

On Generalization of Graph Autoencoders with Adversarial Training

Citation for published version (APA):

Huang, T., Pei, Y., Menkovski, V., & Pechenizkiy, M. (2021). On Generalization of Graph Autoencoders with Adversarial Training. In N. Oliver, F. Pérez-Cruz, S. Kramer, J. Read, & J. A. Lozano (Eds.), *Machine Learning and Knowledge Discovery in Databases. Research Track: European Conference, ECML PKDD 2021, Bilbao, Spain, September 13–17, 2021, Proceedings* (Vol. II, pp. 367-382). (Lecture Notes in Computer Science (LNCS); Vol. 12976), (Lecture Notes in Artificial Intelligence (LNAI); Vol. 12976). Springer.
https://doi.org/10.1007/978-3-030-86520-7_23

Document license:

TAVERNE

DOI:

[10.1007/978-3-030-86520-7_23](https://doi.org/10.1007/978-3-030-86520-7_23)

Document status and date:

Published: 10/09/2021

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.



On Generalization of Graph Autoencoders with Adversarial Training

Tianjin Huang^(✉), Yulong Pei, Vlado Menkovski, and Mykola Pechenizkiy

Department of Mathematics and Computer Science,
Eindhoven University of Technology, 5600 Eindhoven, MB, The Netherlands
{t.huang,y.pei.1,v.menkovski,m.pechenizkiy}@tue.nl

Abstract. Adversarial training is an approach for increasing model’s resilience against adversarial perturbations. Such approaches have been demonstrated to result in models with feature representations that generalize better. However, limited works have been done on adversarial training of models on graph data. In this paper, we raise such a question – does adversarial training improve the generalization of graph representations. We formulate L_2 and L_∞ versions of adversarial training in two powerful node embedding methods: graph autoencoder (GAE) and variational graph autoencoder (VGAE). We conduct extensive experiments on three main applications, i.e. link prediction, node clustering, graph anomaly detection of GAE and VGAE, and demonstrate that both L_2 and L_∞ adversarial training boost the generalization of GAE and VGAE.

Keywords: Graph autoencoders · Variational graph autoencoders · Adversarial training · Node embedding · Generalization

1 Introduction

Networks are ubiquitous in a plenty of real-world applications and they contain relationships between entities and attributes of entities. Modeling such data is challenging due to its non-Euclidean characteristic. Recently, graph embedding that converts graph data into low dimensional feature space has emerged as a popular method to model graph data, For example, DeepWalk [14], node2vec [7] and LINE [23] learn graph embedding by extracting patterns from the graph. Graph Convolutions Networks (GCNs) [9] learn graph embedding by repeated multiplication of normalized adjacency matrix and feature matrix. In particular, graph autoencoder (GAE) [10, 24, 27] and graph variational autoencoder (VGAE) [10] have been shown to be powerful node embedding methods as unsupervised learning. They have been applied to many machine learning tasks, e.g. node clustering [16, 20, 24], link prediction [10, 19], graph anomaly detection [4, 13] and etc.

Adversarial training is an approach for increasing model’s resilience against adversarial perturbations by including adversarial examples in the training set [11]. Several recent studies demonstrate that adversarial training improves feature representations leading to better performance for downstream tasks [18, 26].

However, little work in this direction has been done for GAE and VGAE. Besides, real-world graphs are usually highly noisy and incomplete, which may lead to a sub-optimal results for standard trained models [32]. Therefore, we are interested to seek answers to the following two questions:

- *Does adversarial training improve generalization, i.e. the performance in applications of node embeddings learned by GAE and VGAE?*
- *Which factors influence this improvement?*

In order to answer the first question above, we firstly formulate L_2 and L_∞ adversarial training for GAE and VGAE. Then, we select three main tasks of VGAE and GAE: link prediction, node clustering and graph anomaly detection for evaluating the generalization performance brought by adversarial training. Besides, we empirically explore which factors affect the generalization performance brought by adversarial training.

Contributions: To the best of our knowledge, we are the first to explore generalization for GAE and VGAE using adversarial training. We formulate L_2 and L_∞ adversarial training, and empirically demonstrate that both L_2 and L_∞ adversarial training boost the generalization with a large margin for the node embeddings learned by GAE and VGAE. An additional interesting finding is that the generalization performance of the proposed adversarial training is more sensitive to attributes perturbation than adjacency matrix perturbation and not sensitive to the degree of nodes.

2 Related Work

Adversarial training has been extensively studied in images. It has been important issues to explore whether adversarial training can help generalization. Tsipras et al. [25] illustrates that adversarial robustness could conflict with model’s generalization by a designed simple task. However, Stutz et al. [21] demonstrates that adversarial training with on-manifold adversarial examples helps the generalization. Besides, Salman et al. [18] and Utrera et al. [26] show that the latent features learned by adversarial training are improved and boost the performance of their downstream tasks.

Recently, few works bring adversarial training in graph data. Deng, Dong and Zhu [3] and Sun et al. [22] propose virtual graph adversarial training to promote the smoothness of model. Feng et al. [5] propose graph adversarial training by inducing dynamical regularization. Dai et al. [2] formulate an interpretable adversarial training for DeepWalk. Jin and Zhang [8] introduce latent adversarial training for GCN, which train GCN based on the adversarial perturbed output of the first layer. Besides, several studies explored adversarial training based on adversarial perturbed edges for graph data [1, 28, 31]. Among these works, part of studies pay attention to achieving model’s robustness while ignoring the effect of generalization [1, 2, 8, 28, 31] and the others simply utilize perturbations on nodal attributes while not explore the effect of perturbation on edges [3, 5, 22].

The difference between these works and ours is two-fold: (1) We extend both L_∞ and L_2 adversarial training for graph models while the previous studies only explore L_2 adversarial training. (2) We focus on the generalization effect brought by adversarial training for unsupervised deep learning graph models, i.e. GAE and VGAE while most of the previous studies focus on adversarial robustness for supervised/semi-supervised models.

3 Preliminaries

We first summarize some notations and definitions used in this paper. Following the commonly used notations, we use bold uppercase characters for matrices, e.g. \mathbf{X} , bold lowercase characters for vectors, e.g. \mathbf{b} , and normal lowercase characters for scalars, e.g. c . The i^{th} row of a matrix \mathbf{A} is denoted by $\mathbf{A}_{i,:}$, and $(i, j)^{th}$ element of matrix \mathbf{A} is denoted as $\mathbf{A}_{i,j}$. The i^{th} row of a matrix \mathbf{X} is denoted by \mathbf{x}_i . We use \mathbf{KL} for Kullback-Leibler divergence.

We consider an attributed network $\mathcal{G} = \{V, E, \mathbf{X}\}$ with $|V| = n$ nodes, $|E| = m$ edges and \mathbf{X} node attributed matrix. \mathbf{A} is the binary adjacency matrix of \mathcal{G} .

3.1 Graph Autoencoders

Graph autoencoders is a kind of unsupervised learning models on graph-structure data [10], which aim at learning low dimensional representations for each node by reconstructing inputs. It has been demonstrated to achieve competitive results in multiple tasks, e.g. link prediction [10, 16, 17], node clustering [16, 20, 24], graph anomaly detection [4, 13]. Generally, graph autoencoder consists of a graph convolutional network for encoder and an inner product for decoder [10]. Formally, it can be expressed as follows:

$$\mathbf{Z} = GCN(\mathbf{A}, \mathbf{X}) \quad (1)$$

$$\hat{\mathbf{A}} = \sigma(\mathbf{Z}\mathbf{Z}^T), \quad (2)$$

where σ is the sigmoid function, GCN is a graph convolutional network, \mathbf{Z} is the learned low dimensional representations and $\hat{\mathbf{A}}$ is the reconstructed adjacency matrix.

During the training phase, the parameters will be updated by minimizing the reconstruction loss. Usually, the reconstruction loss is expressed as cross-entropy loss between \mathbf{A} and $\hat{\mathbf{A}}$ [10]:

$$\mathcal{L}^{ae} = -\frac{1}{n^2} \sum_{(i,j) \in V \times V} \left[\mathbf{A}_{i,j} \log \hat{\mathbf{A}}_{i,j} + (1 - \mathbf{A}_{i,j}) \log (1 - \hat{\mathbf{A}}_{i,j}) \right]. \quad (3)$$

3.2 Graph Variational Autoencoders

Kipf and Welling [10] introduced variational graph autoencoder (VGAE) which is a probabilistic model. VGAE is consisted of inference model and generative

model. In their approach, the inference model, i.e. corresponding to the encoder of VGAE, is expressed as follows:

$$q(\mathbf{Z}|\mathbf{X}, \mathbf{A}) = \prod_{i=1}^N q(\mathbf{z}_i|\mathbf{X}, \mathbf{A}), \text{ with } q(\mathbf{z}_i|\mathbf{X}, \mathbf{A}) = \mathcal{N}(\mathbf{z}_i|\boldsymbol{\mu}_i, \text{diag}(\boldsymbol{\sigma}_i^2)), \quad (4)$$

where $\boldsymbol{\mu}_i$ and $\boldsymbol{\sigma}_i$ are learned by a graph neural network respectively. That is, $\boldsymbol{\mu} = GCN_{\boldsymbol{\mu}}(\mathbf{X}, \mathbf{A})$ and $\log\boldsymbol{\sigma} = GCN_{\delta}(\mathbf{X}, \mathbf{A})$, with $\boldsymbol{\mu}$ is the matrix of stacking vectors $\boldsymbol{\mu}_i$; likewise, $\boldsymbol{\sigma}$ is the matrix of stacking vectors δ_i .

The generative model, i.e. corresponding to the decoder of autoencoder, is designed as an inner product between latent variables \mathbf{Z} , which is formally expressed as follows:

$$p(\mathbf{A}|\mathbf{Z}) = \prod_{i=1}^n \prod_{j=1}^n p(\mathbf{A}_{i,j}|\mathbf{z}_i, \mathbf{z}_j), \text{ with } p(\mathbf{A}_{i,j} = 1|\mathbf{z}_i, \mathbf{z}_j) = \sigma(\mathbf{z}_i^T \mathbf{z}_j). \quad (5)$$

During the training phase, the parameters will be updated by minimizing the variational lower bound \mathcal{L}^{vae} :

$$\mathcal{L}^{vae} = \mathbf{E}_{q(\mathbf{Z}|\mathbf{X}, \mathbf{A})}[\log p(\mathbf{A}|\mathbf{Z})] - \mathbf{KL}[q(\mathbf{Z}|\mathbf{X}, \mathbf{A})||p(\mathbf{Z})], \quad (6)$$

where a Gaussian prior is adopted for $p(\mathbf{Z}) = \prod_i p(\mathbf{z}_i) = \prod_i \mathcal{N}(\mathbf{z}_i|0, \mathbf{I})$.

3.3 Adversarial Training

By now, multiple variants of adversarial training has been proposed and most of them are built on supervised learning and Euclidean data, e.g. FGSM-adversarial training [6], PGD-adversarial training [11], Trades [33], MART [29] and etc. Here we introduce Trades that will be extended to GAE and VGAE settings in Sect. 4. Trades [33] separates loss function into two terms: 1) Cross-Entropy Loss for achieving natural accuracy; 2) Kullback-Leibler divergence for achieving adversarial robustness. Formally, given inputs (X, Y) , it can be expressed as follows [33]:

$$\min_{\theta} \mathbf{E}_{(X, Y)}[L(f_{\theta}(X), Y) + \lambda \cdot \mathbf{KL}(P(Y|X')||P(Y|X))], \quad (7)$$

where f_{θ} is a supervised model, X' is the adversarial examples that maximize \mathbf{KL} divergence and $P(Y|X)$ is the output probability after softmax. λ is a tunable hyperparameter and it controls the strength of the \mathbf{KL} regularization term.

4 Graph Adversarial Training

In this section, we formulate L_2 and L_{∞} adversarial training for GAE and VGAE respectively.

4.1 Adversarial Training in Graph Autoencoder

Considering that: (1) the inputs of GAE contains adjacency matrix and attributes, (2) the latent representation \mathbf{Z} is expected to be invariant to the input perturbation, we reformulate the loss function in Eq. 3 as follows:

$$\min_{\theta} \mathcal{L}^{ae} + \lambda \cdot \mathbf{KL}(P(\mathbf{Z}|\mathbf{A}', \mathbf{X}') || P(\mathbf{Z}|\mathbf{A}, \mathbf{X})) \tag{8}$$

$$\mathbf{X}' = \mathit{arg} \max_{\|\mathbf{X}' - \mathbf{X}\| \leq \epsilon} \mathcal{L}^{ae}(\mathbf{A}, \mathbf{X}), \mathbf{A}' = \mathit{arg} \max_{\|\mathbf{A}' - \mathbf{A}\| \leq \epsilon} \mathcal{L}^{ae}(\mathbf{A}, \mathbf{X}) \tag{9}$$

where \mathbf{A}' is the adversarial perturbed adjacency matrix and \mathbf{X}' is the adversarial perturbed attributes. Here the important question is how to generate the perturbed adjacency matrix \mathbf{A}' and attributes \mathbf{X}' in Eq. 9.

Attributes Perturbation \mathbf{X}' . We generate the perturbed X' by projection gradient descent (PGD) [11]. We denote total steps as T .

For \mathbf{X}' bounded by L_2 norm ball, the perturbed data in t -th step \mathbf{X}^t is expressed as follows:

$$\mathbf{X}^t = \prod_{\mathcal{B}(\mathbf{X}, \epsilon \|X\|_2)} (\mathbf{X}^{t-1} + \alpha \cdot g \cdot \|X\|_2 / \|g\|_2) \tag{10}$$

$$g = \nabla_{\mathbf{X}^{t-1}} \mathcal{L}^{ae}(\mathbf{A}, \mathbf{X}^{t-1}) \tag{11}$$

where \prod is the projection operator and $\mathcal{B}(\mathbf{X}, \epsilon \|X\|_2)$ is the L_2 norm ball of nodal attributes $\mathbf{x}_i : \{\mathbf{x}'_i : \|\mathbf{x}'_i - \mathbf{x}_i\|_2 \leq \epsilon \|\mathbf{x}_i\|_2\}$.

For \mathbf{X}' bounded by L_∞ norm ball, the perturbed data in t -th step \mathbf{X}^t is expressed as follows:

$$\mathbf{X}^t = \prod_{\mathcal{B}(\mathbf{X}, \epsilon)} (\mathbf{X}^{t-1} + \alpha \cdot g) \tag{12}$$

$$g = \mathit{sgn}(\nabla_{\mathbf{X}^{t-1}} \mathcal{L}^{ae}(\mathbf{A}, \mathbf{X}^{t-1})), \tag{13}$$

where $\mathcal{B}(\mathbf{X}, \epsilon)$ is the L_∞ norm ball of nodal attributes $\mathbf{x}_i : \{\mathbf{x}'_i : \|\mathbf{x}'_i - \mathbf{x}_i\|_\infty \leq \epsilon\}$ and $\mathit{sgn}(\cdot)$ is the sign function.

Adjacency Matrix Perturbation \mathbf{A}' . Adjacency matrix perturbation includes two-fold: (1) perturb node connections, i.e. Adding or dropping edges, (2) perturb the strength of information flow between nodes, i.e. the strength of correlation between nodes. Here we choose to perturb the strength of information flow between nodes and leave the perturb of node connections for future work. Specifically, we add weight for each edge and change these weights in order to perturb the strength of information flow. Formally, given the adjacency matrix \mathbf{A} , the weighted adjacency matrix $\tilde{\mathbf{A}}$ is expressed as $\mathbf{A} \odot \mathbf{M}$ where the elements of \mathbf{M} are continuous and its values are initialized as same value as \mathbf{A} . \odot denotes the element-wise product. Formally, \mathbf{A}' is expressed as follows:

$$\mathbf{M}' = \mathit{arg} \max_{\|\mathbf{M}' - \mathbf{M}\| \leq \epsilon} \mathcal{L}^{ae}(\tilde{\mathbf{A}}, \mathbf{X}) \tag{14}$$

$$\mathbf{A}' = \mathbf{A} \odot \mathbf{M}'. \tag{15}$$

For \mathbf{A}' bounded by L_2 norm ball, the perturbed data in t -th step \mathbf{A}^t is expressed as follows:

$$g = \nabla_{M^{t-1}} \mathcal{L}^{ae}(\tilde{\mathbf{A}}^{t-1}, \mathbf{X}) \quad (16)$$

$$\mathbf{M}^t = \prod_{\mathcal{B}(M, \epsilon \|M\|_2)} (\mathbf{M}^{t-1} + \alpha \cdot g \cdot \|M\|_2 / \|g\|_2) \quad (17)$$

$$\mathbf{A}^t = \tilde{\mathbf{A}}^t = \mathbf{A} \odot \mathbf{M}^t. \quad (18)$$

For \mathbf{A}' bounded by L_∞ norm ball, the perturbed data in t -th step \mathbf{A}^t is expressed as follows:

$$g = \text{sgn}(\nabla_{M^{t-1}} \mathcal{L}^{ae}(\tilde{\mathbf{A}}^{t-1}, \mathbf{X})) \quad (19)$$

$$\mathbf{M}^t = \prod_{\mathcal{B}(M, \epsilon)} (\mathbf{M}^{t-1} + \alpha \cdot g) \quad (20)$$

$$\mathbf{A}^t = \tilde{\mathbf{A}}^t = \mathbf{A} \odot \mathbf{M}^t. \quad (21)$$

4.2 Adversarial Training in Variational Graph Autoencoder

Similarly to GAE, we reformulate the loss function for training VGAE (Eq. 6) as follows:

$$\min_{\theta} \mathcal{L}^{vae} + \lambda \cdot \text{KL}(P(\mathbf{Z}|\mathbf{A}', \mathbf{X}') \| P(\mathbf{Z}|\mathbf{A}, \mathbf{X})) \quad (22)$$

$$\mathbf{X}' = \arg \max_{\|\mathbf{X}' - \mathbf{X}\| \leq \epsilon} \mathcal{L}^{vae}(\mathbf{A}, \mathbf{X}), \quad \mathbf{A}' = \arg \max_{\|\mathbf{A}' - \mathbf{A}\| \leq \epsilon} \mathcal{L}^{vae}(\mathbf{A}, \mathbf{X}) \quad (23)$$

We generate \mathbf{A}' and \mathbf{X}' exactly the same way as with GAE (replacing \mathcal{L}^{ae} with \mathcal{L}^{vae} in Eq. 10–21.)

For convenience, we abbreviate L_2 and L_∞ adversarial training as AT-2 and AT-Linf respectively in the following tables and figures where L_2/L_∞ denote both attributes and adjacency matrix perturbation are bounded by L_2/L_∞ norm ball.

In practice, we train models by alternatively adding adjacency matrix perturbation and attributes perturbation¹.

5 Experiments

In this section, we present the results of the performance evaluation of L_2 and L_∞ adversarial training under three main applications of GAE and VGAE: link prediction, node clustering, and graph anomaly detection. Then we conduct parameter analysis experiments to explore which factors influence the performance.

Datasets. We used six real-world datasets: Cora, Citeseer and PubMed for link prediction and node clustering tasks, and BlogCatalog, ACM and Flickr for the

¹ We find that optimizing models by alternatively adding these two perturbation is better than adding these two perturbation together (See Appendix).

graph anomaly detection task. The detailed descriptions of the six datasets are showed in Table 1.

Model Architecture. All our experiments are based on the GAE/VGAE model where the encoder/inference model is consisted with a two-layer GCN by default.

Table 1. Datasets descriptions.

DataSets	Cora	Citeseer	PubMed	BlogCatalog	ACM	Flickr
#Nodes	2708	3327	19717	5196	16484	7575
#Links	5429	4732	44338	171743	71980	239738
#Features	1433	3703	500	8189	8337	12074

5.1 Link Prediction

Metrics. Following [10], we use the area under a receiver operating characteristic curve (AUC) and average precision (AP) as the evaluation metric. We conduct 30 repeat experiments with random splitting datasets into 85%, 5% and 10% for training sets, validation sets and test sets respectively. We report the mean and standard deviation values on test sets.

Parameter Settings. We train models on Cora and Citeseer datasets with 600 epochs, and PubMed with 800 epochs. All models are optimized with Adam optimizer and 0.01 learning rate. The λ is set to 4. For attributes perturbation, the ϵ is set to $3e-1$ and $1e-3$ on Citeseer and Cora, 1 and $5e-3$ on PubMed for L_2 and L_∞ adversarial training respectively. For adjacency matrix perturbation, the ϵ is set to $1e-3$ and $1e-1$ on Citeseer and Cora, and $1e-3$ and $3e-1$ on PubMed for L_2 and L_∞ adversarial training respectively. The steps T is set to 1. The α is set to $\frac{\epsilon}{T}$.

For standard training GAE and VGAE, we run the official Pytorch geometric code² with 600 epochs for Citeseer and Cora datasets, 1000 epochs³ for PubMed dataset. Other parameters are set the same as in [10].

Experimental Results. The results are showed in Table 2. It can be seen that both L_2 and L_∞ Adversarial trained GAE and VGAE models consistently boost their performance for both AUC and AP metrics on Cora, Citeseer and PubMed datasets. Specifically, the improvements on Cora and Citeseer dataset reaches at least 2% for both GAE and VGAE (Table 2). The improvements on PubMed is relative small with around 0.3%.

² https://github.com/rusty1s/pytorch_geometric/blob/master/examples/autoencoder.py.

³ Considering PubMed is big graph data, we use more epochs in order to avoiding underfitting.

Table 2. Results for link prediction.

Methods	Cora		Citeseer		PubMed	
	AUC (in%)	AP (in%)	AUC (in%)	AP (in%)	AUC (in%)	AP (in%)
GAE	90.6 ± 0.9	91.2 ± 1.0	88.0 ± 1.2	89.2 ± 1.0	96.8 ± 0.2	97.1 ± 0.2
AT-L2-GAE	93.0 ± 0.9	93.5 ± 0.6	92.5 ± 0.7	93.2 ± 0.6	97.2 ± 0.2	97.4 ± 0.2
AT-Linf-GAE	92.8 ± 1.1	93.4 ± 1.0	92.3 ± 0.9	92.6 ± 1.1	96.9 ± 0.2	97.3 ± 0.2
VGAE	89.8 ± 0.9	90.3 ± 1.0	86.6 ± 1.4	87.6 ± 1.3	96.2 ± 0.4	96.3 ± 0.4
AT-L2-VGAE	92.8 ± 0.6	93.1 ± 0.6	90.7 ± 1.1	91.1 ± 0.9	96.6 ± 0.2	96.7 ± 0.2
AT-Linf-VGAE	92.2 ± 1.2	92.3 ± 1.3	91.9 ± 0.8	92.0 ± 0.6	96.5 ± 0.2	96.6 ± 0.3

5.2 Node Clustering

Metrics. Following [12,30], we use accuracy (ACC), normalized mutual information (NMI), precision, F-score (F1) and average rand index (ARI) as our evaluation metrics. We conduct 10 repeat experiments. For each experiment, datasets are random split into training sets (85% edges), validation sets (5% edges) and test sets (10% edges). We report the mean and standard deviation values on test sets.

Parameter Settings. We train GAE models on Cora and Citeseer datasets with 400 epochs, and PubMed dataset with 800 epochs. We train VGAE models on Cora and Citeseer datasets with 600 epochs and PubMed dataset with 800 epochs. All models are optimized by Adam optimizer with 0.01 learning rate. The λ is set to 4. For attributes perturbation, the ϵ is set to $5e-1$ and $1e-3$ on both Cora and Citeseer dataset, and 1 and $5e-3$ on PubMed dataset for L_2 and L_∞ adversarial training respectively. For adjacency matrix perturbation, the ϵ is set to $1e-3$ and $1e-1$ on Cora and CiteSeer, $1e-3$ and $3e-1$ on PubMed for L_2 and L_∞ adversarial training respectively. The steps T is set to 1. The α is set to $\frac{\epsilon}{T}$.

Likewise, for standard GAE and VGAE, we run the official Pytorch geometric code with 400 epochs for Citeseer and Cora datasets, 800 epochs for PubMed dataset.

Experimental Results. The results are showed in Table 3, Table 4 and Table 5. It can be seen that both L_2 and L_∞ adversarial trained models consistently outperform the standard trained models for all metrics. In particular, on Cora and Citeseer datasets, both L_2 and L_∞ adversarial training improve the performance with large margin for all metrics, i.e. at least +5.4% for GAE, +6.7% for VGAE on Cora dataset (Table 3), and at least +5.8% for GAE, +5.6% for VGAE on Citeseer dataset (Table 4).

5.3 Graph Anomaly Detection

We exactly follow [4] to conduct experiments for graph anomaly detection. In [4], the authors take reconstruction errors of attributes and links as the anomaly

Table 3. Results for node clustering on cora.

Methods	Acc (in%)	NMI (in%)	F1 (in%)	Precision (in%)	ARI (in%)
GAE	61.6 ± 3.4	44.9 ± 2.3	60.8 ± 3.4	62.5 ± 3.5	37.2 ± 3.2
AT-L2-GAE	67.0 ± 3.0	50.8 ± 1.7	66.6 ± 1.7	69.4 ± 1.7	44.1 ± 4.1
AT-Linf-GAE	67.1 ± 3.8	51.4 ± 1.9	67.5 ± 2.8	70.7 ± 2.2	43.4 ± 4.3
VGAE	58.7 ± 2.7	42.3 ± 2.2	57.3 ± 3.2	58.8 ± 3.5	34.6 ± 2.8
AT-L2-VGAE	67.3 ± 3.8	50.5 ± 2.1	66.1 ± 4.1	67.5 ± 3.8	44.3 ± 3.3
AT-Linf-VGAE	65.4 ± 2.3	49.5 ± 1.6	64.0 ± 2.3	65.8 ± 3.0	42.9 ± 2.8

Table 4. Results for node clustering on citeseer.

Methods	Acc (in%)	NMI (in%)	F1 (in%)	Precision (in%)	ARI (in%)
GAE	51.8 ± 2.6	28.0 ± 1.9	50.6 ± 3.1	55.1 ± 3.1	22.8 ± 2.3
AT-L2-GAE	61.6 ± 2.3	36.3 ± 1.4	58.8 ± 2.1	60.9 ± 1.4	34.6 ± 2.3
AT-Linf-GAE	60.2 ± 2.8	38.0 ± 2.3	57.0 ± 2.7	61.1 ± 1.6	34.1 ± 3.4
VGAE	53.6 ± 3.5	28.4 ± 3.3	51.1 ± 3.8	53.2 ± 4.1	26.1 ± 3.5
AT-L2-VGAE	59.2 ± 2.3	35.1 ± 2.3	57.3 ± 2.3	60.4 ± 3.1	33.0 ± 2.4
AT-Linf-VGAE	60.4 ± 1.5	36.5 ± 1.4	58.2 ± 1.4	61.1 ± 1.4	34.7 ± 2.0

Table 5. Results for node clustering on PubMed.

Methods	Acc (in%)	NMI (in%)	F1 (in%)	Precision (in%)	ARI (in%)
GAE	66.2 ± 2.0	27.9 ± 3.7	65.0 ± 2.3	68.8 ± 2.2	27.1 ± 3.3
AT-L2-GAE	67.5 ± 2.9	30.4 ± 5	66.7 ± 3.3	70.2 ± 3.1	28.9 ± 4.8
AT-Linf-GAE	68.4 ± 1.6	31.9 ± 3.2	67.7 ± 1.9	70.9 ± 1.8	30.2 ± 2.8
VGAE	67.5 ± 2.0	29.4 ± 3.2	66.5 ± 2.2	69.9 ± 2.2	28.4 ± 3.2
AT-L2-VGAE	69.8 ± 2.0	33.2 ± 3.4	69.4 ± 2.3	71.7 ± 2.5	32.5 ± 3.2
AT-Linf-VGAE	68.5 ± 1.2	30.7 ± 2.5	67.4 ± 1.5	70.1 ± 1.5	30.4 ± 2.0

scores. Specifically, the node with larger scores are more likely to be considered as anomalies.

Model Architecture. Different from link prediction and node clustering, the model architecture in graph anomaly detection not only contains structure reconstruction decoder, i.e. link reconstruction, but also contains attribute reconstruction decoder. We adopt the same model architecture as in the official code of [4] where the encoder is consisted of two GCN layers, and the decoder of structure reconstruction decoder is consisted of a GCN layer and a InnerProduction layer, and the decoder of attributes reconstruction decoder is consisted of two GCN layers.

Metrics. Following [4, 13], we use the area under the receiver operating characteristic curve (ROC-AUC) as the evaluation metric.

Parameter Settings. We set the α in anomaly scores to 0.5 where it balances the structure reconstruction errors and attributes reconstruction errors. We train the GAE model on Flickr, BlogCatalog and ACM datasets with 300 epochs. We set λ to 5. For adjacency matrix perturbation, we set ϵ to $3e-1$, $5e-5$ on both BlogCatalog and ACM datasets, $1e-3$ and $1e-6$ on Flickr dataset for L_∞ and L_2 adversarial training respectively. For attributes perturbations, we set ϵ to $1e-3$ on BlogCatalog for both L_∞ and L_2 adversarial training, $1e-3$ and $1e-2$ on ACM for L_∞ and L_2 adversarial training respectively, $5e-1$ and $3e-1$ on Flickr for L_∞ and L_2 adversarial training respectively. We set steps T to 1 and the α to $\frac{\epsilon}{T}$.

Anomaly Generation. Following [4], we inject two kinds of anomaly by perturbing structure and nodal attributes respectively:

- Structure anomalies. We randomly select s nodes from the network and then make those nodes fully connected, and then all the s nodes forming the clique are labeled as anomalies. t cliques are generated repeatedly and totally there are $s \times t$ structural anomalies.
- Attribute anomalies. We first randomly select $s \times t$ nodes as the attribute perturbation candidates. For each selected node v_i , we randomly select another k nodes from the network and calculate the Euclidean distance between v_i and all the k nodes. Then the node with largest distance is selected as v_j and the attributes of node v_j is changed to the attributes of v_i .

In this experiments, we set $s = 15$ and $t = 10, 15, 20$ for BlogCatalog, Flickr and ACM respectively which are the same to [4, 13].

Experimental Results. From Table 6, it can be seen that both L_2 and L_∞ adversarial training boost the performance in detecting anomalous nodes. Since adversarial training tend to learn feature representations that are less sensitive to perturbations in the inputs, we conjecture that the adversarial trained node embeddings are less influenced by the anomalous nodes, which helps the graph anomaly detection. A similar claim are also made in image domain [15] where they demonstrate adversarial training of autoencoders are beneficial to novelty detection.

Table 6. Results w.r.t. AUC (in%) for graph anomaly detection.

Methods	Flickr	BlogCatalog	ACM
GAE	80.2 ± 1.3	82.9 ± 0.3	72.5 ± 0.6
AT-L2-GAE	84.9 ± 0.2	84.7 ± 1.4	74.2 ± 1.7
AT-Linf-GAE	81.1 ± 1.1	82.8 ± 1.3	75.3 ± 0.9

6 Understanding Adversarial Training

In this section, we explore the impact of three hyper-parameters on the performance of GAE and VGAE with adversarial training, i.e. the ϵ , λ and T in generating \mathbf{A}' and \mathbf{X}' . These three hyper-parameters are commonly considered to control the strength of regularization for adversarial robustness [33]. Besides, we explore the relationship between the improvements achieved by adversarial training and node degree.

6.1 The Impact of ϵ

The experiments are conducted on link prediction and node clustering tasks based on Cora dataset. We fix ϵ to $5e-1$ and $1e-3$ on adjacency matrix perturbation for L_∞ and L_2 adversarial training respectively when vary ϵ on attributes perturbation. We fix ϵ to $1e-3$ and $3e-1$ on attributes perturbation for L_∞ and L_2 adversarial training respectively when vary ϵ on adjacency matrix perturbation.

The results are showed in Fig. 1. From Fig. 1, we can see that the performance are less sensitive to adjacency matrix perturbation and more sensitive to attributes perturbation. Besides, it can be seen that there is an increase and then a decrease trend when increasing ϵ for attributes perturbation. We conjecture that it is because too large perturbation on attributes may destroy useful information in attributes. Therefore, it is necessary to carefully adapt the perturbation magnitude ϵ when we apply adversarial training for improving the generalization of model.

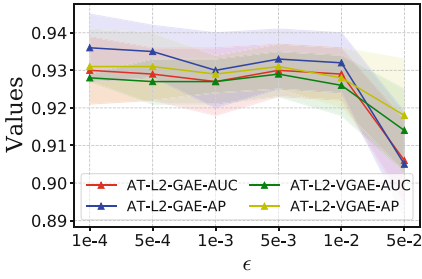
6.2 The Impact of T

The experiments are conducted on link prediction and node clustering tasks based on Cora dataset. For L_2 adversarial training, we set ϵ to $1e-3$ and $5e-1$ for adjacency matrix perturbation and attributes perturbation respectively. For L_∞ adversarial training, we set ϵ to $1e-1$ and $1e-3$ for adjacency matrix perturbation and attributes perturbation respectively. We set λ to 4.

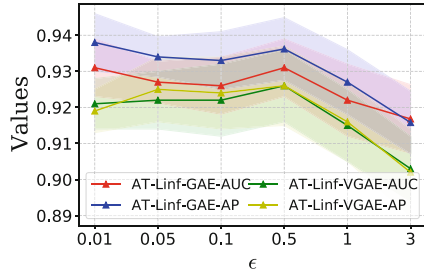
Results are showed in Fig. 2. From Fig. 2, we can see that there is a slightly drop on both link prediction and node clustering tasks when increasing T from 2 to 4, which implies that a big T is not helpful to improve the generalization of node embeddings learned by GAE and VGAE. We suggest that one step is good choice for generating adjacency matrix perturbation and attributes perturbation in both L_2 and L_∞ adversarial training.

6.3 The Impact of λ

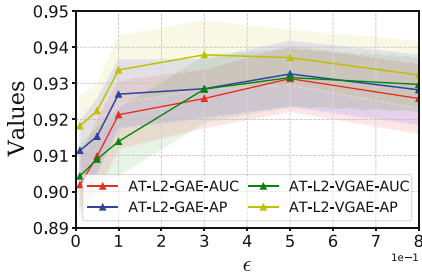
The experiments are conducted on link prediction and node clustering task based on Cora dataset. Likewise, for L_2 adversarial training, ϵ is set to $1e-3$ and $5e-1$ for adjacency matrix perturbation and attributes perturbation respectively. For L_∞ adversarial training, ϵ is set to $1e-1$ and $1e-3$ for adjacency matrix perturbation and attributes perturbation respectively. T is set to 1.



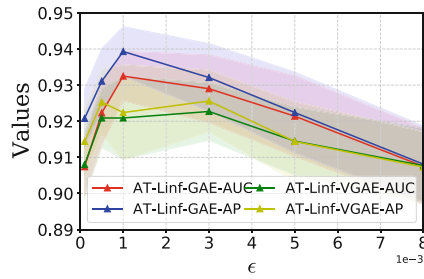
(a) Adjacency-Matrix-Perturbation



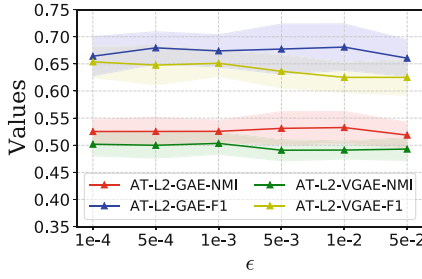
(b) Adjacency-Matrix-Perturbation



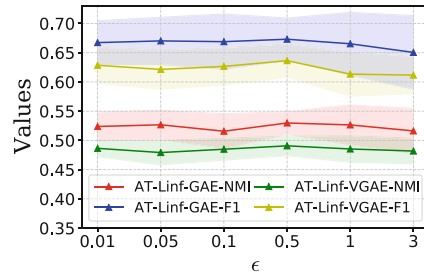
(c) Attributes-Perturbation



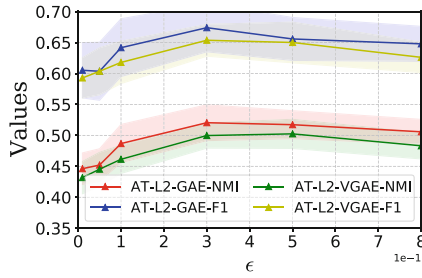
(d) Attributes-Perturbation



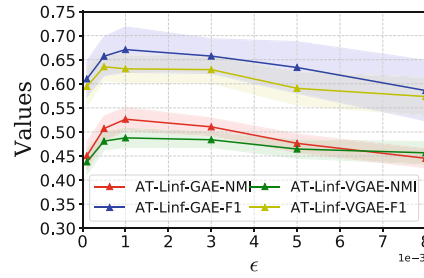
(e) Adjacency-Matrix-Perturbation



(f) Adjacency-Matrix-Perturbation



(g) Attributes-Perturbation



(h) Attributes-Perturbation

Fig. 1. The impact of ϵ in adjacency matrix perturbation and attributes perturbation. (a)–(d) show AUC/AP values for link prediction task and (e)–(h) show NMI/F1 values for node clustering task. Dots denote mean values with 30 repeated runs.

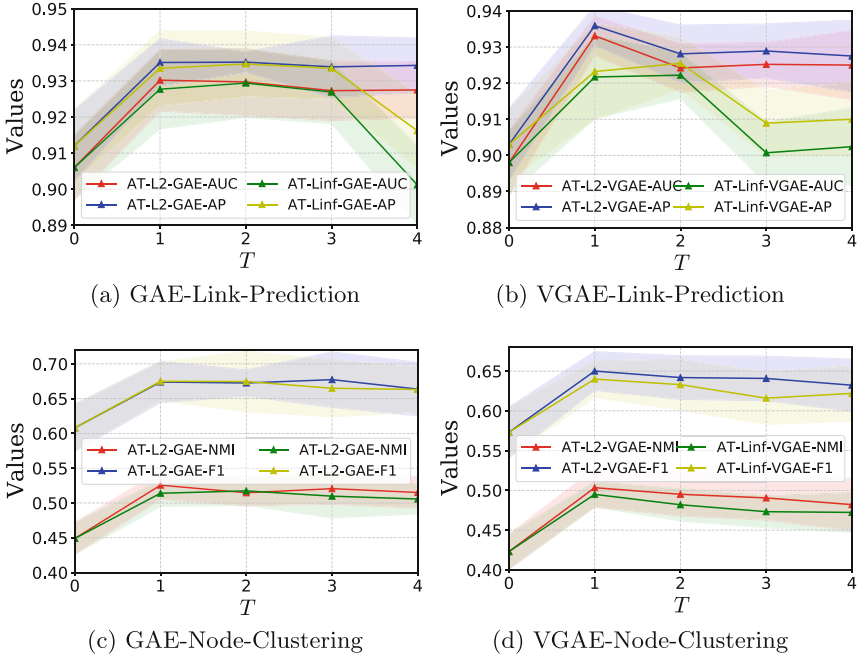


Fig. 2. The impact of steps T . Dots denote mean AUC/AP values for link prediction task and mean NMI/F1 values for node clustering task.

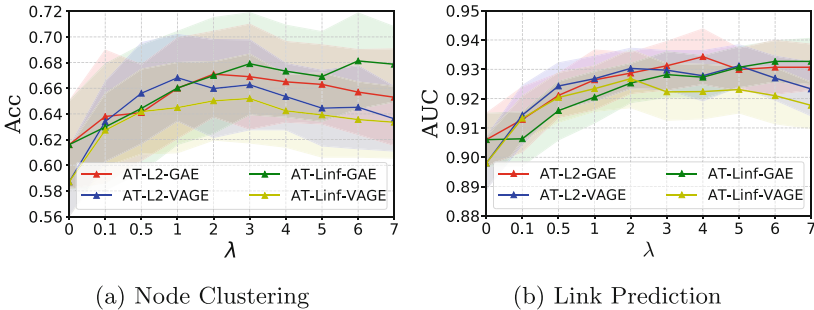


Fig. 3. The impact of λ . λ is varied from 0 to 7. Dots denote mean Acc for node clustering task and mean AUC for link prediction task. Experiments are conducted with 30 repeated runs.

Results are showed in Fig. 3. From Fig. 3, it can be seen that there is a significant increasing trend with the increase of λ , which indicates the effectiveness of both L_2 and L_∞ adversarial training in improving the generalization of GAE and VGAE. Besides, we also notice that a too large λ is not necessary and may lead to a negative effect in generalization of GAE and VGAE.

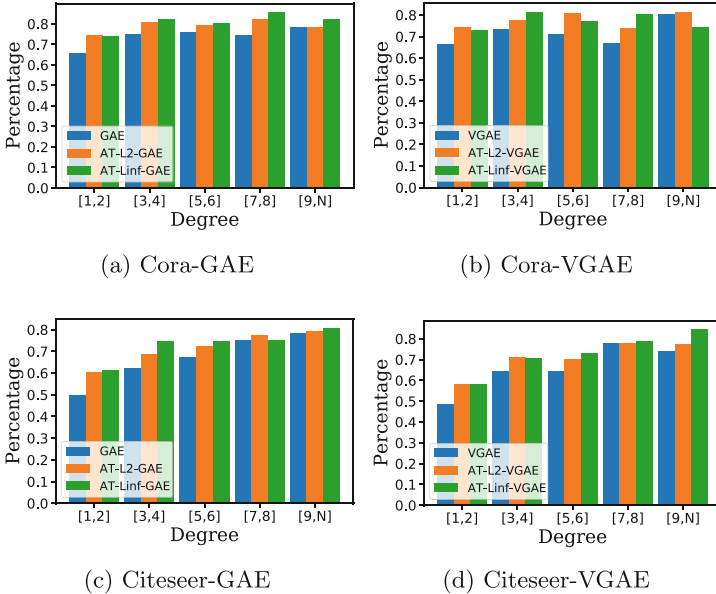


Fig. 4. Performance of GAE/VGAE and adversarial trained GAE/VGAE w.r.t. degrees in Cora and Citeseer datasets.

6.4 Performance w.r.t. Degree

In this section, we explore whether the performance of adversarial trained GAE/VGAE is sensitive to the degree of nodes. To conduct this experiments, we firstly learn node embeddings from Cora and Citeseer datasets by GAE/VGAE with L_2/L_∞ adversarial training and standard training respectively. The hyper-parameters are set the same as in the Node clustering task. Then we build a linear classification based on the learned node embeddings. Their accuracy can be found in Appendix. The accuracy with respect to degree distribution are showed in Fig. 4.

From Fig. 4, it can be seen that for most degree groups, both L_2 and L_∞ adversarial trained models outperform standard trained models, which indicates that both L_2 and L_∞ adversarial training improve the generalization of GAE and VGAE with different degrees. However, we also notice that adversarial training does not achieve a significant improvement on [9, N] group. We conjecture that it is because node embeddings with very large degrees already achieve a high generalization.

7 Conclusion

In this paper, we formulated L_2 and L_∞ adversarial training for GAE and VGAE, and studied their impact on the generalization performance. We conducted experiments on link prediction, node clustering and graph anomaly detection tasks.

The results show that both L_2 and L_∞ adversarial trained GAE and VGAE outperform GAE and VGAE with standard training. This indicates that L_2 and L_∞ adversarial training improve the generalization of GAE and VGAE. Besides, we showed that the generalization performance achieved by the L_2 and L_∞ adversarial training is more sensitive to attributes perturbation than adjacency matrix perturbation, and not sensitive to node degree. In addition, the parameter analysis suggest that a too large λ , ϵ and T would lead to a negative effect on the performance w.r.t. generalization.

References

1. Chen, J., Wu, Y., Lin, X., Xuan, Q.: Can adversarial network attack be defended? arXiv preprint [arXiv:1903.05994](https://arxiv.org/abs/1903.05994) (2019)
2. Dai, Q., Shen, X., Zhang, L., Li, Q., Wang, D.: Adversarial training methods for network embedding. In: The World Wide Web Conference, pp. 329–339 (2019)
3. Deng, Z., Dong, Y., Zhu, J.: Batch virtual adversarial training for graph convolutional networks. arXiv preprint [arXiv:1902.09192](https://arxiv.org/abs/1902.09192) (2019)
4. Ding, K., Li, J., Bhanushali, R., Liu, H.: Deep anomaly detection on attributed networks. In: Proceedings of the 2019 SIAM International Conference on Data Mining, pp. 594–602. SIAM (2019)
5. Feng, F., He, X., Tang, J., Chua, T.S.: Graph adversarial training: dynamically regularizing based on graph structure. IEEE Trans. Knowl. Data Eng. **PP**, 1 (2019)
6. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint [arXiv:1412.6572](https://arxiv.org/abs/1412.6572) (2014)
7. Grover, A., Leskovec, J.: node2vec: scalable feature learning for networks. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 855–864 (2016)
8. Jin, H., Zhang, X.: Latent adversarial training of graph convolution networks. In: ICMML Workshop on Learning and Reasoning with Graph-Structured Representations (2019)
9. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
10. Kipf, T.N., Welling, M.: Variational graph auto-encoders. arXiv preprint [arXiv:1611.07308](https://arxiv.org/abs/1611.07308) (2016)
11. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. arXiv preprint [arXiv:1706.06083](https://arxiv.org/abs/1706.06083) (2017)
12. Pan, S., Hu, R., Long, G., Jiang, J., Yao, L., Zhang, C.: Adversarially regularized graph autoencoder for graph embedding. arXiv preprint [arXiv:1802.04407](https://arxiv.org/abs/1802.04407) (2018)
13. Pei, Y., Huang, T., van Ipenburg, W., Pechenizkiy, M.: ResGCN: attention-based deep residual modeling for anomaly detection on attributed networks. arXiv preprint [arXiv:2009.14738](https://arxiv.org/abs/2009.14738) (2020)
14. Perozzi, B., Al-Rfou, R., Skiena, S.: DeepWalk: online learning of social representations. In: Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 701–710 (2014)
15. Salehi, M., ARAE: adversarially robust training of autoencoders improves novelty detection. arXiv preprint [arXiv:2003.05669](https://arxiv.org/abs/2003.05669) (2020)
16. Salha, G., Hennequin, R., Vazirgiannis, M.: Simple and effective graph autoencoders with one-hop linear models. arXiv preprint [arXiv:2001.07614](https://arxiv.org/abs/2001.07614) (2020)

17. Salha, G., Limnios, S., Hennequin, R., Tran, V.A., Vazirgiannis, M.: Gravity-inspired graph autoencoders for directed link prediction. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management, pp. 589–598 (2019)
18. Salman, H., Ilyas, A., Engstrom, L., Kapoor, A., Madry, A.: Do adversarially robust ImageNet models transfer better? arXiv preprint [arXiv:2007.08489](https://arxiv.org/abs/2007.08489) (2020)
19. Schlichtkrull, M., Kipf, T.N., Bloem, P., van den Berg, R., Titov, I., Welling, M.: Modeling relational data with graph convolutional networks. In: Gangemi, A., Navigli, R., Vidal, M.-E., Hitzler, P., Troncy, R., Hollink, L., Tordai, A., Alam, M. (eds.) ESWC 2018. LNCS, vol. 10843, pp. 593–607. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-93417-4_38
20. Shi, H., Fan, H., Kwok, J.T.: Effective decoding in graph auto-encoder using triadic closure. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, pp. 906–913 (2020)
21. Stutz, D., Hein, M., Schiele, B.: Disentangling adversarial robustness and generalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 6976–6987 (2019)
22. Sun, K., Lin, Z., Guo, H., Zhu, Z.: Virtual adversarial training on graph convolutional networks in node classification. In: Lin, Z., et al. (eds.) PRCV 2019. LNCS, vol. 11857, pp. 431–443. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-31654-9_37
23. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: Line: large-scale information network embedding. In: Proceedings of the 24th International Conference on World Wide Web, pp. 1067–1077 (2015)
24. Tian, F., Gao, B., Cui, Q., Chen, E., Liu, T.Y.: Learning deep representations for graph clustering. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 28 (2014)
25. Tsipras, D., Santurkar, S., Engstrom, L., Turner, A., Madry, A.: Robustness may be at odds with accuracy. arXiv preprint [arXiv:1805.12152](https://arxiv.org/abs/1805.12152) (2018)
26. Utrera, F., Kravitz, E., Erichson, N.B., Khanna, R., Mahoney, M.W.: Adversarially-trained deep nets transfer better. arXiv preprint [arXiv:2007.05869](https://arxiv.org/abs/2007.05869) (2020)
27. Wang, D., Cui, P., Zhu, W.: Structural deep network embedding. In: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1225–1234 (2016)
28. Wang, X., Liu, X., Hsieh, C.J.: GraphDefense: towards robust graph convolutional networks. arXiv preprint [arXiv:1911.04429](https://arxiv.org/abs/1911.04429) (2019)
29. Wang, Y., Zou, D., Yi, J., Bailey, J., Ma, X., Gu, Q.: Improving adversarial robustness requires revisiting misclassified examples. In: International Conference on Learning Representations (2019)
30. Xia, R., Pan, Y., Du, L., Yin, J.: Robust multi-view spectral clustering via low-rank and sparse decomposition. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 28 (2014)
31. Xu, K., et al.: Topology attack and defense for graph neural networks: an optimization perspective. arXiv preprint [arXiv:1906.04214](https://arxiv.org/abs/1906.04214) (2019)
32. Yu, D., Zhang, R., Jiang, Z., Wu, Y., Yang, Y.: Graph-revised convolutional network. arXiv preprint [arXiv:1911.07123](https://arxiv.org/abs/1911.07123) (2019)
33. Zhang, H., Yu, Y., Jiao, J., Xing, E., El Ghaoui, L., Jordan, M.: Theoretically principled trade-off between robustness and accuracy. In: International Conference on Machine Learning, pp. 7472–7482. PMLR (2019)