

# Adversarial Training Methods for Network Embedding

Quanyu Dai

The Hong Kong Polytechnic University  
Kowloon, Hong Kong  
dqym100@hotmail.com

Xiao Shen

The Hong Kong Polytechnic University  
Kowloon, Hong Kong  
shenxiaocam@163.com

Liang Zhang

JD.com  
Beijing, China  
zhangliang16@jd.com

Qiang Li

Y-tech, Kwai  
Beijing, China  
leetsiang.cloud@gmail.com

Dan Wang

The Hong Kong Polytechnic University  
Kowloon, Hong Kong  
csdwang@comp.polyu.edu.hk

## ABSTRACT

Network Embedding is the task of learning continuous node representations for networks, which has been shown effective in a variety of tasks such as link prediction and node classification. Most of existing works aim to preserve different network structures and properties in low-dimensional embedding vectors, while neglecting the existence of noisy information in many real-world networks and the overfitting issue in the embedding learning process. Most recently, generative adversarial networks (GANs) based regularization methods are exploited to regularize embedding learning process, which can encourage a global smoothness of embedding vectors. These methods have very complicated architecture and suffer from the well-recognized non-convergence problem of GANs. In this paper, we aim to introduce a more succinct and effective local regularization method, namely adversarial training, to network embedding so as to achieve model robustness and better generalization performance. Firstly, the adversarial training method is applied by defining adversarial perturbations in the embedding space with an adaptive  $L_2$  norm constraint that depends on the connectivity pattern of node pairs. Though effective as a regularizer, it suffers from the interpretability issue which may hinder its application in certain real-world scenarios. To improve this strategy, we further propose an interpretable adversarial training method by enforcing the reconstruction of the adversarial examples in the discrete graph domain. These two regularization methods can be applied to many existing embedding models, and we take DEEPWALK as the base model for illustration in the paper. Empirical evaluations in both link prediction and node classification demonstrate the effectiveness of the proposed methods.

## KEYWORDS

Network Embedding, Adversarial Training, Robustness

### ACM Reference Format:

Quanyu Dai, Xiao Shen, Liang Zhang, Qiang Li, and Dan Wang. 2019. Adversarial Training Methods for Network Embedding. In *Proceedings of the 2019 World Wide Web Conference (WWW '19)*, May 13–17, 2019, San Francisco, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3308558.3313445>

This paper is published under the Creative Commons Attribution 4.0 International (CC-BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW '19, May 13–17, 2019, San Francisco, CA, USA

© 2019 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC-BY 4.0 License.

ACM ISBN 978-1-4503-6674-8/19/05.

<https://doi.org/10.1145/3308558.3313445>

## 1 INTRODUCTION

Network embedding strategies, as an effective way for extracting features from graph structured data automatically, have gained increasing attention in both academia and industry in recent years. The learned node representations from embedding methods can be utilized to facilitate a wide range of downstream learning tasks, including some traditional network analysis tasks such as link prediction and node classification, and many important applications in industry such as product recommendation in e-commerce website and advertisement distribution in social networks. Therefore, under such great application interest, substantial efforts have been devoted to designing effective and scalable network embedding models.

Most of the existing works focus on preserving network structures and properties in low-dimensional embedding vectors [4, 35, 37]. Firstly, DEEPWALK [35] defines random walk based neighborhood for capturing node dependencies, and node2vec [13] extends it with more flexibility in balancing local and global structural properties. LINE [35] preserves both first-order and second-order proximities through considering existing connection information. Further, GraRep [4] manages to learn different high-order proximities based on different  $k$ -step transition probability matrix. Aside from the above mentioned structure-preserving methods, several research works investigate the learning of property-aware network embeddings. For example, network transitivity, as the driving force of link formation, is considered in [26], and node popularity, as another important factor affecting link generation, is incorporated into RaRE [14] to learn social-rank aware and proximity-preserving embedding vectors. However, the existence of noisy information in real-world networks and the overfitting issue in the embedding learning process are neglected in most of these methods, which leaves the necessity and potential improvement space for further exploration.

Most recently, adversarial learning regularization method is exploited for improving model robustness and generalization performance in network embedding [8, 41]. ANE [8] is the first try in this direction, which imposes a prior distribution on embedding vectors through adversarial learning. Then, the adversarially regularized autoencoder is adopted in NETRA [41] to overcome the mode-collapse problem in ANE method. These two methods both encourage the global smoothness of the embedding distribution based on generative adversarial networks (GANs) [11]. Thus, they have very complicated frameworks and suffer from the well-recognized hard training problems of GANs [3, 29].

In this paper, we aim to leverage the adversarial training (AdvT) method [12, 34] for network embedding to achieve model robustness and better generalization ability. AdvT is a local smoothness regularization method with more succinct architecture. Specifically, it forces the learned classifier to be robust to adversarial examples generated from clean ones with small crafted perturbation [34]. Such designed noise with respect to each input example is dynamically obtained through finding the direction to maximize model loss based on current model parameters, and can be approximately computed with fast gradient method [12]. It has been demonstrated to be extremely useful for some classification problems [12, 23].

However, how to adapt AdvT for graph representation learning remains an open problem. It is not clear how to generate adversarial examples in the discrete graph domain since the original method is designed for continuous inputs. In this paper, we propose an adversarial training DEEPWALK model, which defines the adversarial examples in the embedding space instead of the original discrete relations and obtains adversarial perturbation with fast gradient method. We also leverage the dependencies among nodes based on connectivity patterns in the graph to design perturbations with different  $L_2$  norm constraints, which enables more reasonable adversarial regularization. The training process can be formulated as a two-player game, where the adversarial perturbations are generated to maximize the model loss while the embedding vectors are optimized against such designed noises with stochastic gradient descent method. Although effective as a regularization technique, directly generating adversarial perturbation in embedding space with fast gradient method suffers from interpretability issue, which may restrict its application areas. Further, we manage to restore the interpretability of adversarial examples by constraining the perturbation directions to embedding vectors of other nodes, such that the adversarial examples can be considered as the substitution of nodes in the original discrete graph domain.

Empirical evaluations show the effectiveness of both adversarial and interpretable adversarial training regularization methods by building network embedding method upon DEEPWALK. It is worth mentioning that the proposed regularization methods, as a principle, can also be applied to other embedding models with embedding vectors as model parameters such as node2vec and LINE. The main contributions of this paper can be summarized as follows:

- We introduce a novel, succinct and effective regularization technique, namely adversarial training method, for network embedding models which can improve both model robustness and generalization ability.
- We leverage the dependencies among node pairs based on network topology to design perturbations with different  $L_2$  norm constraints for different positive target-context pairs, which enables more flexible and effective adversarial training regularization.
- We also equip the adversarial training method with interpretability for discrete graph data by restricting the perturbation directions to embedding vectors of other nodes, while maintaining its usefulness in link prediction and only slightly sacrificing its regularization ability in node classification.
- We conduct extensive experiments to evaluate the effectiveness of the proposed methods.

## 2 BACKGROUND

### 2.1 Framework of Network Embedding

The purpose of network embedding is to transform discrete network structure information into compact embedding vectors, which can be further used to facilitate downstream learning tasks, such as node classification and link prediction. The research problem can be formally formulated as follows: Given a weighted (unweighted) directed (undirected) graph  $G(V, E, A)$  ( $N = |V|$ ), with  $V = \{v_i\}_{i=1}^N$  as the node set,  $E = \{e_{ij}\}_{i,j=1}^N$  as the edge set, and  $A$  as the weighted adjacency matrix with  $A_{ij}$  quantifying the strength of the relationship between node  $v_i$  and  $v_j$ , network embedding is aimed at learning a mapping function  $f : V \mapsto U$ , where  $U \in R^{N \times d}$  ( $d \ll N$ ) is the embedding matrix with the  $i$ th row  $\mathbf{u}_i^T$  as the embedding vector of node  $v_i$ . Note that for many network embedding models, a context embedding matrix  $U'$  will also be learned. For these methods, embedding matrix  $U$  is also called target embedding matrix.

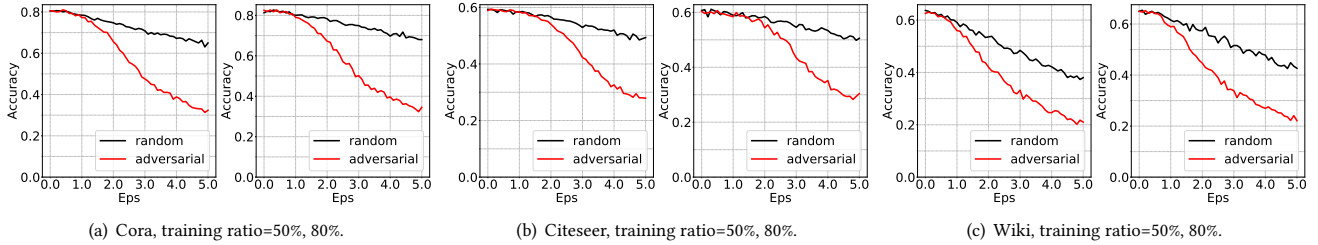
The learning framework of many famous network embedding methods, such as DEEPWALK [27], LINE [35] and node2vec [13], can be summarized into two phases: a sampling phase that determines node pairs with strong relationships, and an optimization phase that tries to preserve pairwise relationships in the embedding vectors through the negative sampling approach [22]. In particular, in the first phase, these three methods capture structural information by defining different neighborhood structures, such as random walk explored neighborhood in [13, 27], first-order and second-order proximities in [35]. We denote the generalized neighbors (not restricted to directly connected nodes) of node  $v_i$  as  $\mathcal{N}(v_i)$ , i.e., nodes in this set are closely related with  $v_i$  and should be close with  $v_i$  in the embedding space. The loss function of this framework can be abstracted as follows:

$$\begin{aligned} \mathcal{L}(G|\Theta) = & - \sum_{v_i \in V} \sum_{v_j \in \mathcal{N}(v_i)} \{\log \sigma(s(v_i, v_j|\Theta))\} \\ & + \sum_{k=1}^K \mathbb{E}_{v_k \sim P_k(v)} [\log \sigma(-s(v_i, v_k|\Theta))], \end{aligned} \quad (1)$$

where  $\Theta$  represents model parameters such as target and context embedding matrices,  $s(v_i, v_j|\Theta)$  represents the similarity score of node  $v_i$  and  $v_j$  based on model parameters  $\Theta$ , and  $\sigma(\cdot)$  is the sigmoid function.  $P_k(v)$  denotes the distribution for sampling negative nodes, and a simple variant of unigram distribution is usually utilized, i.e.,  $P_k(v) \propto d_v^{3/4}$ , where  $d_v$  is the out-degree of node  $v$ . Eq. (1) is actually a cross entropy loss with closely related node pair  $(v_i, v_j)$  as positive samples and  $(v_i, v_k)$  as negative samples, and thus network embedding can be considered as a classification problem.

### 2.2 Adversarial Training

Adversarial training [12, 34] is a newly proposed effective regularization method for classifiers which can not only improve the robustness of the model against adversarial attacks, but also achieve better generalization performance on learning tasks. It augments the original clean data with the dynamically generated adversarial examples, and then trains the model with the newly mixed examples. Denote the input as  $\mathbf{x}$  and model parameters as  $\theta$ . The loss on adversarial examples can be considered as a regularization term in



**Figure 1: Impact of applying adversarial and random perturbations to the embedding vectors learned by DEEPWALK on Cora, Citeseer and Wiki on multi-class classification with training ratio as 50% and 80%. Note that "random" represents random perturbations (noises generated from a normal distribution), while "adversarial" represents adversarial perturbations.**

the trained classifier  $p(y|\cdot)$ , which is as follows:

$$-\log p(y|\mathbf{x} + \mathbf{n}_{adv}; \theta), \quad (2)$$

$$\text{where } \mathbf{n}_{adv} = \arg \min_{\mathbf{n}, \|\mathbf{n}\| \leq \epsilon} \log p(y|\mathbf{x} + \mathbf{n}; \hat{\theta}), \quad (3)$$

where  $\mathbf{n}$  is the perturbation on the input,  $\epsilon$  represents the norm constraint of  $\mathbf{n}$ , and  $\hat{\theta}$  are current model parameters but fixed as constants. We employ  $L_2$  norm in this paper, while  $L_1$  norm has also been used in the literature [12]. Eq. (2) means that the model should be robust on the adversarial perturbed examples. Before each batch training, the adversarial noise  $\mathbf{n}$  with respect to the input  $\mathbf{x}$  is firstly generated by solving optimization problem (3) to make it resistant to current model. Since it is difficult to calculate Eq.(3) exactly in general, fast gradient descent method [12] is widely used to obtain the adversarial noise approximately by linearizing  $\log p(y|\mathbf{x}; \theta)$  around  $\mathbf{x}$ . Specifically, the adversarial perturbation with  $L_2$  norm constraint can be obtained as follows:

$$\mathbf{n}_{adv} = -\epsilon \cdot \frac{\mathbf{g}}{\|\mathbf{g}\|_2} \text{ where } \mathbf{g} = \nabla_{\mathbf{x}} \log p(y|\mathbf{x}; \hat{\theta}). \quad (4)$$

It can be easily calculated with backpropagation method.

### 2.3 Motivation

To improve the generalization ability of network embedding models, two ways have been used: firstly, some denoising autoencoder based methods [5, 8] improve model robustness by adding random perturbation to input data or hidden layers of deep models; secondly, some existing methods [8, 41] regularize embedding vectors from a global perspective through GAN-based method, i.e., encouraging the global smoothness of the distribution of embeddings. In this paper, we aim to introduce a novel, more succinct and effective regularization method for network embedding models, i.e., adversarial training (AdvT) [12]. AdvT generates crafted adversarial perturbations to model inputs and encourages local smoothness for improving model robustness and generalization performance, which can be expected to be more effective than the random perturbation methods [5] and global regularization methods [8, 41]. In the following, we would like to compare the impact of adversarial and random perturbation on embedding vectors to better motivate this new regularization method.

However, it is not clear how to integrate adversarial training into existing network embedding methods. Graph data is discrete, and the continuous adversarial noise can not be directly imposed on the discrete connected information. To bypass this difficulty, we

seek to define the adversarial perturbation on embedding vectors instead of the discrete graph domain as inspired by [23]. We define the adversarial perturbation on node embeddings as follows:

$$\mathbf{n}_{adv} = \arg \max_{\mathbf{n}, \|\mathbf{n}\| \leq \epsilon} \mathcal{L}(G|\hat{\Theta} + \mathbf{n}), \quad (5)$$

which can be further approximated with fast gradient method as presented in Eq. (4).

Take DEEPWALK [27] with negative sampling loss as an illustrative example. We explore the effect of adversarial perturbations on embedding vectors by adding them to the learned embedding vectors from DEEPWALK, and then perform multi-class classification with the perturbed embeddings on several datasets. Besides, we choose random perturbations as the compared baseline, i.e., noises generated from a normal distribution. Figure 1 displays node classification results with varying  $L_2$  norm constraints on the perturbations. We can find that embedding vectors are much more vulnerable to adversarial perturbations than random ones. For example, when  $\epsilon$  is set to 2.0, the performance of node classification with training ratio as 80% on Cora drops 3.35% under random perturbation, while that decreases 16.25% under adversarial perturbation which is around 4 times more serious. If the embedding vectors can be trained to be more robust on adversarial noises, we can expect more significant improvements in generalization performance.

## 3 PROPOSED METHODS

In this section, we first describe the adapted adversarial training method for network embedding models, and present the algorithm based on DEEPWALK. Then, we will tackle its interpretability issue by designing a new adversarial perturbation generation method.

### 3.1 Adversarial Training DEEPWALK

Figure 2 shows the framework of DEEPWALK with adversarial training regularization. It consists of two phases: a sampling phase that determines node pairs with strong relationships, and an optimization phase that tries to preserve pairwise relationships in the embedding vectors based on negative sampling approach. Note that in this paper we take DEEPWALK as the illustrative example, and the proposed framework can be applied to the network embedding methods, such as LINE and node2vec, with the main difference in the sampling phase only.

In the first phase, DEEPWALK transforms the network into node sequences by truncated random walk. For each node  $v_i \in V$ ,  $\eta$

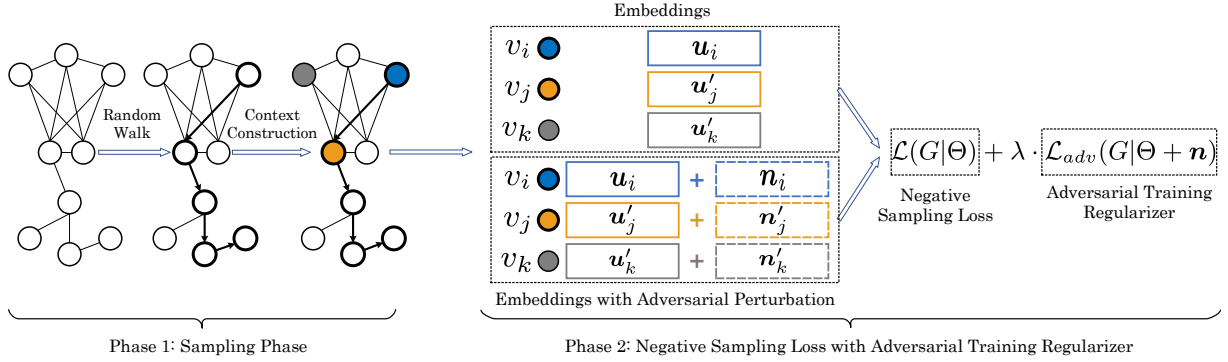


Figure 2: DEEPWALK with Adversarial Training Regularization

sequences each with  $l$  nodes will be randomly sampled based on network structure with  $v_i$  as the starting point. In every walking step, the next node  $v_j$  will be sampled from the neighbors of current node  $v_k$  with the probability proportional to the edge strength  $A_{kj}$  between  $v_k$  and  $v_j$ . In practice, the alias table method [19] is usually leveraged for node sampling given the weight distribution of neighbors of current node, which only takes  $O(1)$  time in a single sampling step. Then in the context construction process, closely related node pairs will be determined based on the sampled node sequences. Denote a node sequence as  $S$  with the  $i$ th node as  $s_i$ . The positive target-context pairs from  $S$  is defined as  $\mathcal{P} = \{(s_i, s_j) : |i - j| \leq c\}$ , where  $c$  represents the window size. With the constructed node pairs, the negative sampling loss will be optimized, which is defined as follows:

$$\begin{aligned} \mathcal{L}(G|\Theta) = & - \sum_{v_i \in V} \sum_{v_j \in \mathcal{N}(v_i)} \{\log \sigma(\mathbf{u}'_j{}^T \cdot \mathbf{u}_i)\} \\ & + \sum_{k=1}^K \mathbb{E}_{v_k \sim P_k(v)} [\log \sigma(-\mathbf{u}'_k{}^T \cdot \mathbf{u}_i)], \end{aligned} \quad (6)$$

where  $(v_i, v_j)$  is from the constructed positive target-context pairs, and  $\mathbf{u}_i$  and  $\mathbf{u}'_j$  are the target embedding of node  $v_i$  and context embedding of node  $v_j$  respectively.

For the adversarial version of DEEPWALK, an adversarial training regularization term is added to the original loss to help learn robust node representations against adversarial perturbations. The regularization term shares the same set of model parameters with the original model, but with the perturbed target and context embeddings as input. Existing methods consider the input examples independently, and impose the unique  $L_2$  norm constraint on all adversarial perturbations [12, 23]. For graph structured data, entities often correlate with each other in a very complicated way, so it is inappropriate to treat all positive target-context relations equally without discrimination. Adversarial regularization helps alleviate overfitting issue, but it may also bring in some noises that can hinder the preservation of structural proximities, i.e., adding noises to those inherently closely-related node pairs will prevent them from having similar embeddings. Thus, we take advantages of the dependencies among nodes to assign different  $L_2$  norm constraints to different positive target-context relations adaptively. Specifically, the more closely two nodes are connected, the smaller the constraint should be. The intuition is that less noises should be added to those node pairs which are inherently strongly-connected in the

original network, thus they can be pushed closer in the embedding space with high flexibility, while for those weakly-connected pairs larger constraint can help alleviate the overfitting issue.

We obtain the similarity score of two nodes through computing the shifted positive pointwise mutual information matrix [18]:

$$M_{ij} = \max\{\log(\frac{\hat{M}_{ij}}{\sum_k \hat{M}_{kj}}) - \log(\beta), 0\}, \quad (7)$$

where  $\hat{M} = \hat{A} + \hat{A}^2 + \dots + \hat{A}^t$  captures different high-order proximities,  $\hat{A}$  is the 1-step probability transition matrix obtained from  $A$  after the row-wise normalization, and  $\beta$  is a shift factor. We set  $t$  to 4 and  $\beta$  to  $\frac{1}{N}$  in the experiments. Then, the adaptive scale factor for the  $L_2$  norm constraint of the target-context pair  $v_i$  and  $v_j$  is calculated as below:

$$\Phi_{ij} = 1 - M_{ij} / \max\{M\}, \quad (8)$$

where  $\max\{M\}$  represents the maximum entry of matrix  $M$ . Since  $M_{ij} > 0$  ( $\forall i, j$ ),  $\Phi_{ij} \in [0, 1]$ . For those strongly-connected target-context pairs, the adaptive scale factor can help scale down the  $L_2$  norm constraint of the adversarial perturbation, and thus alleviate the negative effect from the noises.

Then, the adversarial training regularizer with scale factor for  $L_2$  norm constraint is defined as follows:

$$\begin{aligned} \mathcal{L}_{adv}(G|\Theta + \mathbf{n}_{adv}) = & - \sum_{v_i \in V} \sum_{v_j \in \mathcal{N}(v_i)} \{\log \sigma((\mathbf{u}'_j + \Phi_{ij} \cdot (\mathbf{n}'_j)_{adv})^T \cdot (\mathbf{u}_i + \Phi_{ij} \cdot (\mathbf{n}_i)_{adv}))\} \\ & + \sum_{k=1}^K \mathbb{E}_{v_k \sim P_k(v)} [\log \sigma(-(\mathbf{u}'_k + (\mathbf{n}'_k)_{adv})^T \cdot (\mathbf{u}_i + (\mathbf{n}_i)_{adv}))], \end{aligned} \quad (9)$$

where  $(\mathbf{n}_i)_{adv}$  and  $(\mathbf{n}'_j)_{adv}$  represent the original adversarial perturbation for target embedding of node  $v_i$  and context embedding of node  $v_j$  respectively.

Finally, one key problem is how to compute the adversarial perturbation for the given embedding vector of a node  $v$ . Here, we follow the famous adversarial training method directly [12, 34], and generate the perturbation noises to maximize model loss under current model parameters. The adversarial perturbation for node  $v$  is defined as follows:

$$\mathbf{n}_{adv} = \arg \max_{\mathbf{n}, \|\mathbf{n}\| \leq \epsilon} \mathcal{L}(G|\hat{\Theta} + \mathbf{n}). \quad (10)$$

---

**Algorithm 1:** The adversarial training DEEPWALK

---

**Input** : graph  $G(V, E, A)$ , window size  $c$ , embedding size  $d$ , walks per node  $\eta$ , negative size  $K$ , walk length  $l$ , adversarial noise level  $\epsilon$ , adversarial regularization strength  $\lambda$ , batch size  $b$

**Output**: Embedding matrix  $U$

- 1 Initialize target and context embeddings with DEEPWALK;
- 2 **while** *not converge* **do**
- 3     Generate a set of positive target-context pairs  $\mathcal{P}$  with random walk based method;
- 4     **repeat**
- 5         Sample a batch  $\mathcal{B}$  of target-context pairs from  $\mathcal{P}$ ;
- 6         // Generate adversarial perturbations
- 7          $\mathbf{n}_{adv} = \epsilon \frac{\mathbf{g}}{\|\mathbf{g}\|_2}$  where  $\mathbf{g} = \nabla_{\mathbf{u}} \mathcal{L}(G|\hat{\Theta})$  for each node  $v$  in the batch;
- 8         // Optimize model parameters
- 9         Update target and context embeddings by applying gradient descent technique to Eq. (12) ;
- 10     **until**  $[\mathcal{P}|/b]_+$  *times*;
- 11 **end**

---

It can be further approximated with fast gradient method as follows:

$$\mathbf{n}_{adv} = \epsilon \frac{\mathbf{g}}{\|\mathbf{g}\|_2} \text{ where } \mathbf{g} = \nabla_{\mathbf{u}} \mathcal{L}(G|\hat{\Theta}). \quad (11)$$

Therefore, the overall loss for the proposed adversarial training DEEPWALK is defined as follows:

$$\mathcal{L}(G|\Theta) + \lambda \cdot \mathcal{L}_{adv}(G|\Theta + \mathbf{n}_{adv}), \quad (12)$$

where  $\lambda$  is a hyperparameter to control the importance of the regularization term.

In this paper, we utilize DEEPWALK with negative sampling loss as the base model for building the adversarial version of network embedding methods. Since the original implementation is based on the well encapsulated library, which lacks flexibility for further adaption, we re-implement the model with tensorflow [1] and utilize a slightly different training strategy. Specifically, in each training epoch, we independently construct positive target-context pairs with random walk based method, and then optimize model parameters with mini-batch stochastic gradient descent technique. Algorithm 1 summarizes the training procedure for the adversarial training DEEPWALK. The model parameters are firstly initialized by training DEEPWALK with the method introduced above. For each batch training, adversarial perturbations are generated with fast gradient method for each node in the batch as presented in Line 7. Then, the target and context embeddings will be updated by optimizing the negative sampling loss with adversarial training regularization as shown in Line 9. Asynchronous version of stochastic gradient descent [25] can be utilized to accelerate the training as DEEPWALK. Note that we ignore the derivative of  $\mathbf{n}_{adv}$  with respect to model parameters. The adversarial perturbations can be computed with simple back-propagation method, which enjoys low computational cost. Thus, the adversarial training DEEPWALK is scalable as the base model.

## 3.2 Interpretable Adversarial Training DEEPWALK

Adversarial examples refer to examples that are generated by adding viciously designed perturbations with norm constraint to the clean input, which can significantly increase model loss and probably induce prediction error [34]. Take an example from [12] for illustration, a ‘‘panda’’ image with imperceptibly small adversarial perturbation is assigned to be ‘‘gibbon’’ by the well-trained classification model with high confidence, while the original image can be correctly classified. Such adversarial examples can be well interpreted since the perturbations are imposed on the input space. For the adversarial training DEEPWALK, adversarial perturbations are added to node embeddings instead of the discrete nodes and connected relations, and thus can not be easily reconstructed in the discrete graph domain. Though effective as a regularizer for improving model generalization performance, it suffers from lack of interpretability, which may create a barrier for its adoption in some real-world applications.

In this section, we propose an interpretable adversarial DEEPWALK model by restoring the interpretability of adversarial perturbations. Instead of pursuing the worst perturbation direction only, we restrict the direction of perturbations toward a subset of nodes in the graph in the embedding space, such as the neighbors of the considered node. In this way, the adversarial perturbations in the node embedding space might be interpreted as the substitution of nodes in the original input space, i.e., the discrete target-context relations. However, there might be a certain level of sacrifice on the regularization performance because of the restriction on perturbation directions.

The *direction vector* from node  $v_t$  to  $v_k$  in the embedding space is defined as follows:

$$\mathbf{v}_k^{(t)} = \frac{\tilde{\mathbf{v}}_k^{(t)}}{\|\tilde{\mathbf{v}}_k^{(t)}\|_2}, \text{ where } \tilde{\mathbf{v}}_k^{(t)} = \mathbf{u}_k - \mathbf{u}_t. \quad (13)$$

Denote  $V^{(t)} \subseteq V$  ( $|V^{(t)}| = T$ ,  $|V^{(t)}| \ll |V|$ ) as a set of nodes for generating adversarial perturbation for node  $v_t$ . We define  $V^{(t)}$  as the top  $T$  nearest neighbors of node  $v_t$  in the embedding space based on current model parameters. To improve model efficiency, we can also obtain  $V^{(t)}$  based on the pretrained model parameters, and fix it for all training epochs. We use the latter strategy for experiments in this paper. Denote  $\mathbf{w}^{(t)} \in \mathcal{R}^T$  as the weight vector for node  $v_t$  with  $w_k^{(t)}$  representing the weight associated with direction vector  $\mathbf{v}_k^{(t)}$ , where  $v_k$  is the  $k$ th node in  $V^{(t)}$ . The interpretable perturbation for  $v_t$  is defined as the weighted sum of the direction vectors starting from  $v_t$  and ending with nodes in  $V^{(t)}$ :

$$\mathbf{n}(\mathbf{w}^{(t)}) = \sum_{k=1}^T w_k^{(t)} \mathbf{v}_k^{(t)}, \quad v_k \in V^{(t)}, \quad \forall k = 1, \dots, T. \quad (14)$$

The adversarial perturbation is obtained by finding the weights that can maximize the model loss:

$$\mathbf{w}_{iAdv}^{(t)} = \arg \max_{\mathbf{w}^{(t)}, \|\mathbf{w}^{(t)}\| \leq \epsilon} \mathcal{L}_{iAdv}(G|\Theta + \mathbf{n}(\mathbf{w}^{(t)})), \quad (15)$$

where  $\mathcal{L}_{iAdv}$  is obtained by replacing  $\mathbf{n}_{adv}$  in Eq. (9) with  $\mathbf{n}(\mathbf{w}^{(t)})$ . In consideration of model efficiency, the above regularization term is approximated with first-order Taylor series for easy computation as

**Table 1: Statistics of benchmark datasets**

Name	Cora	Citeseer	Wiki	CA-GrQc	CA-HepTh
$ V $	2,708	3,264	2,363	5,242	9,877
$ E $	5,278	4,551	11,596	14,484	25,973
Avg. degree	1.95	1.39	4.91	2.76	2.63
#Labels	7	6	17	-	-

in [12]. Thus, the weights for constructing interpretable adversarial perturbation for node  $v_t$  can be computed as follows:

$$\mathbf{w}_{iAdv}^{(t)} = \epsilon \frac{\mathbf{g}}{\|\mathbf{g}\|_2} \text{ where } \mathbf{g} = \nabla_{\mathbf{w}^{(t)}} \mathcal{L}_{iAdv}(G|\Theta + \mathbf{n}(\mathbf{w}^{(t)})). \quad (16)$$

Substituting  $\mathbf{w}_{iAdv}^{(t)}$  into Eq. (14), we can get the adversarial perturbations  $\mathbf{n}(\mathbf{w}_{iAdv}^{(t)})$ . Further, by replacing  $\mathbf{n}_{adv}$  with  $\mathbf{n}(\mathbf{w}_{iAdv}^{(t)})$  in Eq. (9), we can have the interpretable adversarial training regularizer for DEEPWALK. The algorithm for interpretable adversarial training DEEPWALK is different from Algorithm 1 in the way of generating adversarial perturbations, and thus we do not present it in this paper due to space limitation. Since  $|V^{(t)}| \ll |V|$ , the computation of adversarial perturbation for one node takes constant time. Therefore, the time complexity of this model is also linear to the number of nodes in the graph as DEEPWALK.

## 4 EXPERIMENTS

In this section, we empirically evaluate the proposed methods through performing link prediction and node classification on several benchmark datasets.

### 4.1 Experiment Setup

**4.1.1 Datasets.** We conduct experiments on several benchmark datasets from various real-world applications. Table 1 shows some statistics of them. Note that we do some preprocessing on the original datasets by deleting self-loops and nodes with zero degree. Some descriptions of these datasets are summarized as follows:

- **Cora, Citeseer** [21]: Paper citation networks. Cora consists of 2708 papers with 7 categories, and Citeseer consists 3264 papers including 6 categories.
- **Wiki** [31]: Wiki is a network with nodes as web pages and edges as the hyperlinks between web pages.
- **CA-GrQc, CA-HepTh** [17]: Author collaboration networks. They describe scientific collaborations between authors with papers submitted to General Relativity and Quantum Cosmology category, and High Energy Physics, respectively.

**4.1.2 Baseline Models.** The descriptions of the baseline models are as follows:

- **Graph Factorization (GF)** [2]: GF directly factorizes the adjacency matrix with stochastic gradient descent technique to obtain the embeddings, which enables it scale to large networks.
- **DEEPWALK** [27]: DEEPWALK regards node sequence obtained from truncated random walk as word sequence, and then uses skip-gram model to learn node representations. We directly use the publicly available source code with hierarchical softmax approximation for experiments.

- **LINE** [35]: LINE preserves network structural proximities through modeling node co-occurrence probability and node conditional probability, and leverages the negative sampling approach to alleviate the expensive computation.
- **node2vec** [13]: node2vec differs from DEEPWALK by proposing more flexible method for sampling node sequences to strike a balance between local and global structural properties.
- **GraRep** [4]: GraRep applies SVD technique to different  $k$ -step probability transition matrix to learn node embeddings, and finally obtains global representations through concatenating all  $k$ -step representations.
- **AIDW** [8]: AIDW is an inductive version of DEEPWALK with GAN-based regularization method. A prior distribution is imposed on node representations through adversarial learning to achieve a global smoothness in the distribution.

Our implemented version of DEEPWALK is based on negative sampling approach, thus we denote it as DwNS to avoid confusion. We also include a baseline, namely DwNS\_rand, with noises from a normal distribution as perturbations in the regularization term. Following existing work [30], we denote the adversarial training DEEPWALK as DwNS\_AdvT, and the interpretable adversarial training DEEPWALK as DwNS\_iAdvT in the rest of the paper.

**4.1.3 Parameter Settings.** For DwNS and its variants including DwNS\_rand, DwNS\_AdvT and DwNS\_iAdvT, the walk length, walks per node, window size, negative size, regularization strength, batch size and learning rate are set to 40, 1, 5, 5, 1, 1024 and 0.001, respectively. The adversarial noise level  $\epsilon$  has different settings in DwNS\_AdvT and DwNS\_iAdvT, while DwNS\_rand follows the settings of DwNS\_AdvT. For DwNS\_AdvT,  $\epsilon$  is set to different value for different datasets. Specifically,  $\epsilon$  is set to 0.9 for Cora, 1.1 for Citeseer in both link prediction and node classification, and 0.6 and 0.5 for Wiki in node classification and link prediction respectively, while it is set to 0.5 for all other datasets in these two learning tasks. For DwNS\_iAdvT,  $\epsilon$  is set to 5 for all datasets in both node classification and link prediction tasks, and the size of the nearest neighbor set  $T$  is set to 5. Besides, the dimension of embedding vectors are set to 128 for all methods.

### 4.2 Impact of Adversarial Training Regularization

In this section, we conduct link prediction and multi-class classification on adversarial training DEEPWALK, i.e., DwNS\_AdvT, to study the impact of adversarial training regularization on network representation learning from two aspects: model performance on different training epochs and model performance under different model size.

Node classification is conducted with support vector classifier in Liblinear package<sup>1</sup> [10] in default settings with the learned embedding vectors as node features. In link prediction, network embedding is first performed on a sub-network, which contains 80% of edges in the original network, to learn node representations. Note that the degree of each node is ensured to be greater than or equal to 1 during subsampling process to avoid meaningless embedding vectors. We use AUC score as the performance measure, and treat

<sup>1</sup><https://www.csie.ntu.edu.tw/~cjlin/liblinear/>




Figure 3: Training curves of node classification (left, training ratio 10%) and link prediction (right).




Figure 4: Performance comparison between Dwms and Dwms\_AdvT on multi-class classification with training ratio as 10% (left) and 50% (right) respectively under varying embedding size.

link prediction as a classification problem. Specifically, a  $L_2$ -SVM classifier is trained with edge feature inputs obtained from the Hadamard product of embedding vectors of two endpoints as many other works [13, 40], positive training samples as the observed 80% edges, and the same number of negative training samples randomly sampled from the network, i.e., node pairs without direct edge connection. The testing set consists of the hidden 20% edges and two times of randomly sampled negative edges. All experimental results are obtained by making an average of 10 different runs.

**4.2.1 Training Process.** We train Dwms model for 100 epochs, and evaluate the generalization performance of the learned embedding vectors in each epoch with node classification and link prediction on Cora, Citeseer and Wiki. We also conduct similar experiments on Dwms\_AdvT for 90 epochs with the model parameters initialized from those of Dwms after 10<sup>th</sup> training epochs. Figures 3 shows the experimental results.

In general, adversarial training regularization can bring a significant improvement in generalization ability to Dwms through the observation of training curves in both node classification and link prediction. Specifically, after 10 training epochs, the evaluation performance has little improvements for all datasets in two learning tasks with further training for Dwms, while adversarial training regularization leads to an obvious performance increase. In Figure 3, the blue line is drew by setting its vertical coordinates as the maximum value of the metrics of Dwms in the corresponding experiments. We can find that the training curve of Dwms\_AdvT is continuously above the blue line in different training epochs. Particularly, there is an impressive 7.2% and 9.2% relative performance improvement in link prediction for Cora and Citeseer respectively. We notice that the performance of Dwms\_AdvT drops slightly after

about 40 training epochs for Cora in link prediction, and about 20 training epochs for Wiki in node classification. The reason might be that some networks are more vulnerable to overfitting, and deeper understanding of this phenomenon needs further exploration.

**4.2.2 Performance vs. Embedding Size.** We explore the effect of adversarial regularization under different model size with multi-class classification. Figure 4 demonstrates the classification results on Cora, Citeseer and Wiki with training ratio as 10% and 50%. In general, adversarial training regularization is essential for improving model generalization ability. Across all tested embedding size, our proposed adversarial training DEEPWALK can consistently outperform the base model. For two models, when varying embedding size from 2 to 512, the classification accuracy firstly increases in a relatively fast speed, then grows slowly, and finally becomes stable or even drops slightly. The reason is that model generalization ability is improved with the increase of model capacity firstly until some threshold, since more network structural information can be captured with larger model capacity. However, when the model capacity becomes too large, it can easily result in overfitting, and thus cause performance degradation. We notice that the performance improvement of Dwms\_AdvT over Dwms is quite small when the embedding size is 2. It is probably because model capacity is the main reason limiting model performance and model robustness is not a serious issue when embedding size is too small.

### 4.3 Link Prediction

Link prediction is essential for many applications such as extracting missing information and identifying spurious interaction [20]. In this section, we conduct link prediction on five real-world networks,

and compare our proposed methods with the state-of-the-art methods. The experimental settings have been illustrated in Section 4.2. Table 2 summarizes the experimental results.

It can be easily observed that both our proposed methods, including DwNS\_AdvT and DwNS\_iAdvT, performs better than DwNS in all five datasets, which demonstrates that two types of adversarial regularization methods can help improve model generalization ability. Specifically, there is a 4.62% performance improvement for DwNS\_AdvT over DwNS on average across all datasets, and that for DwNS\_iAdvT is 4.60%, which are very impressive.

We noticed that AIDW has a poor performance in link prediction. The reasons can be two-folds: firstly, AIDW encourages the smoothness of embedding distribution from a global perspective by imposing a prior distribution on them, which can result in over-regularization and thus cause performance degradation; secondly, AIDW suffers from mode-collapse problem because of its generative adversarial network component, which can also result in model corruption. Besides, DwNS\_rand has similar performance with DwNS, which means that the regularization term with random perturbation contributes little to model generalization ability. By comparison, our proposed novel adversarial training regularization method is more stable and effective.

It can be observed that the performance of DwNS\_AdvT and DwNS\_iAdvT are comparable. Either DwNS\_AdvT or DwNS\_iAdvT achieves the best results across the five datasets, which shows the remarkable usefulness of the proposed regularization methods. For Cora and CA-GrQc, DwNS\_iAdvT has better performance, although we restrict the perturbation directions toward the nearest neighbors of the considered node. It suggests that such restriction of perturbation directions might provide useful information for representation learning.

#### 4.4 Node Classification

Node classification can be conducted to dig out missing information in a network. In this section, we conduct multi-class classification on three benchmark datasets, including Cora, Citeseer and Wiki, with the training ratio ranging from 1% to 90%. Tables 3, 4 and 5 summarize the experimental results.

Firstly, DwNS\_rand and DwNS have similar performance in all three datasets. For example, the average improvement of DwNS\_rand over DwNS is 0.16% across all training ratios in Wiki, which can be negligible. It validates that random perturbation for the regularization term contributes little to the model generalization performance again. It is understandable, since the expected dot product between any reference vector and the random perturbation from a zero mean gaussian distribution is zero, and thus the regularization term will barely affect the embedding learning.

Secondly, DwNS\_AdvT and DwNS\_iAdvT consistently outperform DwNS across all different training ratios in the three datasets, with the only exception of DwNS\_iAdvT in Citeseer when the training ratio is 3%. Specifically, DwNS\_AdvT achieves 5.06%, 6.45% and 5.21% performance gain over DwNS on average across all training ratios in Cora, Citeseer and Wiki respectively, while the improvement over DwNS for DwNS\_iAdvT are 2.35%, 4.50% and 2.62% respectively. It validates that adversarial perturbation can provide useful

direction for generating adversarial examples, and thus brings significant improvements to model generalization ability after the adversarial training process. For DwNS\_iAdvT, it brings less performance gain compared with DwNS\_AdvT, which might be because the restriction on perturbation direction limit its regularization ability for classification tasks. In this case, there is a tradeoff between interpretability and regularization effect.

Thirdly, AIDW achieves better results than DEEPWALK, LINE and GraRep, which shows that global regularization on embedding vectors through adversarial learning can help improve model generalization performance. Our proposed methods, especially DwNS\_AdvT, demonstrate superiority over all the state-of-the-art baselines, including AIDW and node2vec, based on experimental results comparison. We can summarize that the adversarial training regularization method has advantages over the GAN-based global regularization methods in three aspects, including more succinct architecture, better computational efficiency and more effective performance contribution.

#### 4.5 Parameter Sensitivity

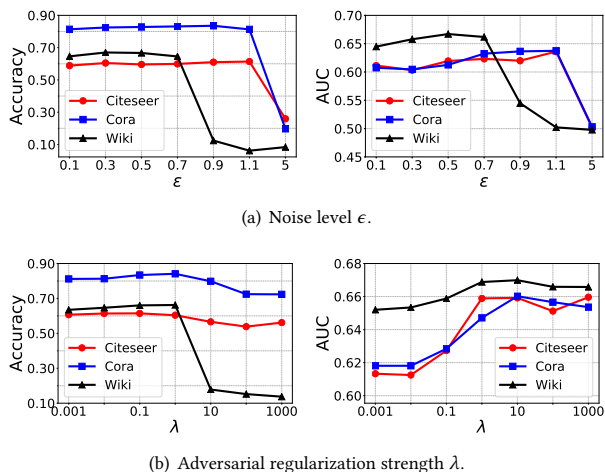
We conduct parameter sensitivity analysis with link prediction and multi-class classification on Cora, Citeseer and Wiki in this section. Here we only present the results for DwNS\_AdvT due to space limitation. Adversarial training regularization method is very succinct. DwNS\_AdvT only has two more hyperparameters compared with DwNS, which are noise level  $\epsilon$  and adversarial regularization strength  $\lambda$ . Note that when studying one hyper-parameter, we follow default settings for other hyper-parameters. The experimental settings of link prediction and node classification have been explained in Section 4.2.

Fig. 5(a) presents the experimental results when varying  $\epsilon$  from 0.1 to 5.0. For both learning tasks, we can find that the performance in these three datasets first improves with the increase of  $\epsilon$ , and then drops dramatically after  $\epsilon$  passing some threshold. It suggests that appropriate setting of  $\epsilon$  improves the model robustness and generalization ability, while adversarial perturbation with too large norm constraint can destroy the learning process of embedding vectors. Besides, it can be easily noticed that the best settings of  $\epsilon$  are different for different datasets in general. Specifically, Citeseer has the best results in both link prediction and node classification when  $\epsilon = 1.1$ , Cora achieves the best results when  $\epsilon = 0.9$ , while the best setting of  $\epsilon$  for Wiki is around 0.5. Based on the experimental results on these three datasets only, it seems that the denser the network is, the smaller the best noise level parameter  $\epsilon$  should be.

We conduct link prediction and node classification on three datasets with the adversarial regularization strength  $\lambda$  from the set  $\{0.001, 0.01, 0.1, 1, 10, 100, 1000\}$ . Fig. 5(b) displays the experimental results. For node classification, the best result is obtained when  $\lambda$  is set to around 1, larger values can result in performance degradation. For example, the classification accuracy on Wiki drops dramatically when  $\lambda$  reaches 10, and larger setting produces worse results. For link prediction, the performance is quite consistent among the three datasets. Specifically, when  $\lambda$  increases from 0.001 to 10, the AUC score shows apparent increase for all datasets, and then tends to saturate or decrease slightly. Empirically, 1 is an appropriate value for the adversarial regularization strength  $\lambda$ .







**Figure 5: Impact of hyperparameters on node classification (left, training ratio 50%) and link prediction (right).**

first-order and second-order proximities in network, and resorts to negative sampling for capturing them.

Further, some works [4, 26, 39] tried to preserve various network structural properties in embedding vectors based on matrix factorization technique. GraRep [4] can preserve different  $k$ -step proximities between nodes independently, HOPE [26] aims to capture asymmetric transitivity property in node embeddings, while N-NMF [39] learns community structure preserving embedding vectors by building upon the modularity based community detection model [24]. Meanwhile, deep learning embedding models [5, 32, 33, 37] have also been proposed to capture highly non-linear structure. DNCR [5] takes advantages of deep denoising autoencoder for learning compact node embeddings, which can also improve model robustness. SDNE [37] modifies the framework of stacked autoencoder to learn both first-order and second-order proximities simultaneously. DNE-SBP [33] utilizes a semi-supervised SAE to preserve the structural balance property of the signed networks. Both GraphGAN [38] and A-RNE [9] leverage generative adversarial networks to facilitate network embedding, with the former unifies the generative models and discriminative models of network embedding to boost the performance while the latter focuses on sampling high-quality negative nodes to achieve better similarity ranking among node pairs.

However, the above mentioned models mainly focus on learning different network structures and properties, while neglecting the existence of noisy information in real-world networks and the overfitting issue in embedding learning process. Most recently, some methods, including ANE [8] and NETRA [41], try to regularize the embedding learning process for improving model robustness and generalization ability based on generative adversarial networks (GANs). They have very complicated frameworks and suffer from the well-recognized hard training problems of GANs. Furthermore, these two methods both encourage the global smoothness of the embedding distribution, while in this paper we utilize a more succinct and effective local regularization method.

**Adversarial Machine Learning.** It was found that several machine learning models, including both deep neural network and shallow classifiers such as logistic regression, are vulnerable to examples with imperceptibly small designed perturbations, called adversarial examples [12, 34]. This phenomenon was firstly observed in areas like computer vision with continuous input vectors. To improve model robustness and generalization ability, adversarial training method [12] is shown to be effective. It generates adversarial perturbations for original clean input with the aim of maximizing current model loss, and further approximates the difficult optimization objective with first-order Taylor Series. Such method has also been applied to text classification problem in [23, 30] by defining the perturbation on continuous word embeddings, and recommendation in [15] by generating adversarial perturbations on model parameters. However, to the best of our knowledge, there is no practice of adversarial training regularization for graph representation learning.

For graph structured data, they are fundamentally different from images because of their discrete and indifferentiable characteristics. Some existing works [6, 7, 42] aimed to explore how to generate the adversarial examples in the discrete, binary graph domain, and whether similar vulnerability exists in graph analysis applications. In [7], adversarial attacks are generated by modifying combinatorial structure of graph with a reinforcement learning based method, which is shown to be effective in Graph Neural Network models. Both [42] and [6] designed attack methods to Graph Convolutional Network [16]. Particularly, NETTACK [42] focuses on attributed graph classification problem and FGA [6] tackles network representation learning. However, all of them studied adversarial attack methods without providing any defense algorithms for improving the robustness of existing methods against these attacks. Differently, in this paper, we aim to propose adversarial regularization method for network embedding algorithms to improve both model robustness and generalization ability.

## 6 CONCLUSION

In this paper, we proposed two adversarial training regularization methods for network embedding models to improve the robustness and generalization ability. Specifically, the first method is adapted from the classic adversarial training method by defining the perturbation in the embedding space with adaptive  $L_2$  norm constraint. Though it is effective as a regularizer, the lack of interpretability may hinder its adoption in some real-world applications. To tackle this problem, we further proposed an interpretable adversarial training method by restricting the perturbation directions to embedding vectors of other nodes, such that the crafted adversarial examples can be reconstructed in the discrete graph domain. Both methods can be applied to the existing embedding models with node embeddings as model parameters, and DEEPWALK is used as the base model in the paper for illustration. Extensive experiments prove the effectiveness of the proposed adversarial regularization methods for improving model robustness and generalization ability. Future works would include applying adversarial training method to the parameterized network embedding methods such as deep learning embedding models.

## 7 ACKNOWLEDGMENTS

Parts of the work were supported by HK ITF UIM/363.

## REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek Gordon Murray, Benoit Steiner, Paul A. Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *OSDI*. USENIX Association, 265–283.
- [2] Amr Ahmed, Nino Shervashidze, Shravan M. Narayanamurthy, Vanja Josifovski, and Alexander J. Smola. 2013. Distributed large-scale natural graph factorization. In *WWW*. 37–48.
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein Generative Adversarial Networks. In *ICML*. 214–223.
- [4] Shaosheng Cao, Wei Lu, and Qionghai Xu. 2015. GraRep: Learning Graph Representations with Global Structural Information. In *CIKM*. 891–900.
- [5] Shaosheng Cao, Wei Lu, and Qionghai Xu. 2016. Deep Neural Networks for Learning Graph Representations. In *AAAI*. 1145–1152.
- [6] Jinyin Chen, Yangyang Wu, Xuanheng Xu, Yixian Chen, Haibin Zheng, and Qi Xuan. 2018. Fast Gradient Attack on Network Embedding. *CoRR* abs/1809.02797 (2018).
- [7] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. 2018. Adversarial Attack on Graph Structured Data. In *ICML (JMLR Workshop and Conference Proceedings)*, Vol. 80. 1123–1132.
- [8] Quanyu Dai, Qiang Li, Jian Tang, and Dan Wang. 2018. Adversarial Network Embedding. In *AAAI*.
- [9] Quanyu Dai, Qiang Li, Liang Zhang, and Dan Wang. 2019. Ranking Network Embedding via Adversarial Learning. In *PAKDD*.
- [10] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. 2008. LIBLINEAR: A Library for Large Linear Classification. *JMLR* 9 (2008), 1871–1874.
- [11] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. Generative Adversarial Nets. In *NIPS*. 2672–2680.
- [12] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. 2014. Explaining and Harnessing Adversarial Examples. *CoRR* abs/1412.6572 (2014).
- [13] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable Feature Learning for Networks. In *KDD*. 855–864.
- [14] Yupeng Gu, Yizhou Sun, Yanen Li, and Yang Yang. 2018. RaRE: Social Rank Regulated Large-scale Network Embedding. In *WWW*. ACM, 359–368.
- [15] Xiangnan He, Zhankui He, Xiaoyu Du, and Tat-Seng Chua. 2018. Adversarial Personalized Ranking for Recommendation. In *SIGIR*. ACM, 355–364.
- [16] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *ICLR*.
- [17] Jure Leskovec, Jon M. Kleinberg, and Christos Faloutsos. 2007. Graph evolution: Densification and shrinking diameters. *TKDD* 1, 1 (2007), 2.
- [18] Omer Levy and Yoav Goldberg. 2014. Neural Word Embedding as Implicit Matrix Factorization. In *NIPS*. 2177–2185.
- [19] Aaron Q. Li, Amr Ahmed, Sujith Ravi, and Alexander J. Smola. 2014. Reducing the sampling complexity of topic models. In *KDD*. 891–900.
- [20] Linyuan Lv and Tao Zhou. 2011. Link prediction in complex networks: A survey. *Physica A: Statistical Mechanics and its Applications* 390, 6 (2011), 1150 – 1170.
- [21] Andrew McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. 2000. Automating the Construction of Internet Portals with Machine Learning. *Inf. Retr.* 3, 2 (2000), 127–163.
- [22] Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *NIPS*. 3111–3119.
- [23] Takeru Miyato, Andrew M. Dai, and Ian Goodfellow. 2017. Adversarial Training Methods for Semi-Supervised Text Classification. In *ICLR*.
- [24] M. E. J. Newman. 2006. Finding community structure in networks using the eigenvectors of matrices. *Phys. Rev. E* 74 (Sep 2006), 036104. Issue 3.
- [25] Feng Niu, Benjamin Recht, Christopher Ré, and Stephen J. Wright. 2011. HOG-WILD!: A Lock-Free Approach to Parallelizing Stochastic Gradient Descent. *Advances in Neural Information Processing Systems* (2011), 693–701.
- [26] Mingdong Ou, Peng Cui, Jian Pei, Ziwei Zhang, and Wenwu Zhu. 2016. Asymmetric Transitivity Preserving Graph Embedding. In *KDD*. 1105–1114.
- [27] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. DeepWalk: online learning of social representations. In *KDD*. 701–710.
- [28] Sam T. Roweis and Lawrence K. Saul. 2000. Nonlinear Dimensionality Reduction by Locally Linear Embedding. *Science* 290 (2000), 2323–2326. Issue 5500.
- [29] Tim Salimans, Ian J. Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved Techniques for Training GANs. In *NIPS*. 2226–2234.
- [30] Motoki Sato, Jun Suzuki, Hiroyuki Shindo, and Yuji Matsumoto. 2018. Interpretable Adversarial Perturbation in Input Embedding Space for Text. In *IJCAL*. 4323–4330.
- [31] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Galligher, and Tina Eliassi-Rad. 2008. Collective Classification in Network Data. *AI Magazing* 29, 3 (2008), 93–106.
- [32] Xiao Shen and Fu-Lai Chung. 2017. Deep Network Embedding with Aggregated Proximity Preserving. In *Proceedings of the IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. 40–43.
- [33] Xiao Shen and Fu-Lai Chung. 2018. Deep Network Embedding for Graph Representation Learning in Signed Networks. *IEEE Transactions on Cybernetics* (2018).
- [34] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing properties of neural networks. In *ICLR*.
- [35] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. 2015. LINE: Large-scale Information Network Embedding. In *WWW*. 1067–1077.
- [36] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. 2000. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science* 290 (2000), 2319–2323. Issue 5500.
- [37] Daixin Wang, Peng Cui, and Wenwu Zhu. 2016. Structural Deep Network Embedding. In *KDD*. 1225–1234.
- [38] Hongwei Wang, Jia Wang, Jialin Wang, Miao Zhao, Weinan Zhang, Fuzheng Zhang, Xing Xie, and Minyi Guo. 2018. GraphGAN: Graph Representation Learning with Generative Adversarial Nets. In *AAAI*.
- [39] Xiao Wang, Peng Cui, Jing Wang, Jian Pei, Wenwu Zhu, and Shiqiang Yang. 2017. Community Preserving Network Embedding. In *AAAI*. 203–209.
- [40] Zhitao Wang, Chengyao Chen, and Wenjie Li. 2017. Predictive Network Representation Learning for Link Prediction. In *SIGIR*. 969–972.
- [41] Wenchao Yu, Cheng Zheng, Wei Cheng, Charu C. Aggarwal, Dongjin Song, Bo Zong, Haifeng Chen, and Wei Wang. 2018. Learning Deep Network Representations with Adversarially Regularized Autoencoders. In *KDD*. 2663–2671.
- [42] Daniel Zügner, Amir Akbarnejad, and Stephan Günnemann. 2018. Adversarial Attacks on Neural Networks for Graph Data. In *KDD*. 2847–2856.