A Systematic Evaluation of Node Embedding Robustness

Anonymous Author(s)
Anonymous Affiliation
Anonymous Email

Abstract

Node embedding methods map network nodes to low dimensional vectors that can be subsequently used in a variety of downstream prediction tasks. The popularity of these methods has significantly increased in recent years, yet, their robustness to perturbations of the input data is still poorly understood. In this paper, we assess the empirical robustness of node embedding models to random and adversarial poisoning attacks. Our systematic evaluation covers representative embedding methods based on Skip-Gram, matrix factorization, and deep neural networks. We compare edge addition, deletion and rewiring strategies computed using network properties as well as node labels. We also investigate the effect of label homophily and heterophily on robustness. We report qualitative results via embedding visualization and quantitative results in terms of downstream node classification and network reconstruction performances. We found that node classification suffers from higher performance degradation as opposed to network reconstruction, and that degree-based and label-based attacks are on average the most damaging.

1 Introduction

In recent years, the design of robust machine learning models has become an important topic and attracted significant amounts of research attention [1–4]. The term 'robust' refers to the ability of a model to provide consistent and accurate predictions under small perturbations in the input data. These perturbations can appear in the form of random noise, out of distribution (OOD) data, or partially observed inputs [5]. They can affect models at train or evaluation times and be random or adversarial in nature. For a more complete overview of robustness in machine learning we refer the reader to [6]. In this manuscript, we empirically study both random and adversarial attack scenarios where perturbations are either a consequence of noise or specifically crafted to reduce model performance. We further focus our analysis on attacks affecting the models at training time exclusively, also know as the poisoning scenario [7].

Simultaneously, node representation learning or node embedding models have become increasingly popular for bridging the gap between traditional machine learning and network structured data [8–10]. These approaches map network nodes to real-valued vectors that can be subsequently used in downstream prediction tasks, such as classification [11] and regression [12]. Training of these models can be performed in a semi-supervised or unsupervised fashion. In the former, embeddings are optimized for a particular downstream task while in the latter, general purpose embeddings are obtained. Robustness is an important feature for representation learning models as well. One would generally prefer small changes in the input networks to have a minimal impact on the vector representations learned and subsequently on downstream performance. Moreover, with the deployment of these models in safety-critical environments (e.g., [13]) and on the web (where adversaries are common) [14, 15], robustness evaluation has become ever more essential. Unfortunately, the robustness of unsupervised node embedding approaches is poorly understood. Some recent studies (e.g., [16]) have analyzed specific cases of semi-supervised models based on the Graph Neural Network or particular shallow models (e.g., [17]). Other studies evaluate the performance of unsupervised random walk approaches under poison attacks [7]. However, methods from other paradigms for unsupervised network embeddings, such as matrix factorization and deep neural models, have not received much attention yet.

Submitted to the First Learning on Graphs Conference (LoG 2022). Do not distribute.

We perform a systematic empirical analysis of the robustness of foundational works in the field of node embeddings. The 10 unsupervised approaches we evaluate include Node2vec [18], GraRep [19], and SDNE [20], which have inspired many other methods based on the same principles, e.g., [21–23]. The evaluated models can be categorized into Skip-Gram, matrix factorization, and deep neural approaches, and we compare their robustness on two important downstream tasks: node classification and network reconstruction. We evaluate robustness under randomized and adversarial attacks targeting the network edges. Within the adversarial attacks we direct our analysis to heuristic-based approaches, as opposed to optimization based attacks. The former heuristically target network properties such as assortativity or degree. These attacks are significantly more computationally efficient that optimization attacks, where an optimization problem must be solved to identify the most damaging changes to the network. Additionally, heuristic-based attacks provide more interpretable results in terms of the edges being targeted and have also been shown to effectively lead to structural collapse in networks [24]. Further, we focus our evaluation on the important global attack scenario, where changes can be performed anywhere on the graph structure provided a fixed attack budget.

Contributions. Our main contribution is a systematic analysis of node embedding robustness. We evaluate a total of 10 unsupervised node embedding approaches based on three different learning paradigms. We employ a total of 6 small and mid-sized networks and compare 14 different poison attack strategies. Further, we investigate the differences between randomized and adversarial attacks and compare edge addition, deletion and rewiring strategies. We also investigate adversarial robustness under node label homophily, where nodes with similar labels tend to be connected to each other, and heterophily, where nodes of different labels are more often connected. This constitutes the first empirical evaluation of this magnitude on node embedding robustness.

The remainder is organized as follows: in Section 2 we present the related work and in Section 3 we introduce the embedding methods and attack strategies evaluated. In Section 4, we discuss the experimental evaluation and results and finally, in Section 5 we outline our main conclusions.

2 Related Work

A large body of research has shown that traditional machine learning models and more recently deep neural models can be easily misled into providing wrong answers with high confidence [25, 26]. Work on identifying and protecting against these adversarial attacks has particularly developed in the computer vision field. Here, several works including [27, 28], have shown how changes unperceivable to the human eye can result in dramatic performances drops or misclassifications. Later, adversarial attacks were introduced in the field of network science [5]. In [24], the authors show how structural properties of networks can collapse as a result of attacks. The authors further provide a framework for simulating attacks and defenses on networks. With the popularization of node embedding methods authors have also investigated adversarial attacks on semi-supervised [16, 29] and unsupervised [17] approaches. While there are some empirical studies comparing the performance of these types of methods (e.g., [30]), there is little research comparing their robustness. With the present work, our aim is to fill this gap and provide a fist empirical study and overview on the robustness to random and adversarial attacks of unsupervised node embedding approaches.

3 Methods

In this section we introduce the node embeddings approaches evaluated and the attack strategies used to poison the input networks. Regarding notation, in what follows we will use $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ to refer to an undirected graph with vertex set $\mathbf{V} = \{v_1, \ldots, v_N\}$, $N = |\mathbf{V}|$ and edge set $\mathbf{E} \subseteq (\mathbf{V} \times \mathbf{V})$, $M = |\mathbf{E}|$. We will represent edges or connected node-pairs as unordered pairs $\{v_i, v_j\} \in \mathbf{E}$. And refer to pairs $\{v_i, v_j\} \notin \mathbf{E}$ as non-edges or unconnected node-pairs. Node embeddings are denoted as $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N)$, $\mathbf{X} \in \mathbb{R}^{\mathrm{N} \times \mathrm{d}}$ where $\mathbf{x}_i$ is the d-dimensional vector representation corresponding to node v_i.

3.1 Node embedding methods

For our experimental evaluation we have selected 10 representative methods spanning three different embedding learning paradigms, namely Skip-Gram, matrix factorization and deep neural networks. Next, we introduce each paradigm and the corresponding methods.

Table 1: Poison attacks evaluated and their types: (D) deterministic, (ND) non-deterministic.

Edge addition		Edge deletion		Edge rewiring	
Name	Type	Name	Type	Name	Type
add_rand	ND	del_rand	ND	rew_rand	ND
add_deg	ND	del_deg	D	-	-
add_pa	ND	del_pa	D	-	-
add_da	ND	del_da	D	-	-
add_dd	ND	del_dd	D	-	-
add_ce	ND	del_di	ND	DICE	ND

Skip-Gram. These approaches capture node similarities in the graph through random walks and leverage the Skip-Gram model [31] to obtain node representations that maximize the posterior probability of observing neighboring nodes in the walks. From this category we evaluate: Deepwalk [32], the seminal work that proposed fixed length random walks to capture node similarities and Skip-Gram (approximated via hierarchical softmax) for learning the embedding matrix $\mathbf{X}$; Node2vec [18], which introduced more flexible random walks controlled by in/out and return parameters and approximates Skip-Gram via negative sampling; LINE [33], where the authors leverage first and second order proximities to learn representations; And finally, VERSE [11], which minimizes the KL-divergence between a similarity metric on $\mathbf{G}$ (by default Personalized PageRank) and a vector similarity on $\mathbf{X}$.

Matrix Factorization. Factorization methods take as input node similarities encoded in the graph Laplacian, incidence matrices, adjacency matrices ($\mathbf{A}$) and their polynomials, etc. and compute low dimensional embeddings by factorizing the selected matrix. We evaluate the following methods that are based on this paradigm: GraRep [19], HOPE [34], NetMF [35] and M-NMF [36]. GraRep factorizes high order polynomials of $\mathbf{A}$, HOPE can factorize different similarity matrices provided they can be expressed as a composition of two sparse proximity matrices. NetMF decomposes the DeepWalk transition matrix via SVD and lastly, M-NMF computes embeddings via non-negative matrix factorization and incorporates community structure in this process.

Deep Neural Networks. Deep neural models, from auto-encoders to Siamese networks or CNNs, have also been used to obtain node representations from a graph's link structure in an unsupervised fashion. Among these types of methods we evaluate SDNE [20], a deep neural model that captures first and second order proximity in the graph, and PRUNE [12], a Siamese network architecture able to preserve global node ranking and community structure.

3.2 Network attacks

We subdivide network attacks into randomized and adversarial and further into three main types based on the changes to the network structure. These changes are edge addition, edge deletion and edge rewiring. Table 1 summarizes all attacks and below we briefly describe each.

Randomized Attacks. These attacks are designed to simulate random errors or failures in the networks. We consider edge addition (*add_rand*), deletion (*del_rand*) and rewiring (*rew_rand*). In the first case, pairs of nodes, $v_i, v_j \in \mathbf{V}$ are selected uniformly at random and added to $\mathbf{E}$ iff $v_i \neq v_j$ and $\{v_i, v_j\} \notin \mathbf{E}$. For deletion attacks edges $\{v_i, v_j\} \in \mathbf{E}$ are selected uniformly at random and removed from $\mathbf{E}$ iff $d_i \geq 2 \wedge d_j \geq 2$. Here d_i and d_j represent the degrees of node v_i and v_j, respectively. In rewire attacks we use *del_rand* to remove a certain number of edges $\{v_i, v_j\} \in \mathbf{E}$ and then reconnect each v_i to a new node v_k such that $v_k \neq v_j$ and $\{v_i, v_k\} \notin \mathbf{E}$.

Adversarial Attacks. We also consider a particular type of heuristic-based adversarial attacks which target specific network properties such as node degrees, network assortativity, and node labels. As already discussed, the importance of these attacks steams from their higher computational efficiency, explainability and capacity to collapse the network structure. Other optimization-based adversarial attacks are beyond the scope of the present paper.

For all edge addition attacks we ensure that the newly generated pairs do not exist already in the graph, i.e., $\{v_i, v_j\} \notin \mathbf{E}$, and that they do not form selfloops, i.e., $v_i \neq v_j$. For the degree-based

(*add_deg*) and preferential attachment (*add_pa*) edge addition strategies we sample nodes uniformly and based on degree, respectively, and connect them to destination nodes sampled based on degree. For the degree assortativity (*add_da*) and disassortativity (*add_dd*) attacks, we generate edges which increase and respectively, decrease this property. We use the definition of assortativity from [37] (Eq. (21)). For *add_da* we generate edges $\{v_i, v_j\}$ where $d_i \approx d_j$ and for *add_dd* edges where $d_i \nsim d_j$. The *add_ce* strategy applies only to attributed graphs and adds a set of random edges connecting nodes of dissimilar labels, exclusively.

Unless otherwise specified, edge deletion attacks ensure that the input networks do not become disconnected after the attack. For *del_deg* and *del_pa* we first sort all edges based on the appropriate metric, i.e., $d_i + d_j$ for *del_deg* and $d_i \times d_j$ for *del_pa*, and later remove the top edges that can be removed without disconnecting the network. For *del_da* and *del_dd* we compute the assortativity contribution of each edge according to Eq. (21) in [37] and once again take the top candidates while avoiding disconnections. Lastly the *add_di* strategy applies exclusively to attributed graphs and randomly selects edges for removal where the incident nodes share the same label.

Finally, *DICE* [7] is an adversarial attack where one randomly decides if an edge will be removed or added to the network with equal probability. If an edge is to be removed, this is done according to the *add_di* strategy, and if it is to be added, it is done following *add_ce*.

It is important to note that all edge deletion attacks with the exception of *add_di* are deterministic while the remaining addition and rewire attacks are non-deterministic.

4 Experiments

In this section we present the experimental setup, networks used and the results obtained. All our experiments were carried out on a single machine equipped with two 12 Core Intel(R) Xeon(R) Gold processors, 1TB of RAM and an RTX 3090 GPU.

To ensure reproducibility of results, we have employed and extended the capabilities of the EvalNE toolbox [38]. This Python framework allows users to asses the performance and robustness of network embedding approaches for downstream node classification, network reconstruction, link prediction and sign prediction. In the framework we have integrated a variety of random and adversarial poison attack strategies, including those introduced in Section 3.2 and Table 1. In EvalNE, complete evaluation pipelines and hyperparameters are specified through configuration files which can be used at any time to replicate results. These configuration files together with the rest of our code are available online at `https://tinyurl.com/5n8tsmrs`.

4.1 Preliminaries and Setup

As pointed out in Section 1, the main goal of this paper is to investigate the robustness of node embedding approaches to poison attacks. To this end we report changes in downstream node classification and network reconstruction performances for different attacks on the input graphs. Next, we summarize the main goals and evaluation pipelines for both tasks and the overall evaluation setup.

Node Classification. Given an input graph and labels for a subset of the vertices, the goal in node classification is to infer the labels of the remaining vertices. To evaluate node classification robustness we proceed as follows. (1) We start by attacking an input network $\mathbf{G}$ with a specific strategy (from Table 1) and budget b. The budget defines the number of edges an attacker can add, delete or rewire in the network, expressed as a fraction of the total edges. For example, $b = 0.1$ indicates 10% of all edges in $\mathbf{E}$. (2) The attacked network $\hat{\mathbf{G}} = (\mathbf{V}, \hat{\mathbf{E}})$ is then provided as input to different node embedding approaches, which yield a representation matrix $\mathbf{X}$ containing the vertex representations as its rows. As shown by Mara et. al. [30], gains from optimizing the hyperparameters of these models are marginal, and thus, we resort to fixed default values [1]. We also fix the embedding dimensionality $d = 128$. (3) Given a number of training nodes N_{tr} (also defined as a fraction of all nodes in $\mathbf{V}$), a multi-class one-versus-rest logistic regression model with 5-fold cross validation is trained to predict node labels from node representations. (4) We repeat the previous step 3 times with different node samples and report average results. For some experiments we will report results independent of

[1] Exact hyperparameter values for each method as well as the implementations used are reported in our EvalNE configuration files.

Table 2: Main statistics of the networks used for evaluation. The average degree is indicated by $\langle k \rangle$ and the assortativity coefficient by r.

Network	Type	Task	# Nodes	# Edges	# Labels	$\langle k \rangle$	r
Citeseer	Citation	NC	2110	3668	6	3.48	0.01
Cora	Citation	NC	2485	5069	7	4.08	-0.07
PolBlogs	Web	NR	1222	16714	-	27.35	-0.22
Facebook	Social	NR	4039	88234	-	43.69	0.06
IIP	Collaboration	Viz	219	630	3	5.75	-0.22
StudentDB	Relational	Viz	395	3423	7	17.33	-0.34

the value of N_{tr}. In these cases we additionally average results over several values of N_{tr}. (5). Finally, and unless otherwise specified, for the non-deterministic attacks listed in Table 1 we repeat the complete process 3 times with varying random seeds resulting in different sets of edges being removed in step 1). We report node classification performance in terms of f1_micro and f1_macro.

Network Reconstruction. In this task the aim is to investigate how well the link structure of an input network can be recovered from the node representations. To this end node representations are first learned from the input network. Then, node pair representations are derived by applying a binary operator on the node representations. Finally, a binary classifier is trained to discriminate edges from non-edges. High quality representations are expected to result in the classifier scores of edges being higher than those of non-edges.

We evaluate robustness on this task akin to node classification. (1) We attack the input network $\mathbf{G}$ with a given strategy and budget b. (2) We compute node representations for $\hat{\mathbf{G}}$ with different methods for which we use fixed default hyperparameters. (3) Representations of node pairs $\{v_i, v_j\}$ are combined into node-pair representations using the Hadamard product, i.e., $\mathbf{x}_{i,j} = \mathbf{x}_i \cdot \mathbf{x}_j$. [2] (4) A binary Logistic Regression with 5-fold cross validation is trained using representations corresponding to edges and non-edges in $\hat{\mathbf{G}}$. (5) The classifiers performance is tested using representations of edges and non-edges of the original unattacked graph $\mathbf{G}$. For computational efficiency, we approximate the performance using 5% of all possible node-pairs in $\mathbf{G}$. (6) We again repeat the complete process 3 times for non-deterministic attacks. For this task we report AUC and average precision scores.

Experimental Setup. Our evaluation setup is structured as follows. First, in Section 4.3.1 we investigate the performance of node embedding approaches under random attacks. In this case, we use the *add_rand* and *del_rand* strategies and vary the attack budget $b \in [0.1, 0.2, ..., 0.9]$. For node classification specifically, we report average results over $N_{tr} \in [0.1, 0.5, 0.9]$, 3 node shuffles for each N_{tr} value, and 3 experiment repetitions for non-deterministic attacks. For network reconstruction we only perform the 3 experiment repetitions for non-deterministic attacks. We then also investigate the effect of the number of labeled nodes for node classification by comparing the results obtained for $N_{tr} = 0.1$ to $N_{tr} = 0.5$ and $N_{tr} = 0.9$. Second, in Section 4.3.2 we evaluate adversarial robustness. We use a similar setup with the following exceptions: we compare all attacks from Table 1 (random attacks are used as baselines) and the budget is fixed to $b = 0.2$. Third, in Section 4.3.3 we compare addition, deletion and rewiring attacks. For both downstream tasks we compare *add_rand*, *del_rand* and *rew_rand* and for node classification we additionally compare *add_ce*, *del_di* and *DICE*. Other parameters are set as for the adversarial attack experiment. In this section we also investigate differences between deletion attacks that disconnect and those that do not disconnect the input networks. Lastly, in Section 4.3.4 we compare adversarial attacks on node classification under homophily and heterophily of node labels.

4.2 Data

To conduct our experiments we use a total of 6 small and mid sized networks from different domains. Specifically, for node classification we use Citeseer [39] and Cora [40], two citation networks where nodes denote publications, edges represent citations between them and node labels indicate the main research field of each paper. For network reconstruction we use PolBlogs [41], a network of political blogs connected to each other via hyperlinks, and Facebook [42], a network of individuals and

[2] With the exception of PRUNE, where we use $\mathbf{x}_{i,j} = (\mathbf{x}_i + \mathbf{x}_j)/2$.

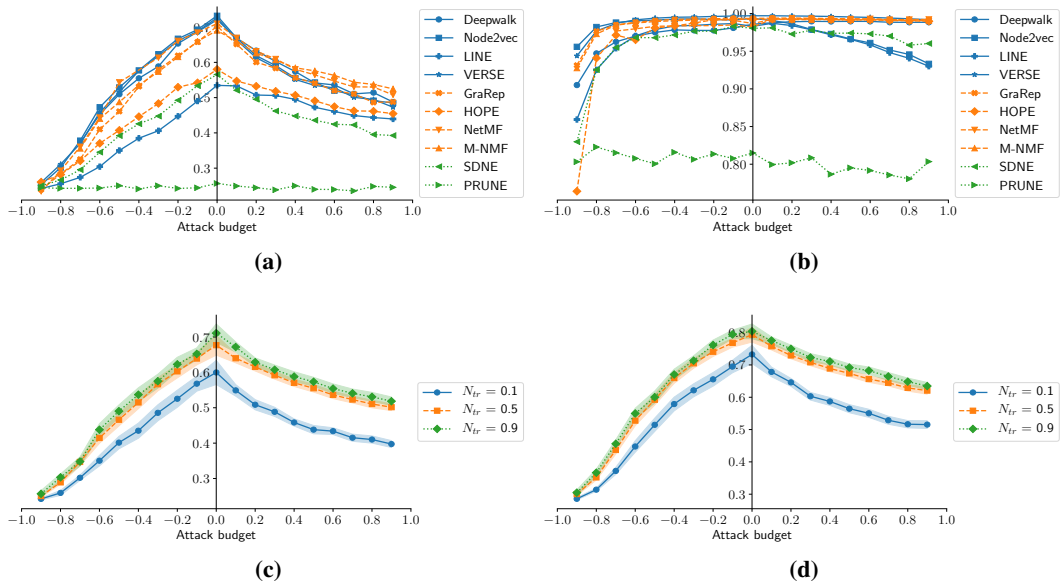

Figure 1: Robustness to randomized attacks for different budget values. The x-axis shows budgets as a fraction of all edges in the graph. Negative values represent edge deletion and positives edge addition attacks. Figure 1a presents f1_micro scores for node classification on Citeseer. Figure 1b shows AUCs for network reconstruction of Facebook. In Figures 1c and 1d we show average node classification performances for different fractions of labeled nodes N_{tr} on Citeseer and Cora, respectively. Shaded areas denote 95% confidence intervals and the y-axis present f1_micro scores.

their social relations on the platform. Lastly, we perform qualitative and visualization experiments on the internet industry partnership (IIP) [43] and the StudentDB [44] networks. In the former, nodes represent companies, edges represent relations such as alliance or partnership and node labels indicate the company's main business area, i.e., user content, infrastructure or commerce. The latter, StudentDB, is a k-partite network representing a snapshot of the Antwerp University relational database. Nodes represent entities such as students, courses, tracks, etc., and edges are binary relations, e.g., student-in-track, course-in-track, etc.[3] Node labels indicate the type of each nodes. In Table 2 we summarize the main statistics of each network.

4.3 Experimental Results

4.3.1 Randomized attacks

We start in Figure 1a with the node classification performance under random edge attacks and varying attack budgets. In the chart, negative budget values indicate edge deletion and positives indicate edge addition. In this case we allow edge deletions to disconnect the original networks. We report f1_micro scores for the Citeseer network (f1_macro results as well as those for the Cora network are similar and provided in Appendix A.2). From the figure we first note different general behaviors for edge deletion and addition attacks. Deletions cause a consistent performance degradation until complete collapse at 90% edge deletion. Addition presents a sharper loss in performance for relatively low budget values ($b \leq 0.2$) which starts to become less severe around ($b = 0.4$). Thus, in the low budget regime $-0.2 < b < 0.2$, commonly analyzed in the literature, edge addition attacks degrade performance more than edge deletions. Outside of this range, however, edge deletions are more severe. From Figure 1a, we also observe that Skip-Gram methods are in general more robust to edge addition attacks than other approaches while for edge deletion the results are similar across the board.

In Figure 1b we present the AUC scores for reconstructing the original Facebook network $\mathbf{G}$, from the attacked graph $\hat{\mathbf{G}}$. The plot indicates almost perfect edge recovery under random attacks with AUCs close to 1. Most of the evaluated methods show high robustness for a wide range of budget values.

[3]Further details on the IIP and StudentDB networks are provided in Appendix A.1.

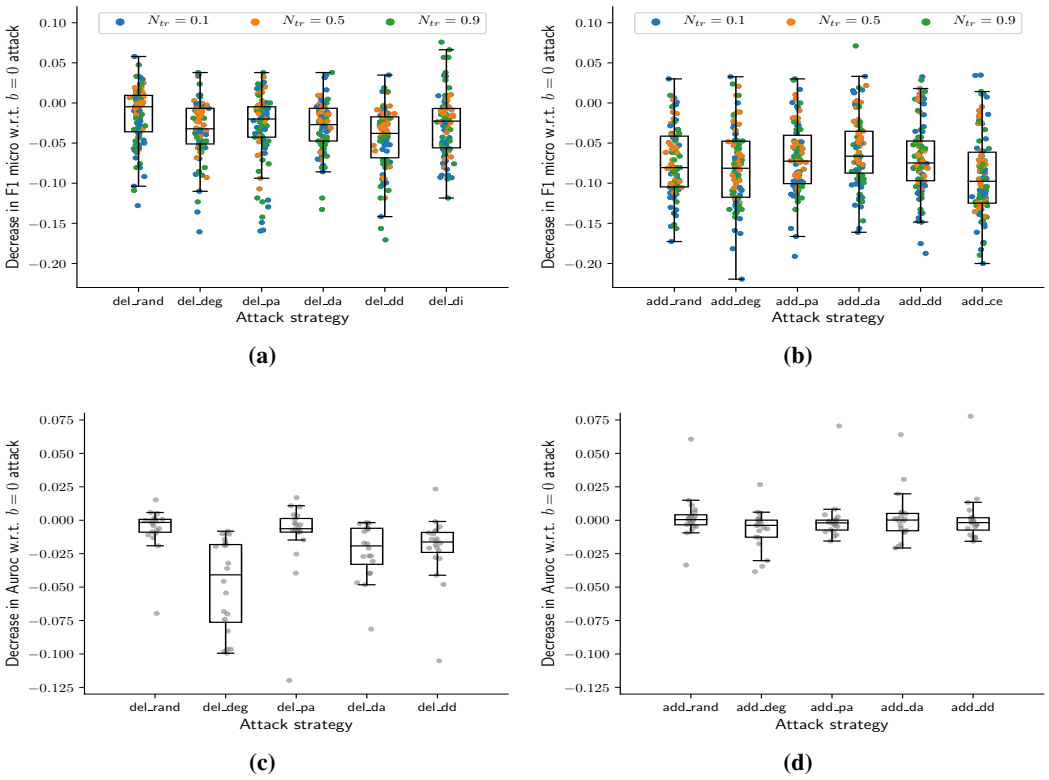

Figure 2: Comparison of adversarial edge deletion and addition attacks on node classification and network reconstruction for $b = 0.2$. Figures 2a and 2b show deletion and addition attacks on node classification for Citeseer. Colors indicate the fraction of train nodes N_{tr}. Figures 2a and 2b show similar results for network reconstruction on both Facebook and PolBlogs networks combined.

Some notable exceptions are Node2vec, LINE and SDNE, which consistently lose performance as more edges are added to the network. For the PolBlogs dataset presented in Appendix A.2, we observe similar robustness to edge addition and deletion. A notable exception in this case is HOPE which significantly degrades performance for strong edge deletion attacks $b \leq -0.6$. The overall high robustness of the evaluated methods on network reconstruction is an interesting finding particularly given the fact that attacks on this downstream task affect methods twice. First at embedding learning time and later while training the binary classifier (the edge and non-edge train labels are extracted from the attacked graph $\hat{\mathbf{G}}$).

We now focus our attention to the impact of the number of labeled nodes available for node classification (N_{tr}). In Figures 1c and 1d we compare the average performance over all methods and experiment repetitions for $N_{tr} \in [0.1, 0.5, 0.9]$. For both Citeseer and Cora we observe similarly low performances when only a relatively small amount of labeled nodes are available i.e., $N_{tr} = 0.1$. For larger values ($N_{tr} \geq 0.5$) the performances are very similar. We also observe that as networks become denser (as we move right on the x-axis in each plot) the difference between low and high values of N_{tr} become more significant. This indicates that node embedding methods will generally not provide robust predictions when few labeled nodes are available and this situation will worsen the denser the network is.

4.3.2 Adversarial attacks

We now compare the effect of different heuristic-based adversarial attacks on node classification. Figures 2a and 2b summarize the results on the Citeseer network for edge deletion and addition attacks, respectively. In both cases we present decreases in f1_micro caused by different attacks with budget $b = 0.2$, as compared to the performance on the non-attacked graph. Firstly, if we compare

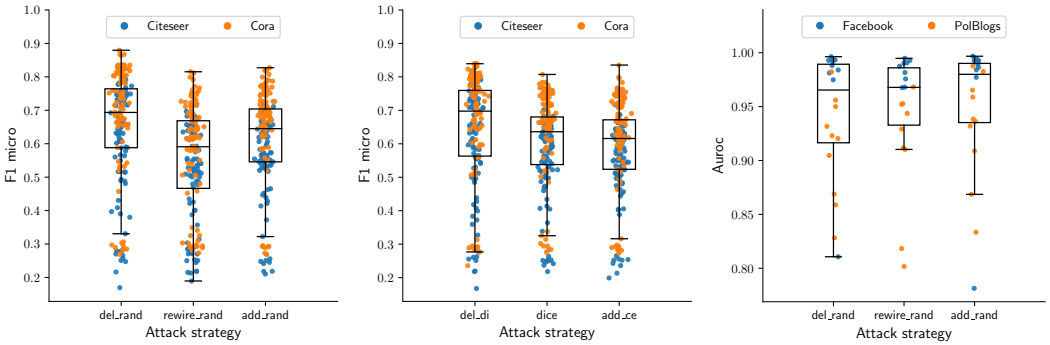

Figure 3: Comparison of edge addition, rewiring and deletion attacks for both downstream tasks. The leftmost and center figures present f1_micro scores for random and node label based attacks on node classification. The rightmost figure shows AUC results for random attacks on network reconstruction.

across graphs we observe that edge additions decrease performance more than deletions across all methods for this particular budget value. This is also consistent with our observations from Figure 1a for random attacks on node classification. Among the edge deletion attacks we see that *del_dd* is, from an adversarial perspective, the most effective strategy. With this attack, we are targeting edges from high degree to low degree nodes further increasing the uncertainty regarding the latter. On the other hand, for edge addition the most effective strategies are connecting edges with different labels together (*add_ce*) or connecting nodes with similar degrees to each other (*add_deg*). It is interesting to note that attacks with full knowledge of the node labels *del_di* and *add_ce* are not significantly stronger than others e.g., degree based attacks. The colors in both figures indicate different fractions of labeled nodes. We observe that most of the variance in performance comes from the experiments with $N_{tr} = 0.1$ (blue points) and that these are also mostly concentrated in the lower ends of the boxplots. The results for $N_{tr} \geq 0.5$ are in general very similar.

In Figures 2c and 2d we present similar results for network reconstruction. In this case we show the combined performances for both Facebook and PolBlogs datasets. The experiments reveal that edge deletion attacks are marginally stronger than edge additions. In particular, deleting edges based on degree is the most effective adversarial technique of the ones we have evaluated. Overall, we also observe much less variance in performance compared to the results on node classification.

4.3.3 Other attacks

In Figure 3 we compare edge addition, rewiring and deletion attacks on both downstream tasks. The attack budget is fixed to 0.2 and we show combined results for the two networks used in each task (marker color denotes the data used). We observe that for node classification rewiring attacks perform best (central boxes in the left and middle plots in Figure 3). This is also the case if we look at each individual dataset with results for Cora (orange dots) being significantly higher than those on Citeseer (blue dots). For network reconstruction we have much less data available, considering that we do not need to test different train sizes and shuffles per size. In this case the results indicate similar performances for all attack types. We further observe that results on the Facebook network are overall higher than on PolBlogs. The f1_macro and average precision scores for each task also corroborate this findings and are presented in Appendix A.3.

We further investigate how strong a role network connectivity plays in adversarial attacks. We compare random and degree attacks constrained to not disconnecting the input networks and their unconstrained counterparts. We find that constrained attacks are on average, over all methods and networks 5% less effective. Specifically, for random attacks the f1_micro performance without disconnections is 0.651 ± 0.166 (mean and standard deviation) and with disconnections 0.612 ± 0.161. Similarly, for degree based attacks average performance reaches 0.637 ± 0.163 when disconnections are prevented and 0.606 ± 0.164 when they are not.

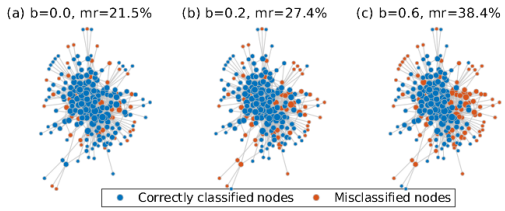

Figure 4: Correctly and incorrectly classified nodes for the homophilic IIP network for varying attack budgets.

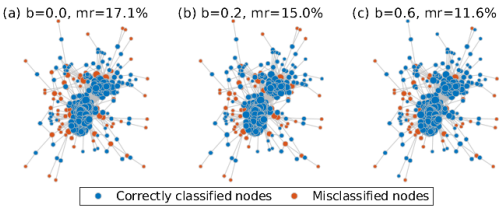

Figure 5: Correctly and incorrectly classified nodes for the heterophilic StudentDB network for varying attack budgets.

4.3.4 Homophily and heterophily assumptions

In this section we investigate adversarial attacks on node classification under the assumptions of label homophily and label heterophily. For this experiment we make use of the IIP and StudentDB datasets. The former is an example of a homophilic network where 70.9% of all edges connect nodes of the same label. On the other hand, the StudentDB is an example of a strongly heterophilic network where no edges connect nodes sharing the same label. We restrict our evaluation to the best performing node embedding approach i.e., node2vec, and discuss results for the strongest adversarial attack from Section 4.3.2 i.e., *add_ce* (results for other attacks are presented in Appendix A.4).

We start our evaluation by attacking both networks with budgets $b \in [0.0, 0.2, 0.6]$. We then learn node embeddings and perform downstream node classification for each network and attack budget. We proceed to record the correctly and incorrectly classified nodes at validation time in each case. Figure 4 presents three identical spring-layout representations of the IIP network, one for each attack budget. The nodes in each subplot are colored according to their status as correctly or incorrectly classified for that budget. From the Figure we can visually confirm that, as the attack strength increases, the misclassification rate (mr in the figure) also increases. This is also confirmed numerically by the mr value presented above each plot.

In Figure 5 we present the same information for the StudentDB network. In this case, as the attack strength increases the misclassification rate decreases (as can be seen both visually and through the mr values). This seemingly counter intuitive behavior can be explained by the fact that our attack introduces new edges in the graph. This results in a more compact representation and, in turn, lower margins for the node label classifier. It is important to note that classification performance increases despite the fact that the edges added are uninformative (they connect nodes with different labels further reinforcing heterophily). When other attacks are used such as random or degree-based edge additions, which can potentially be more informative as they might increase homophily, the misclassification margin decreases even further (see results in Appendix A.4). In our experiments we have also seen that edge deletion attacks do not have the same effect of lowering the misclassification rate as they further reduce the network density.

5 Conclusions

In this paper we have demonstrated that node embedding approaches, regardless of their underlying representation mechanisms, are sensitive to random and adversarial poison attacks. We have shown that results on downstream node classification are significantly less robust compared to those on network reconstruction. Our experiments also revealed that for low attacks budgets (below 20% of edges in the graph) edge addition attacks are generally stronger than edge deletions. Outside of this range, the opposite is true. Surprisingly, our empirical evaluation showed no significant differences between different heuristic-based adversarial attacks. Even leveraging full knowledge of the node labels when attacking node classification does not yield significantly stronger attacks. Finally, we have also shown that the number of labeled nodes plays a fundamental role in node classification robustness, that rewiring attacks are generally stronger than addition or deletion independently, and that attacks under heterophily assumption can unexpectedly result in better model performance. With this work and our the extension to robustness evaluation for the EvalNE software we hope to lay the foundations for further research in this area.

References

[1] Kush Bhatia, Prateek Jain, Parameswaran Kamalaruban, and Purushottam Kar. Consistent robust regression. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. 1

[2] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4334–4343. PMLR, 10–15 Jul 2018.

[3] Dimitris Bertsimas, Jack Dunn, Colin Pawlowski, and Ying Daisy Zhuo. Robust classification. *INFORMS Journal on Optimization*, 1(1):2–34, 2019. doi: 10.1287/ijoo.2018.0001.

[4] Wei Jin, Yao Ma, Xiaorui Liu, Xianfeng Tang, Suhang Wang, and Jiliang Tang. Graph structure learning for robust graph neural networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, KDD '20, page 66–74, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379984. doi: 10.1145/3394486.3403049. 1

[5] Stephan Günnemann. *Graph Neural Networks: Adversarial Robustness*, pages 149–176. Springer Nature Singapore, Singapore, 2022. ISBN 978-981-16-6054-2. doi: 10.1007/978-981-16-6054-2_8. 1, 2

[6] Cho-Jui Hsieh Pin-Yu Chen. *Adversarial Robustness for Machine Learning*. Elsevier, 2022. ISBN 9780128242575. 1

[7] Aleksandar Bojchevski and Stephan Günnemann. Adversarial attacks on node embeddings via graph poisoning. In *Proc. of ICML*, pages 695–704, 2019. 1, 4

[8] Bo Kang, Jefrey Lijffijt, and Tijl De Bie. Conditional network embeddings. In *Proc. of ICLR*, 2019. 1

[9] J. Qiu, Yuxiao Dong, Hao Ma, Jun Yu Li, Chi Wang, Kuansan Wang, and Jie Tang. Netsmf: Large-scale network embedding as sparse matrix factorization. *Proc. of WWW*, 2019.

[10] Alexandru Mara, Yoosof Mashayekhi, Jefrey Lijffijt, and Tijl de Bie. Csne: Conditional signed network embedding. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, CIKM '20, page 1105–1114. Association for Computing Machinery, 2020. ISBN 9781450368599. doi: 10.1145/3340531.3411959. 1

[11] Anton Tsitsulin, Davide Mottin, Panagiotis Karras, and Emmanuel Muller. Verse: Versatile graph embeddings from similarity measures. In *Proc. of WWW*, page 539–548, 2018. ISBN 9781450356398. doi: 10.1145/3178876.3186120. 1, 3

[12] Yi-An Lai, Chin-Chi Hsu, Wen Hao Chen, Mi-Yen Yeh, and Shou-De Lin. Prune: Preserving proximity and global ranking for network embedding. In *Proc. of NIPS*, pages 5257–5266, 2017. 1, 3

[13] Laila Rasmy, Yang Xiang, Ziqian Xie, Cui Tao, and Degui Zhi. Med-BERT: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. In *npj Digital Medicine*, 2021. doi: 10.1038/s41746-021-00455-y. 1

[14] Lei Guo, Yufei Wen, and Xinhua Wang. Exploiting pre-trained network embeddings for recommendations in social networks. *Journal of Computer Science and Technology*, 33:682–696, 2018. 1

[15] Federico Monti, Fabrizio Frasca, Davide Eynard, Damon Mannion, and Michael M. Bronstein. Fake news detection on social media using geometric deep learning. *ArXiv*, abs/1902.06673, 2019. 1

[16] Daniel Zügner, Oliver Borchert, Amir Akbarnejad, and Stephan Günnemann. Adversarial attacks on graph neural networks: Perturbations and their patterns. *ACM Trans. Knowl. Discov. Data*, 14(5), jun 2020. ISSN 1556-4681. doi: 10.1145/3394520. 1, 2

[17] Xi Chen, Bo Kang, Jefrey Lijffijt, and Tijl De Bie. Adversarial robustness of probabilistic network embedding for link prediction. In *PKDD/ECML Workshops*, 2021. 1, 2

[18] Aditya Grover and Jure Leskovec. node2vec: Scalable feature learning for networks. In *Proc. of KDD*, pages 855–864, 2016. 2, 3

[19] Shaosheng Cao, Wei Lu, and Qiongkai Xu. GraRep: Learning graph representations with global structural information. In *Proc. of CIKM*, pages 891–900, 2015. 2, 3

[20] Daixin Wang, Peng Cui, and Wenwu Zhu. Structural deep network embedding. In *Proc. of KDD*, pages 1225–1234, 2016. 2, 3

[21] Yiyue Qian, Yiming Zhang, Qianlong Wen, Yanfang Ye, and Chuxu Zhang. Rep2vec: Repository embedding via heterogeneous graph adversarial contrastive learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '22, page 1390–1400, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393850. doi: 10.1145/3534678.3539324. 2

[22] Jinxin Cao, Weizhong Xu, Di Jin, Xiaofeng Zhang, Anthony Miller, Lu Liu, and Weiping Ding. A network embedding-enhanced nmf method for finding communities in attributed networks. *IEEE Access*, pages 1–1, 2022. doi: 10.1109/ACCESS.2022.3198979.

[23] Asan Agibetov. Neural graph embeddings as explicit low-rank matrix factorization for link prediction. *Pattern Recognition*, 133:108977, 2023. ISSN 0031-3203. doi: https://doi.org/10.1016/j.patcog.2022.108977. 2

[24] Scott Freitas, Diyi Yang, Srijan Kumar, Hanghang Tong, and Duen Horng Chau. Evaluating graph vulnerability and robustness using tiger. *ACM International Conference on Information and Knowledge Management*, 2021. 2

[25] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian J. Goodfellow, and Rob Fergus. Intriguing properties of neural networks. In Yoshua Bengio and Yann LeCun, editors, *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. 2

[26] Ian J. Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *CoRR*, abs/1412.6572, 2015. 2

[27] Ivan Evtimov, Kevin Eykholt, Earlence Fernandes, Tadayoshi Kohno, Bo Li, Atul Prakash, Amir Rahmati, and Dawn Song. Robust physical-world attacks on machine learning models. *CoRR*, abs/1707.08945, 2017. 2

[28] Yevgeniy Vorobeychik and Murat Kantarcioglu. Adversarial machine learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 12(3):1–169, 2018. 2

[29] Hanjun Dai, Hui Li, Tian Tian, Xin Huang, Lin Wang, Jun Zhu, and Le Song. Adversarial attack on graph structured data. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1115–1124. PMLR, 10–15 Jul 2018. 2

[30] Alexandru Mara, Jefrey Lijffijt, and Tijl De Bie. An empirical evaluation of network representation learning methods. *Big Data*, 00, 2022. doi: 10.1089/big.2021.0107. 2, 4

[31] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In Yoshua Bengio and Yann LeCun, editors, *Proc. of ICLR*, 2013. 3

[32] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. Deepwalk: Online learning of social representations. In *Proc. of KDD*, pages 701–710, 2014. 3

[33] Jian Tang, Meng Qu, Mingzhe Wang, Ming Zhang, Jun Yan, and Qiaozhu Mei. LINE: Large-scale information network embedding. In *Proc. of WWW*, pages 1067–1077, 2015. 3

[34] Mingdong Ou, Peng Cui, Jian Pei, Ziwei Zhang, and Wenwu Zhu. Asymmetric transitivity preserving graph embedding. In *Proc. of KDD*, pages 1105–1114, 2016. 3

[35] Jiezhong Qiu, Yuxiao Dong, Hao Ma, Jian Li, Kuansan Wang, and Jie Tang. Network embedding as matrix factorization: Unifying deepwalk, line, pte, and node2vec. In *Proc. of WSDM*, page 459–467, 2018. ISBN 9781450355810. doi: 10.1145/3159652.3159706. 3

[36] Xiao Wang, Peng Cui, Jing Wang, Jian Pei, Wenwu Zhu, and Shiqiang Yang. Community preserving network embedding. In *Proc. of AAAI*, pages 203–209, 2017. 3

[37] M. E. J. Newman. Mixing patterns in networks. *Physical Review E*, 67 026126:1024–1034, 2003. 4

[38] Alexandru Mara, Jefrey Lijffijt, and Tijl De Bie. Evalne: A framework for network embedding evaluation. *SoftwareX*, 17, 2022. ISSN 100997. doi: 10.1016/j.softx.2022.100997. 4

[39] C. Lee Giles, Kurt D. Bollacker, and Steve Lawrence. Citeseer: an automatic citation indexing system. In *INTERNATIONAL CONFERENCE ON DIGITAL LIBRARIES*, pages 89–98. ACM Press, 1998. 5

[40] Andrew Kachites Mccallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. Automating the construction of internet portals with machine learning. *Information Retrieval*, 3:127–163, 2000. 5

[41] Lada A. Adamic and Natalie Glance. The political blogosphere and the 2004 u.s. election: Divided they blog. In *Proceedings of the 3rd International Workshop on Link Discovery*, LinkKDD '05, page 36–43, New York, NY, USA, 2005. Association for Computing Machinery. ISBN 1595932151. doi: 10.1145/1134271.1134277. 5

[42] Jure Leskovec and Andrej Krevl. SNAP Datasets: Stanford large network dataset collection, 2015. 5

[43] Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *AAAI*, 2015. URL https://networkrepository.com. 6

[44] Bart Goethals, Wim Le Page, and Michael Mampaey. Mining interesting sets and rules in relational databases. In *Proc. of SAC*, pages 997–1001, 2010. ISBN 978-1-60558-639-7. doi: 10.1145/1774088.1774299. 6

A Appendix

A.1 Further dataset details

The IIP network represents a set of companies competing in the internet industry between 1998 and 2001. Nodes in the graph denote companies and edges represent business relations such as joint venture, strategic alliance or other type of partnership. The associated node labels denote the company's main business area i.e., content, infrastructure of commerce.

The StudentDB network represents a snapshot of Antwerp University's relational student database. Nodes in the network represent entities, more specifically: students, professors, tracks, programs, courses and rooms. Edges constitute binary relations between them, that is, student-in-track, student-in-program, student-takes-course, professor-teaches-course, and course-in-room. Numerical node labels are assigned according to each node's type.

A.2 Randomized attacks: additional results

In this section we present our additional experiments regarding randomized attacks on node embeddings. We start in Figures 6 and 7 by presenting the node classification f1_micro results for the Cora dataset and the network reconstruction AUC scores for PolBlogs.

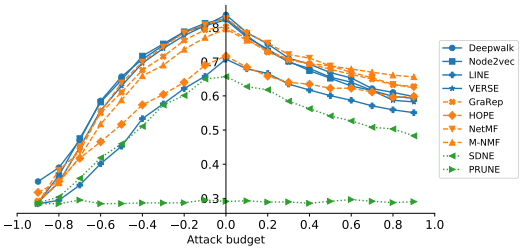

Figure 6: Node classification performance for the Cora network. Y axis indicates f1_micro scores. Negative attack budgets indicate edge deletion.

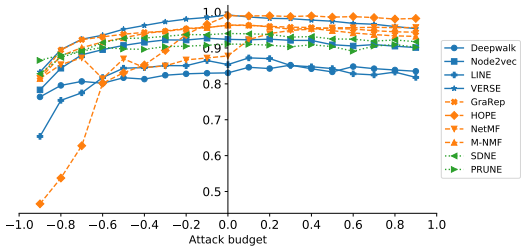

Figure 7: Network reconstruction performance for the PolBlogs network. Y axis indicates AUC scores. Negative attack budgets indicate edge deletion.

In Figures 8 and 9 we summarize the f1_macro scores for both Citeseer and Cora and Figures 8 and 9 present the average precision on Facebook and PolBlogs.

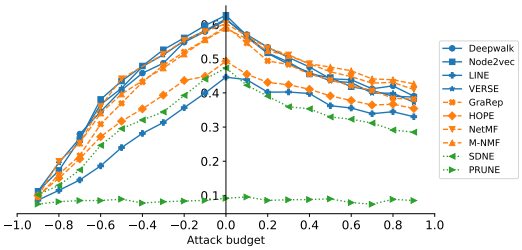

Figure 8: Node classification performance for the Citeseer network. Y axis indicates f1_macro scores. Negative attack budgets indicate edge deletion.

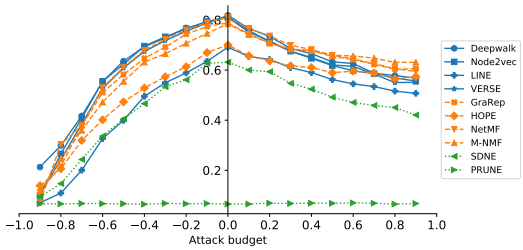

Figure 9: Node classification performance for the Cora network. Y axis indicates f1_macro scores. Negative attack budgets indicate edge deletion.

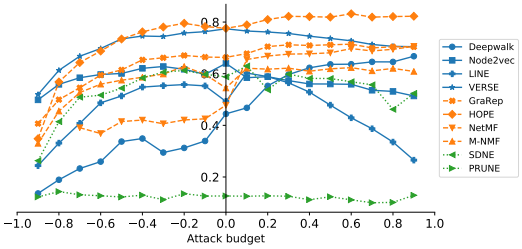

Figure 10: Network reconstruction performance for the Facebook network. Y axis indicates average precision scores. Negative attack budgets indicate edge deletion.

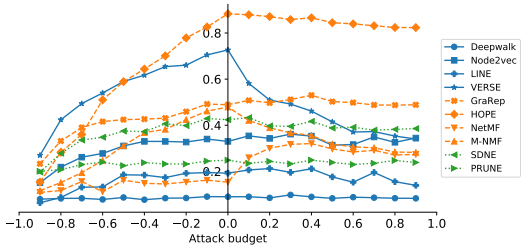

Figure 11: Network reconstruction performance for the PolBlogs network. Y axis indicates average precision scores. Negative attack budgets indicate edge deletion.

A.3 Other attacks: additional results

We also compare the performance of edge addition, rewiring and deletion on both downstream tasks in terms of f1_micro and average precision. These results support our conclusions in Section 4.3.3 (see Figure 12).

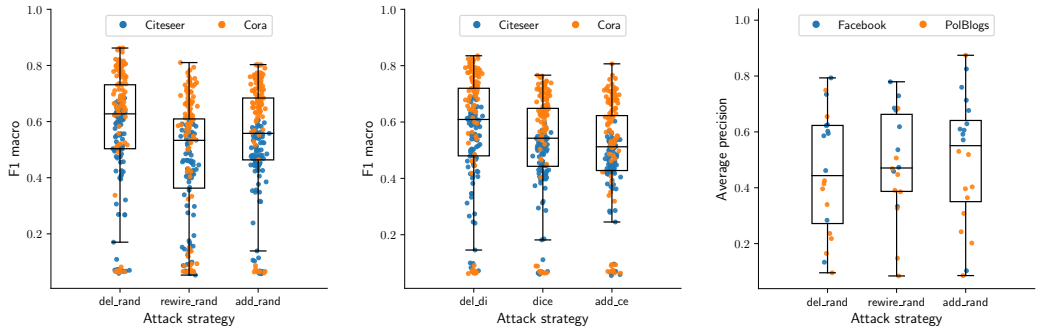

Figure 12: Comparison of edge addition, rewiring and deletion attacks for both downstream tasks. The leftmost and center figures present f1_macro scores for random and node label based attacks on node classification. The rightmost figure shows average precision results for random attacks on network reconstruction.

A.4 Homophily and Heterophily assumptions: additional results

In this section we present two additional plots showing the decrease in misclassification rate after an adversarial attack on the StudentDB network. In Figure 13 we present the correctly and incorrectly classified nodes under random edge deletion (*del_rand*) for different budget values and in Figure 14 we show the results for random edge addition (*add_rand*).

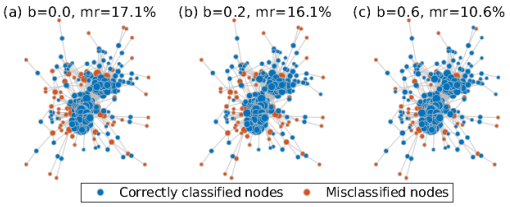

Figure 13: Correctly and incorrectly classified nodes for the heterophilic StudentDB network with attack strategy *del_rand* and varying attack budgets.

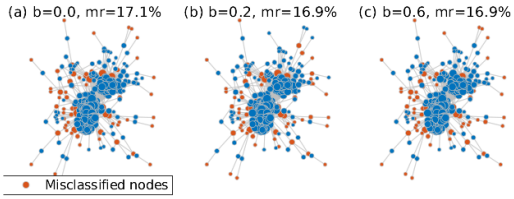

Figure 14: Correctly and incorrectly classified nodes for the heterophilic StudentDB network with attack strategy *add_rand* and varying attack budgets.