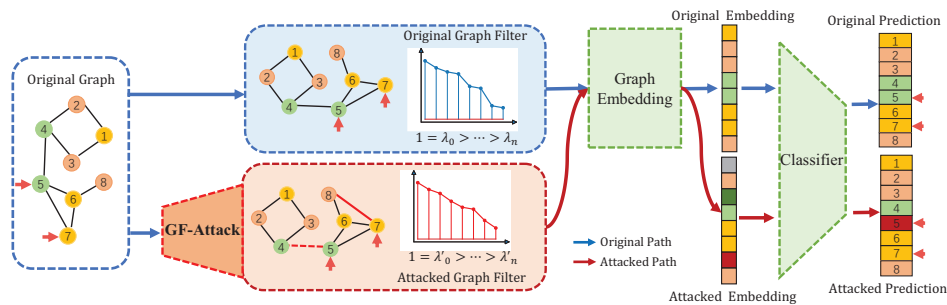


Figure 1: The overview of whole attack procedure of *GF-Attack*. Given target vertices 5 and 7, *GF-Attack* aims to misclassify them by attacking the graph filter and producing adversarial edges (edge e_{45} deleted and edge e_{78} added) on graph structure. The common graph embedding block refers to the general target GNN model and can be any kind of potential GNN models, illustrating the flexibility and extensibility of *GF-Attack*. In this vein, *GF-Attack* would not change the target embedding model.

2019). Therefore, how to perform the effective adversarial attack toward graph embedding model relying on the training input, a.k.a., RBA setting, is still more challenging yet meaningful in practice.

The core task of the adversarial attack on graph embedding model is to damage the quality of output embeddings to harm the performance of downstream tasks within the manipulated features or graph structures, i.e., vertex or edge insertion/deletion. Namely, finding the embedding quality measure to evaluate the damage of embedding quality is vital. For the WBA and PWA attackers, they have enough information to construct this quality measure, such as the loss function of the target model. In this vein, the attack can be performed by simply maximize the loss function reversely, either by gradient ascent (Dai et al. 2018) or a surrogate model (Zügner, Akbarnejad, and Günnemann 2018; Zügner and Günnemann 2019) given the known labels. However, the RBA attacker can not employ the limited information to recover the loss function of the target model, even constructing a surrogate model is impossible. In a nutshell, the biggest challenge of the RBA attacker is: how to figure out the goal of the target model barely by the training input.

In this paper, we try to understand the graph embedding model from a new perspective and propose an attack framework: *GF-Attack*, which can perform adversarial attack on various kinds of graph embedding models. Specifically, we formulate the graph embedding model as a general graph signal processing with corresponding graph filter which can be computed by the input adjacency matrix. Therefore, we employ the graph filter as well as feature matrix to construct the embedding quality measure as a T -rank approximation problem. In this vein, instead of attacking the loss function, we aim to attack the graph filter of given models. It enables *GF-Attack* to perform attack in a restrict black-box fashion. Furthermore, by evaluating this T -rank approximation problem, *GF-Attack* is capable to perform the adversarial attack on any graph embedding models which can be formulate to a general graph signal processing. Meanwhile, we give the quality measure construction for four popular graph embedding models (GCN, SGC, DeepWalk, LINE). Figure 1 provides the overview of whole attack procedure of *GF-Attack*. Empirical results show that our general attacking method

is able to effectively propose adversarial attacks to popular unsupervised/semi-supervised graph embedding models on real-world datasets without access to the classifier.

Related work

For explanation of graph embedding models, (Xu et al. 2018) and (Qiu et al. 2018) show some insights on the understanding of Graph Convolutional Networks and sampling-based graph embedding, respectively. However, they focus on proposing new graph embedding frameworks in each type of methods rather than building up a theoretical connection.

Only recently adversarial attacks on deep learning for graphs have drawn unprecedented attention from researchers. (Dai et al. 2018) exploits a reinforcement learning based framework under RBA setting. However, they restrict their attacks on edge deletions only for node classification, and do not evaluate the transferability. (Zügner, Akbarnejad, and Günnemann 2018) proposes attacks based on a surrogate model and they can do both edge insertion/deletion in contrast to (Dai et al. 2018). But their method utilizes additional information from labels, which is under PWA setting. Further, (Zügner and Günnemann 2019) utilizes meta-gradients to conduct attacks under black-box setting by assuming the attacker uses a surrogate model same as (Zügner, Akbarnejad, and Günnemann 2018). Their performance highly depends on the assumption of the surrogate model, and also requires label information. Moreover, they focus on the global attack setting. (Xu et al. 2019) also proposes a gradient-based method under WBA setting and overcomes the difficulty brought by discrete graph structure data. (Bojchevski and Günnemann 2019) considers a different adversarial attack task on vertex embeddings under RBA setting. Inspired by (Qiu et al. 2018), they maximize the loss obtained by DeepWalk with matrix perturbation theory while only consider the information from adjacent matrix. In contrast, we focus on semi-supervised learning on node classification combined with features. Remarkably, despite all above-introduced works except (Dai et al. 2018) show the existence of transferability in graph embedding methods by experiments, they all lack theoretical analysis on the implicit connection. In this work, for the first time, we theoretically connect different kinds of graph embedding models and propose a general optimization prob-

lem from parametric graph signal processing. An effective algorithm is developed afterwards under RBA setting.

Preliminary

Let $G(\mathcal{V}, \mathcal{E})$ be an attributed graph, where \mathcal{V} is a vertex set with size $n = |\mathcal{V}|$ and \mathcal{E} is an edge set. Denote $A \in \{0, 1\}^{n \times n}$ as an adjacent matrix and $X \in \mathbb{R}^{n \times l}$ as a feature matrix with dimension l . $D_{ii} = \sum_j A_{ij}$ refers the degree matrix. $\text{vol}(G) = \sum_i \sum_j A_{ij} = \sum_i D_{ii}$ denotes the volume of G . For consistency, we denote the perturbed adjacent matrix as A' and the normalized adjacent matrix as $\hat{A} = D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$. Symmetric normalized Laplacian and random walk normalized Laplacian are referred as $L^{sym} = I_n - D^{-\frac{1}{2}} A D^{-\frac{1}{2}}$ and $L^{rw} = I_n - D^{-1} A$.

Given a graph embedding model \mathcal{M}_Θ parameterized by Θ and a graph $G(\mathcal{V}, \mathcal{E})$, the adversarial attack on graph aims to perturb the learned vertex representation $Z = \mathcal{M}_\Theta(A, X)$ to damage the performance of the downstream learning tasks. Three components in graphs can be attacked as targets:

- Attack on \mathcal{V} : Add/delete vertices in graphs. This operation may change the dimension of the adjacency matrix A .
- Attack on A : Add/delete edges in graphs. This operation would lead to the changes of entries in the adjacency matrix A . This kind of attack is also known as *structural attack*.
- Attack on X : Modify the attributes attached on vertices.

Here, we mainly focus on adversarial attacks on graph structure A , since attacking A is more practical than others in real applications (Tong et al. 2012).

Adversarial Attack Definition

Formally, given a fixed budget β indicating that the attacker is only allowed to modify 2β entries in A (undirected), the adversarial attack on a graph embedding model \mathcal{M}_Θ can be formulated as (Bojchevski and Günnemann 2019):

$$\begin{aligned} \arg \max_{A'} \mathcal{L}(A', Z) \\ \text{s.t. } Z = \mathcal{M}_\Theta(A', X), \\ \Theta^* = \arg \min_{\Theta} \mathcal{L}(\Theta; A', X), \|A' - A\| = 2\beta, \end{aligned} \quad (1)$$

where Z is the embedding output of the model \mathcal{M}_Θ and $\mathcal{L}(\cdot, \cdot)$ is the loss function minimized by Θ . $\mathcal{L}(A', Z)$ is defined as the loss measuring the attack damage on output embeddings, lower loss corresponds to higher quality. For the WBA, $\mathcal{L}(A', Z)$ can be defined by the minimization of the target loss, i.e., $\mathcal{L}(A', Z) = \inf_{\Theta} \mathcal{L}(\Theta; A', Z)$. This is a bi-level optimization problem if we need to re-train the model during attack. Here we consider a more practical scenario: $\Theta^* = \arg \min_{\Theta} \mathcal{L}(\Theta; A, X)$ are learned on the clean graph and remains unchanged during attack.

Methodologies

Graph Signal Processing (GSP) focuses on analyzing and processing data points whose relations are modeled as graph (Shuman et al. 2013; Ortega et al. 2018). Similar to Discrete Signal Processing, these data points can be treated as *signals*.

Thus the definition of *graph signal* is a mapping from vertex set \mathcal{V} to real numbers $\mathbf{x} : \mathcal{V} \rightarrow \mathbb{R}$. In this sense, the feature matrix X can be treated as graph signals with l channels. From the perspective of GSP, we can formulate graph embedding model $\mathcal{M} : (A, X) \rightarrow \mathbb{R}^{n \times d}$ as the generalization of signal processing according to graph filter \mathcal{H} together with feature transformation:

$$\tilde{X} = \mathcal{H}(X), X' = \sigma(\tilde{X}\Theta), \quad (2)$$

where \mathcal{H} denotes a graph signal filter, $\sigma(\cdot)$ denotes the activation function, and $\Theta \in \mathbb{R}^{l' \times l}$ denotes a convolution filter from l input channels to l' output channels. \mathcal{H} can be constructed by a polynomial function $h(x) = \sum_{i=0}^L a_i x^i \in \mathbb{R}^{n \times n}$ with graph-shift filter S , i.e., $\tilde{X} = h(S)X$. Here, the graph-shift filter S reflects the locality property of graphs, i.e., it represents a linear transformation of the signals of one vertex and its neighbors. It's the basic building blocks to construct \mathcal{H} . We call this general model *Graph Filter Attack (GF-Attack)*. *GF-Attack* introduces the trainable weight matrix Θ to enable stronger expressiveness which can fuse the structural and non-structural information.

Embedding Quality Measure $\mathcal{L}(A', Z)$ of *GF-Attack*

According to (2), in order to avoid accessing the target model parameter Θ , we can construct the restricted black-box attack loss $\mathcal{L}(A', Z)$ by attacking the graph filter \mathcal{H} . Recent works (Yang et al. 2015; Nar et al. 2019) demonstrate that the output embeddings of graph embedding models can have very low-rank property. Since our goal is to damage the quality of output embedding Z , we establish the general optimization problem accordingly as a T -rank approximation problem inspired from (Qiu et al. 2018):

$$\mathcal{L}(A', Z) = \|h(S')X - h(S')_T X\|_F^2, \quad (3)$$

where $h(S')$ is the polynomial graph filter, S' is the graph shift filter constructed from the perturbed adjacency matrix A' . $h(S')_T$ is the T -rank approximation of $h(S')$. According to low-rank approximation, $\mathcal{L}(A', Z)$ can be rewritten as:

$$\mathcal{L}(A', Z) = \left\| \sum_{i=T+1}^n \lambda'_i \mathbf{u}_i \mathbf{u}_i^T X \right\|_F^2 \leq \sum_{i=T+1}^n \lambda'_i{}^2 \cdot \sum_{i=T+1}^n \|\mathbf{u}_i^T X\|_2^2, \quad (4)$$

where n is the number of vertices. $h(S') = U\Lambda U^T$ is the eigen-decomposition of the graph filter $h(S')$. $h(S')$ is a symmetric matrix. $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, $U = [\mathbf{u}_1^T, \dots, \mathbf{u}_n^T]$ are the eigenvalue and eigenvector of graph filter \mathcal{H} , respectively, in order of $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. λ'_i is the corresponding eigenvalue after perturbation. While $\left\| \sum_{i=T+1}^n \lambda_i \mathbf{u}_i \mathbf{u}_i^T X \right\|$ is hard to optimized, from (4), we can compute the upper bound instead of minimizing the loss directly. Accordingly, the goal of adversarial attack is to maximize the upper bound of the loss reversely. Thus the restrict black-box adversarial attack is equivalent to optimize:

$$\begin{aligned} \arg \max_{A'} \sum_{i=T+1}^n \lambda'_i{}^2 \cdot \sum_{i=T+1}^n \|\mathbf{u}_i^T X\|_2^2, \\ \text{s.t. } \|A' - A\| = 2\beta. \end{aligned} \quad (5)$$

Now our adversarial attack model is a general attacker. Theoretically, we can attack any graph embedding model which can be described by the corresponding graph filter \mathcal{H} . Meanwhile, our general attacker provides theoretical explanation on the transferability of adversarial samples created by (Zügner, Akbarnejad, and Günnemann 2018; Zügner and Günnemann 2019; Bojchevski and Günnemann 2019), since modifying edges in adjacent matrix A implicitly perturbs the eigenvalues of graph filters. In the following, we will analyze two kinds of popular graph embedding methods and aim to perform adversarial attack according to (5).

GF-Attack on Graph Convolutional Networks

Graph Convolution Networks extend the definition of convolution to the irregular graph structure and learn a representation vector of a vertex with feature matrix X . Namely, we generalize the Fourier transform to graphs to define the convolution operation: $g_\theta * \mathbf{x} = U g_\theta U^T \mathbf{x}$. To accelerate calculation, ChebyNet (Defferrard, Bresson, and Vandergheynst 2016) proposed a polynomial filter $g_\theta(\Lambda) = \sum_{k=0}^K \theta_k \Lambda^k$ and approximated $g_\theta(\Lambda)$ by a truncated expansion concerning Chebyshev polynomials $T_k(x)$:

$$g_{\theta'} * \mathbf{x} \approx \sum_{k=0}^K \theta'_k T_k(\tilde{L}) \mathbf{x}, \quad (6)$$

where $\tilde{L} = \frac{2}{\lambda_{\max}} L - I_n$ and λ_{\max} is the largest eigenvalue of Laplacian matrix L . $\theta' \in \mathbb{R}^K$ is now the parameter of Chebyshev polynomials $T_k(x)$. K denotes the K_{th} order polynomial in Laplacian. Due to the natural connection between Fourier transform and single processing, it's easy to formulate ChebyNet to *GF-Attack*:

Lemma 1. *The K -localized single-layer ChebyNet with activation function $\sigma(\cdot)$ and weight matrix Θ is equivalent to filter graph signal X with a polynomial filter $\mathcal{H} = \sum_{k=0}^K T_k(S)$ with graph-shift filter $S = 2\frac{L^{sym}}{\lambda_{max}} - I_n$. $T_k(S)$ represents Chebyshev polynomial of order k . Equation (2) can be rewritten as:*

$$\tilde{X} = \sum_{k=0}^K T_k\left(2\frac{L^{sym}}{\lambda_{max}} - I_n\right) X, \quad X' = \sigma(\tilde{X}\Theta).$$

Proof. The K -localized single-layer ChebyNet with activation function $\sigma(\cdot)$ is $\sigma(\sum_{k=0}^K \theta'_k T_k(2\frac{L^{sym}}{\lambda_{max}} - I_n) X)$. Thus, we can directly write graph-shift filter as $S = 2\frac{L^{sym}}{\lambda_{max}} - I_n$ and linear and shift-invariant filter $\mathcal{H} = \sum_{k=0}^K T_k(S)$. \square

GCN (Kipf and Welling 2017) constructed the layer-wise model by simplifying the ChebyNet with $K = 1$ and the *re-normalization trick* to avoid gradient exploding/vanishing:

$$X^{(l+1)} = \sigma\left(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} X^{(l)} \Theta^{(l)}\right), \quad (7)$$

where $\tilde{A} = A + I_n$ and $\tilde{D}_{ii} = \sum_j \tilde{A}_{ij}$. $\Theta = \{\theta_1^{(l)}, \dots, \theta_n^{(l)}\}$ is the parameters in the l_{th} layer and $\sigma(\cdot)$ is an activation function.

SGC (Wu et al. 2019) further utilized a single linear transformation to achieve computationally efficient graph convolution, i.e., $\sigma(\cdot)$ in SGC is a linear activation function. We can formulate the multi-layer SGC as *GF-Attack* through its theoretical connection to ChebyNet:

Corollary 2. *The K -layer SGC is equivalent to the K -localized single-layer ChebyNet with K_{th} order polynomials of the graph-shift filter $S^{sym} = 2I_n - L^{sym}$. Equation (2) can be rewritten as:*

$$\tilde{X} = (2I_n - L^{sym})^K X, \quad X' = \sigma(\tilde{X}\Theta).$$

Proof. We can write the K -layer SGC as $(2I_n - L^{sym})^K X \Theta$. Since Θ is the learned parameters by the neural network, we can employ the reparameterization trick to use $(2I_n - L^{sym})^K$ to approximate the same order polynomials $\sum_{k=0}^K T_k(2I_n - L^{sym})$ with new $\tilde{\Theta}$. Then we rewrite the K -layer SGC by polynomial expansion as $\sum_{k=0}^K T_k(2I_n - L^{sym}) X \tilde{\Theta}$. Therefore, we can directly write the graph-shift filter $S^{sym} = 2I_n - L^{sym}$ with the same linear and shift-invariant filter \mathcal{H} as K -localized single-layer ChebyNet. \square

Note that SGC and GCN are identical when $K = 1$. Even though non-linearity disturbs the explicit expression of graph-shift filter of multi-layer GCN, the spectral analysis from (Wu et al. 2019) demonstrated that both GCN and SGC share similar graph filtering behavior. Thus, we extend the general attack loss from multi-layer SGC to multi-layer GCN under non-linear activation functions scenario. Our experiments also validate that the attack model for multi-layer SGC also shows excellent performance on multi-layer GCN.

GF-Attack loss for SGC/GCN. As stated in Corollary 2, the graph-shift filter S of SGC/GCN is defined as $S^{sym} = 2I_n - L^{sym} = D^{-\frac{1}{2}} A D^{-\frac{1}{2}} + I_n = \hat{A} + I_n$, where \hat{A} denotes the normalized adjacent matrix. Thus, for K -layer SGC/GCN, we can decompose the graph filter \mathcal{H} as $\mathcal{H}^{sym} = (S^{sym})^K = U_{\hat{A}} (\Lambda_{\hat{A}} + I_n)^K U_{\hat{A}}^T$, where $\Lambda_{\hat{A}}$ and $U_{\hat{A}}$ are eigen-pairs of \hat{A} . The corresponding adversarial attack loss for K_{th} order SGC/GCN can be rewritten as:

$$\arg \max_{A'} \sum_{i=T+1}^n (\lambda'_{\hat{A}',i} + 1)^{2K} \cdot \sum_{i=T+1}^n \|\mathbf{u}_{\hat{A}',i}^T X\|_2^2, \quad (8)$$

where $\lambda'_{\hat{A}',i}$ refers to the i_{th} largest eigenvalue of the perturbed normalized adjacent matrix \hat{A}' .

While each time directly calculating $\lambda'_{\hat{A}',i}$ from attacked normalized adjacent matrix A' will need an eigen-decomposition operation, which is extremely time consuming, eigenvalue perturbation theory is introduced to estimate $\lambda'_{\hat{A}',i}$ in a linear time:

Theorem 3. *Let $A' = A + \Delta A$ be a perturbed version of A by adding/removing edges and ΔD be the respective change in the degree matrix. $\lambda_{\hat{A},i}$ and $\mathbf{u}_{\hat{A},i}$ are the i_{th} eigen-pair of eigenvalue and eigenvector of \hat{A} and also solve the*

generalized eigen-problem $\mathbf{A}\mathbf{u}_{\hat{A},i} = \lambda_{\hat{A},i}\mathbf{D}\mathbf{u}_{\hat{A},i}$. Then the perturbed generalized eigenvalue $\lambda'_{\hat{A},i}$ is approximately as:

$$\lambda'_{\hat{A},i} \approx \lambda_{\hat{A},i} + \frac{\mathbf{u}_{\hat{A},i}^T \Delta \mathbf{A} \mathbf{u}_{\hat{A},i} - \lambda_{\hat{A},i} \mathbf{u}_{\hat{A},i}^T \Delta \mathbf{D} \mathbf{u}_{\hat{A},i}}{\mathbf{u}_{\hat{A},i}^T \mathbf{D} \mathbf{u}_{\hat{A},i}}. \quad (9)$$

Proof. Please kindly refer to (Zhu et al. 2018). \square

With Theorem 3, we can directly derive the explicit formulation of $\lambda'_{\hat{A}}$, perturbed by ΔA on adjacent matrix A .

GF-Attack on Sampling-based Graph Embedding

Sampling-based graph embedding learns vertex representations according to sampled vertices, vertex sequences, or network motifs. For instance, LINE (Tang et al. 2015) with second order proximity intends to learn two graph representation matrices X', Y' by maximizing the NEG loss of the skip-gram model:

$$\mathcal{L} = \sum_{i=1}^{|\mathcal{V}|} \sum_{j=1}^{|\mathcal{V}|} A_{i,j} \left(\log \sigma(x_i^T y_j) + b \mathbb{E}_{j' \sim P_N} [\log \sigma(-x_i^T y_{j'})] \right), \quad (10)$$

where x'_i, y'_j are rows of X', Y' respectively; σ is the sigmoid function; b is the negative sampling parameter; P_N denotes the noise distribution generating negative samples. Meanwhile, DeepWalk (Perozzi, Al-Rfou, and Skiena 2014) adopts the similar loss function except that $A_{i,j}$ is replaced with an indicator function indicating whether vertices v_i and v_j are sampled in the same sequence within given context window size K .

From the perspective of sampling-based graph embedding models, the embedded matrix is obtained by generating training corpus for the skip-gram model from adjacent matrix or a set of random walks. (Yang and Liu 2015; Qiu et al. 2018) show that Point-wise Mutual Information (PMI) matrices are implicitly factorized in sampling-based embedding approaches. It indicates that LINE/DeepWalk can be rewritten into a matrix factorization form:

Lemma 4. (Qiu et al. 2018) *Given context window size K and number of negative sampling b in skip-gram, the result of DeepWalk in matrix form is equivalent to factorize matrix:*

$$M = \log \left(\frac{\text{vol}(G)}{bK} \left(\sum_{k=1}^K (D^{-1}A)^k \right) D^{-1} \right), \quad (11)$$

where $\text{vol}(G) = \sum_i \sum_j A_{ij} = \sum_i D_{ii}$ denotes the volume of graph G . And LINE can be viewed as the special case of DeepWalk with $K = 1$.

For proof of Lemma 4, please kindly refer to (Qiu et al. 2018). Inspired by this insight, we prove that LINE can be viewed from a GSP manner as well:

Theorem 5. *LINE is equivalent to filter a graph signal $X = \frac{1}{b}I_n$ with a polynomial filter \mathcal{H} and fixed parameters $\Theta = \text{vol}(G)D^{-1}$. $\mathcal{H} = S$ is constructed by graph-shift filter $S^{rw} = I_n - L^{rw}$. Equation (2) can be rewritten as:*

$$\tilde{X} = \frac{1}{b}(I_n - L^{rw})D^{-1}I_n, \quad X' = \log(\text{vol}(G)\tilde{X}).$$

Note that LINE is formulated from an optimized unsupervised NEG loss of skip-gram model. Thus, the parameter Θ and value of the NCG loss have been fixed at the optimal point of the model with given graph signals.

We can extend Theorem 5 to DeepWalk since LINE is a 1-window special case of DeepWalk:

Corollary 6. The output of K -window DeepWalk with b negative samples is equivalent to filtering a set of graph signals $X = \frac{1}{b}I_n$ with given parameters $\Theta = \text{vol}(G)D^{-1}$. Equation (2) can be rewritten as:

$$\tilde{X} = \frac{1}{bK} \sum_{k=1}^K (I_n - L^{rw})^k D^{-1}I_n, \quad X' = \log(\text{vol}(G)\tilde{X}).$$

Proof of Theorem 5 and Corollary 6. With Lemma 4, we can explicitly write DeepWalk as $\exp(M) = \frac{\text{vol}(G)}{b} \left(\sum_{k=1}^K \frac{1}{K} (I_n - L^{rw})^k D^{-1}I_n \right)$. Therefore, we can directly have the explicit expression of Equation (2) on LINE/DeepWalk. \square

GF-Attack loss for LINE/DeepWalk. As stated in Corollary 6, the graph-shift filter S of DeepWalk is defined as $S^{rw} = I_n - L^{rw} = D^{-1}A = D^{-\frac{1}{2}}\hat{A}D^{\frac{1}{2}}$. Therefore, graph filter \mathcal{H} of the K -window DeepWalk can be decomposed as $\mathcal{H}^{rw} = \frac{1}{K} \sum_{k=1}^K (S^{rw})^k$, which satisfies $\mathcal{H}^{rw}D^{-1} = D^{-\frac{1}{2}}U_{\hat{A}} \left(\frac{1}{K} \sum_{k=1}^K \Lambda_{\hat{A}}^k \right) U_{\hat{A}}^T D^{-\frac{1}{2}}$.

Since multiplying $D^{-\frac{1}{2}}$ in GF-Attack loss brings extra complexity, (Qiu et al. 2018) provides us a way to well approximate the perturbed $\lambda'_{\mathcal{H}^{rw}D^{-1}}$ without this term.

Inspired by (Qiu et al. 2018), we can find that both the magnitude of eigenvalues and smallest eigenvalue of $\mathcal{H}^{rw}D^{-1}$ are always well bounded. Thus we can approximate $\lambda'_{\mathcal{H}^{rw}D^{-1}} \approx \frac{1}{d_{\min}} \lambda'_{U_{\hat{A}} \left(\frac{1}{K} \sum_{k=1}^K \Lambda_{\hat{A}}^k \right) U_{\hat{A}}^T}$. Therefore, the corresponding adversarial attack loss of K_{th} order DeepWalk can be rewritten as:

$$\arg \max_{A'} \sum_{i=T+1}^n \left(\frac{1}{d_{\min}} \left| \frac{1}{K} \sum_{k=1}^K \lambda'_{A',i} \right| \right)^2 \cdot \sum_{i=T+1}^n \|\mathbf{u}_{A',i}^T X\|_2^2. \quad (12)$$

When $K = 1$, Equation (12) becomes the adversarial attack loss of LINE. Similarly, Theorem 3 is utilized to estimate $\lambda'_{A'}$ in the loss of LINE/DeepWalk.

The Attack Algorithm

Now the general attack loss is established, the goal of our adversarial attack is to misclassify a target vertex t from an attributed graph $G(\mathcal{V}, \mathcal{E})$ given a downstream node classification task. We start by defining the candidate flips then the general attack loss is responsible for scoring the candidates.

We first adopt the hierarchical strategy in (Dai et al. 2018) to decompose the single edge selection into two ends of this edge in practice. Then we let the candidate set \mathcal{C} for edge selection contains all vertices (edges and non-edges) directly accessory to the target vertex, i.e. $\mathcal{C} = \{(v, t) | v \neq t\}$, as (Dai et al. 2018; Bojchevski and Günnemann 2019). Intuitively,

Algorithm 1 Graph Filter Attack (GF-Attack) adversarial attack algorithm under RBA setting

Input:

Adjacent Matrix A ; feature matrix X ; target vertex t ;
 number of top- T smallest singular values/vectors selected T ; order of graph filter K ; fixed budget β .

Output:

Perturbed adjacent Matrix A' .

- 1: Initial the candidate flips set as $\mathcal{C} = \{(v, t) | v \neq t\}$, eigenvalue decomposition of $\hat{A} = U_{\hat{A}} \Lambda_{\hat{A}} U_{\hat{A}}^T$;
 - 2: **for** $(v, t) \in \mathcal{C}$ **do**
 - 3: Approximate $\Lambda'_{\hat{A}}$ resulting by removing/inserting edge (v, t) via Equation (9);
 - 4: Update $Score_{(v, t)}$ from loss Equation (8) or Equation (12);
 - 5: **end for**
 - 6: $\mathcal{C}_{sel} \leftarrow$ edge flips with top- β $Score$;
 - 7: $A' \leftarrow A \pm \mathcal{C}_{sel}$;
 - 8: **return** A'
-

further away the vertices from target t , less influence they impose on t . Meanwhile, experiments in (Zügner, Akbarnejad, and Günnemann 2018; Bojchevski and Günnemann 2019) also showed that they can do significantly more damage compared to candidate flips chosen from other parts of graph. Thus, our experiments are restricted on such choices.

Overall, for a given target vertex t , we establish the target attack by sequentially calculating the corresponding *GF-Attack* loss w.r.t graph-shift filter S for each flip in candidate set as scores. Then with a fixed budget β , the adversarial attack is accomplished by selecting flips with top- β scores as perturbations on the adjacent matrix A of clean graph. Details of the *GF-Attack* adversarial attack algorithm under RBA setting is in Algorithm 1.

Experiments

Datasets. We evaluate our approach on three real-world datasets: Cora, Citeseer and Pubmed. In all three citation network datasets, vertices are documents with corresponding bag-of-words features and edges are citation links. The data preprocessing settings are closely followed the benchmark setup in (Kipf and Welling 2017). Only the largest connected component (LCC) is considered to be consistent with (Zügner, Akbarnejad, and Günnemann 2018). For statistical overview of datasets, please kindly refer to (Zügner, Akbarnejad, and Günnemann 2018).

Baselines. In current literatures, few of studies strictly follow the restricted black-box attack setting. They utilize the additional information to help construct the attackers, such as labels (Zügner, Akbarnejad, and Günnemann 2018), gradients (Dai et al. 2018), etc.

Hence, we compare four baselines with the proposed attacker under RBA setting as follows:

- *Random* (Dai et al. 2018): for each perturbation, randomly choosing insertion or removing of an edge in graph G . We report averages over 10 different seeds to alleviate the

influence of randomness.

- *Degree* (Tong et al. 2012): for each perturbation, inserting or removing an edge based on degree centrality, which is equivalent to the sum of degrees in original graph G .
- *RL-S2V* (Dai et al. 2018): a reinforcement learning based attack method, which learns the generalizable attack policy for GCN under RBA scenario.
- \mathcal{A}_{class} (Bojchevski and Günnemann 2019): a matrix perturbation theory based black-box attack method designed for DeepWalk. Then \mathcal{A}_{class} evaluates the targeted attacks on node classification by learning a logistic regression.

Target Models. To validate the generalization ability of our proposed attacker, we choose four popular graph embedding models: GCN (Kipf and Welling 2017), SGC (Wu et al. 2019), DeepWalk (Perozzi, Al-Rfou, and Skiena 2014) and LINE (Tang et al. 2015) for evaluation. First two of them are Graph Convolutional Networks and the others are sampling-based graph embedding methods. For DeepWalk, the hyperparameters are set to commonly used values: window size as 5, number of negative sampling in skip-gram as 5 and top-128 largest singular values/vectors. A logistic regression classifier is connected to the output embeddings of sampling-based methods for classification. Unless otherwise stated, all Graph Convolutional Networks contain two layers.

Attack Configuration. A small budget β is applied to regulate all the attackers. To make this attacking task more challenging, β is set to 1. Specifically, the attacker is limited to only add/delete a single edge given a target vertex t . For our method, we set the parameter T in our general attack model as $n - T = 128$, which means that we choose the top- T smallest eigenvalues for T -rank approximation in embedding quality measure. Unless otherwise indicated, the order of graph filter in *GF-Attack* model is set to $K = 2$. Following the setting in (Zügner, Akbarnejad, and Günnemann 2018), we split the graph into labeled (20%) and unlabeled vertices (80%). Further, the labeled vertices are splitted into equal parts for training and validation. The labels and classifier is invisible to the attacker due to the RBA setting. The attack performance is evaluated by the decrease of node classification accuracy following (Dai et al. 2018).

Attack Performance Evaluation

In the section, we evaluate the overall attack performance of different attackers.

Attack on Graph Convolutional Networks. Table 1 summaries the attack results of different attackers on Graph Convolutional Networks. Our *GF-Attack* attacker outperforms other attackers on all datasets and all models. Moreover, *GF-Attack* performs quite well on 2 layers GCN with nonlinear activation. This implies the generalization ability of our attacker on Graph Convolutional Networks.

Attack on Sampling-based Graph Embedding. Table 1 also summaries the attack results of different attackers on sampling-based graph embedding models. As expected, our attacker achieves the best performance nearly on all target models. It validates the effectiveness of our method on attacking sampling-based models.

Another interesting observation is that the attack performance on LINE is much better than that on DeepWalk. This

Table 1: Summary of the change in classification accuracy (in percent) compared to the clean/original graph. Single edge perturbation under RBA setting. Lower is better.

Dataset	Cora				Citeseer				Pubmed			
	GCN	SGC	DeepWalk	LINE	GCN	SGC	DeepWalk	LINE	GCN	SGC	DeepWalk	LINE
Models (unattacked)	80.20	78.82	77.23	76.75	72.50	69.68	69.68	65.15	80.40	80.21	78.69	72.12
<i>Random</i>	-1.90	-1.22	-1.76	-1.84	-2.86	-1.47	-6.62	-1.78	-1.75	-1.77	-1.25	-1.01
<i>Degree</i>	-2.21	-4.42	-3.08	-12.40	-4.68	-5.21	-9.67	-12.55	-3.86	-4.44	-2.43	-13.05
<i>RL-S2V</i>	-5.20	-5.62	-5.24	-10.38	-6.50	-4.08	-12.13	-20.10	-6.40	-6.11	-6.10	-13.21
\mathcal{A}_{class}	-3.62	-2.96	-6.29	-7.55	-3.48	-2.83	-12.56	-10.28	-4.21	-2.25	-3.05	-6.75
<i>GF-Attack</i>	-7.06	-6.73	-5.31	-13.27	-7.78	-6.19	-12.50	-22.11	-7.96	-7.20	-7.43	-14.16




Figure 2: Comparison between order K of *GF-Attack* and number of layers in GCN/SGC on Citeseer.

Table 2: Running time (s) comparison over all baseline methods on Citeseer. We report the 10 times average running time of processing single node for each model.

Models	<i>Random</i>	<i>Degree</i>	<i>RL-S2V</i>	\mathcal{A}_{class}	<i>GF-Attack</i>
Citeseer	0.19	42.21	222.80	146.58	12.78

result may due to the deterministic structure of LINE, while the random sampling procedure in DeepWalk may help raise the resistance to adversarial attack. Moreover, *GF-Attack* on all graph filters successfully drop the classification accuracy on both Graph Convolutional Networks and sampling-based models, which again indicates the transferability of our general model in practice.

Evaluation of Multi-layer GCNs

To further inspect the transferability of our attacker, we conduct attack towards multi-layer Graph Convolutional Networks w.r.t the order of graph filter in *GF-Attack* model. Figure 2 presents the attacking results on 2, 3, 4 and 5 layers GCN and SGC with different orders, and the number followed by *GF-Attack* indicates the graph-shift filter order K in general attack loss. From Figure 2, we can observe that: first, the transferability of our general model is demonstrated, since all graph-shift filters in loss with different order K can perform the effective attack on all models. Interestingly, *GF-Attack-5* achieves the best attacking performance in most cases. It implies that the higher order filter contains higher order information and has positive effects on attack to simpler models. Second, the attacking performance on SGC is always better than GCN under all settings. We conjecture that the non-linearity between layers in GCN successively

adding robustness to GCN.

Evaluation under Multi-edge Perturbation Settings

In this section, we evaluate the performance of attackers with multi-edge perturbation, i.e. $\beta \geq 1$. The results of multi-edge perturbations on Cora under RBA setting are reported in Figure 3 for demonstration. Clearly, with increasing of the number of perturbed edges, the attacking performance gets better for each attacker. Our attacker outperforms other baselines on all cases. It validates that our general attacker can still perform well when fixed budget β becomes larger.




Figure 3: Multiple-edge attack results on Cora under RBA setting. Lower is better.

Computational Efficiency Analysis

In this section, we empirically evaluate the computational efficiency of our *GF-Attack*. The running time (s) comparison of 10 times average on Citeseer is demonstrated in Table 2. While being less efficient than two native baselines (*Random*

and *Degree*), our *GF-Attack* is much faster than the developed methods *RL-S2V* and \mathcal{A}_{class} . Joining the performance in Table 1, it reads that *GF-Attack* is not only effective in performance but also efficient computationally.

Conclusion

In this paper, we consider the adversarial attack on different kinds of graph embedding models under restrict black-box attack scenario. From graph signal processing of view, we try to formulate the graph embedding method as a general graph signal process with corresponding graph filter and construct a restricted adversarial attacker which aims to attack the graph filter only by the adjacency matrix and feature matrix. Thereby, a general optimization problem is constructed by measuring embedding quality and an effective algorithm is derived accordingly to solve it. Experiments show the vulnerability of different kinds of graph embedding models to our attack framework.

Acknowledgements

This work is supported by National Natural Science Foundation of China Major Project No. U1611461 and National Program on Key Basic Research Project No. 2015CB352300. We would like to thank the anonymous reviewers for the helpful comments. We also thank Daniel Zügner from Technical University of Munich for the valuable suggestions and discussions.

References

Akhtar, N., and Mian, A. S. 2018. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access* 6:14410–14430.

Bojchevski, A., and Günnemann, S. 2019. Adversarial attacks on node embeddings via graph poisoning. In *International Conference on Machine Learning*, 695–704.

Cui, P.; Wang, X.; Pei, J.; and Zhu, W. 2018. A survey on network embedding. *IEEE Transactions on Knowledge and Data Engineering* 31(5):833–852.

Dai, H.; Li, H.; Tian, T.; Huang, X.; Wang, L.; Zhu, J.; and Song, L. 2018. Adversarial attack on graph structured data. *international conference on machine learning* 1115–1124.

Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. *neural information processing systems* 3844–3852.

Duvenaud, D. K.; Maclaurin, D.; Aguilera-Iparraguirre, J.; Gómez-Bombarelli, R.; Hirzel, T.; Aspuru-Guzik, A.; and Adams, R. P. 2015. Convolutional networks on graphs for learning molecular fingerprints. *neural information processing systems* 2224–2232.

Hamilton, W. L.; Ying, Z.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *neural information processing systems* 1024–1034.

Kipf, T. N., and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. *international conference on learning representations*.

Nar, K.; Ocal, O.; Sastry, S. S.; and Ramchandran, K. 2019. Cross-entropy loss and low-rank features have responsibility for adversarial examples. *arXiv preprint arXiv:1901.08360*.

Ortega, A.; Frossard, P.; Kovačević, J.; Moura, J. M.; and Vandergheynst, P. 2018. Graph signal processing: Overview, challenges, and applications. *Proceedings of the IEEE* 106(5):808–828.

Paranjape, A.; Benson, A. R.; and Leskovec, J. 2017. Motifs in temporal networks. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, 601–610.

Perozzi, B.; Al-Rfou, R.; and Skiena, S. 2014. Deepwalk: Online learning of social representations. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, 701–710. ACM.

Qiu, J.; Dong, Y.; Ma, H.; Li, J.; Wang, K.; and Tang, J. 2018. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, 459–467.

Scarselli, F.; Gori, M.; Tsoi, A. C.; Hagenbuchner, M.; and Monfardini, G. 2009. The graph neural network model. *IEEE Transactions on Neural Networks* 20(1):61–80.

Shuman, D. I.; Narang, S. K.; Frossard, P.; Ortega, A.; and Vandergheynst, P. 2013. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine* 30(3):83–98.

Sun, L.; Wang, J.; Yu, P. S.; and Li, B. 2018. Adversarial attack and defense on graph data: A survey. *arXiv preprint arXiv:1812.10528*.

Tang, J.; Qu, M.; Wang, M.; Zhang, M.; Yan, J.; and Mei, Q. 2015. Line: Large-scale information network embedding. In *Proceedings of the 24th International Conference on World Wide Web*, 1067–1077.

Tong, H.; Prakash, B. A.; Eliassi-Rad, T.; Faloutsos, M.; and Faloutsos, C. 2012. Gelling, and melting, large graphs by edge manipulation. In *Proceedings of the 21st ACM international conference on Information and knowledge management*, 245–254.

Tu, K.; Cui, P.; Wang, X.; Yu, P. S.; and Zhu, W. 2018. Deep recursive network embedding with regular equivalence. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2357–2366. ACM.

Wu, F.; Zhang, T.; Souza Jr., A. H.; Fifty, C.; Yu, T.; and Weinberger, K. Q. 2019. Simplifying graph convolutional networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks. *arXiv preprint arXiv:1810.00826*.

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. *arXiv preprint arXiv:1906.04214*.

Yang, C., and Liu, Z. 2015. Comprehend deepwalk as matrix factorization. *arXiv preprint arXiv:1501.00358*.

Yang, C.; Liu, Z.; Zhao, D.; Sun, M.; and Chang, E. 2015. Network representation learning with rich text information. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

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.

Zügner, D., and Günnemann, S. 2019. Adversarial attacks on graph neural networks via meta learning. *international conference on learning representations*.

Zhu, D.; Cui, P.; Zhang, Z.; Pei, J.; and Zhu, W. 2018. High-order proximity preserved embedding for dynamic networks. *IEEE Transactions on Knowledge and Data Engineering* 30:2134–2144.