

Performance Evaluation of Knowledge Graph Embedding Approaches under Non-adversarial Attacks

Sourabh Kapoor

Arnab Sharma

Michael Röder

Caglar Demir

Axel-Cyrille Ngonga Ngomo

Data Science Research Group, Paderborn University

SOURABHK@MAIL.UNI-PADERBORN.DE

ARNAB.SHARMA@UPB.DE

MICHAEL.ROEDER@UPB.DE

CAGLAR.DEMIR@UPB.DE

AXEL.NGONGA@UPB.DE

Abstract

Knowledge Graph Embeddings (KGEs) transform a discrete Knowledge Graph (KG) into a continuous vector space facilitating its use in various AI-driven applications like Semantic Search, Question Answering, or Recommenders. While KGE approaches are effective in these applications, most existing approaches assume that all information in the given KG is correct. This enables attackers to influence the output of these approaches, e.g., by perturbing the input. Consequently, the robustness of such KGE approaches has to be addressed. Recent work focused on adversarial attacks. However, non-adversarial attacks on all attack surfaces of these approaches have not been thoroughly examined. We close this gap by evaluating the impact of non-adversarial attacks on the performance of 5 state-of-the-art KGE algorithms on 5 datasets with respect to attacks on 3 attack surfaces—graph, parameter, and label perturbation. Our evaluation results suggest that label perturbation has a strong effect on the KGE performance, followed by parameter perturbation with a moderate and graph with a low effect.

Keywords: Knowledge graph embedding, Non-adversarial attack, Robustness

1. Introduction

A KG is a structured representation of knowledge, typically organized as a multi-relational directed graph where nodes represent entities or concepts, and edges represent relationships between them. The knowledge of real-world facts is represented therein in the form of triples denoted as (h, r, t) where h and t correspond to the head and tail entities and r is the relationship between them. Due to their effectiveness in representing knowledge, KGs have been used in various areas such as in information retrieval (Dalton et al., 2014), question answering (Ferrucci et al., 2010), and others. To make efficient use of the knowledge representation in KGs, knowledge graph embedding (KGE) models (Bordes et al., 2013b; Dettmers et al., 2018) are introduced which aim to capture the complex relationships between entities and relations in KGs. This is done by embedding symbolic representations of KGs into continuous vector spaces by preserving their inherent structure.

The demand to develop effective KGE models to be applied in various downstream application tasks is ever increasing and that has led to building KGs, harnessing data from public sources, e.g., DBpedia (Auer et al., 2007). Although this has led to the benefit of high-quality KGE models to be used for downstream tasks, this has also opened a new *attack*

window for malicious users. More specifically, the usage of KGE models by utilizing open source KGs as the basis introduces malicious attempts to poison the KGs and thereby the KGE model as well. In recent years, several researchers, studied different *adversarial* attack strategies on KGE models by poisoning the KGs or by performing adversarial manipulations of the embedding model (Zhang et al., 2019a; Pezeshkpour et al., 2019; Bhardwaj et al., 2021a,b; You et al., 2023). The fundamental concept behind these attacks is to focus on a particular fact and manipulate the KGE model to either increase or decrease its *plausibility* score. This score represents the likelihood of the fact being true: a higher score indicates a higher probability, while a lower score indicates a lower probability. Apart from these works that perform targeted adversarial attacks, an attacker might simply perform non-targeted adversarial (or non-adversarial) attacks on the KGE models. Note that such studies have been carried out for machine learning (ML) models by Hendrycks and Dietterich (2018), but not for the KGE models. These models are frequently used in many critical areas in the web domain. Since the web is a critical point of any country’s information sources, an attacker might attempt to disrupt the performance of some critical services (e.g., knowledge-graph-based chatbots on government webpages), thereby destabilizing the country. This kind of attack does not need to have a concrete target and can simply be an attack to degrade the performance of the critical information sources. We can think of such attacks as being similar to denial of service (DoS) attacks. Studying the security of KGE models is only infancy, and some limited works so far have focused on adversarial attacks. However, we do strongly believe that non-adversarial attacks need to be studied to make KGE models that are robust and trustworthy. Therefore, in this work, we study non-adversarial attacks on the KGE approaches considering different attack surfaces.

In this work, we perform such non-adversarial attacks considering the entire learning framework of the KGE approaches, i.e., performing attacks on the 1. knowledge graph, 2. parameters, and 3. output labels. In case of (1), we attack by perturbing the existing triples selected randomly from the KG. More specifically, k percentage of the triples are chosen randomly, and then for each of the selected triples, based on the random value from a Bernoulli distribution, either the head or the relation of the triple is changed (i.e., replaced by some other entity or relation in the same KG). In (2), the embedding space of the underlying KGE model is targeted where the embedding vectors are perturbed. Herein again a k percentage of the embeddings is selected and then for each of the selected embeddings, either the head or the relation is chosen (based on the Bernoulli distribution). Using a probability distribution, continuous noise is then added to either the head or the relation. Finally, in (3), the labels of randomly selected triples, which in the case of the KGE models are typically the tail entity, are simply flipped, the 0s to 1s and vice-versa. We aim to find out the robustness of the existing state-of-the-art KGE models when these non-adversarial attacks in these three levels are done. Precisely, we want to investigate if some KGE models can perform better than others and if so, in which cases and how much it might depend on the underlying KGs. To this end, we have considered 5 datasets and 5 state-of-the-art KGE algorithms to perform non-adversarial attacks. Our results suggest that the attack by performing the label perturbation causes the worst degradation of the performance of the KGE models, followed by parameter and graph perturbations. Moreover, in graph perturbation, for some models, which do not perform well, initially, perturbations can act as a regularizer, thereby improving their overall performance.

Preliminaries and some formalizations that are used throughout the paper are given in Section 3. Section 4 describes the three different attack approaches that are considered in this work. Section 5 shows details about the experiments and the computational results. We discuss related studies in Section 2.

2. Related Work

In the context of performing malicious attacks on KGE approaches, not many works can be found in the literature, and most importantly, most of them focused on performing adversarial attacks. For instance, Zhang et al. (2019a) first introduced *data poisoning* attack strategies to perform adversarial attacks on the KGE models by adding or deleting specific triples. To this end, their strategies follow a two-step process, (a) shifting the embedding of either of the head or tail entities of a target triple to maximize the attack goal, and then (b) adding and/or removing triples from the KG which would facilitate in achieving the goal in (a). The aim in this setting is to degrade or promote the plausibility of a specific fact (i.e., the target triple). A later work by Pezeshkpour et al. (2019) followed a similar sort of setting where they used a gradient-based approach to find out the most influential neighboring triples of the target fact, the removal of which would maximize the attack objective. Searching is performed in the embedding space and then an auto-encoder is used to generate the triples of KG. Bhardwaj et al. (2021a) attempted to leverage the inductive capabilities of the KGE models, which are encapsulated by relationship patterns such as symmetry, inversion, and composition within the knowledge graph to perform adversarial attacks. Their approach aims to decrease or increase the model’s confidence in predicting target facts by enhancing its confidence in predicting a set of *decoy* triples. A further work by them (Bhardwaj et al., 2021b) used instance attribution methods from the domain of interpretable ML to perform data poisoning attacks on KGE models. Such attribution methods are first used to identify a (set of) triple(s) in the training set, which contributes most to the prediction of a specific target triple. Then the triple from the training set is either removed or added by replacing one of the two entities of the influential triple. You et al. (2023) recently proposed approaches for data poisoning attacks by considering several aspects: black-box attack, poisoning by adding semantically preserving triples, and stealthiness by showing good performance on the *cleaned* triples.

Finally, apart from these works which focused on adversarial robustness, there exists a line of works focusing on building KGE models that are *robust* to noise in KGs, by Xie et al. (2018); Shan et al. (2018); Nayyeri et al. (2021); Zhang et al. (2023), amongst others. To this end, Xie et al. (2018) first proposed the idea of global and local *confidence* scores to identify a tripe as a correct (positive) or a noisy (negative) triple. Assigning scores to triples would help the KGE model to distinguish the correct triple from the noisy ones, thereby dictating the model to learn correctly with the help of the adjusted loss function. Shan et al. (2018) proposed *dissimilarity* measure and *support* score alongside confidence score to categorize noisy triples. Cheng et al. (2020) proposed to use an adversarial training setup, extending the previous works to improve over the works by Xie et al. (2018). Precisely, they came up with a loss function that makes the KGE models aware of noisy triples. In a recent work, Zhang et al. (2023) proposed a reinforcement learning framework to identify the noisy triples before the training and then remove them. Thus, the KGE model generated

in this way would be robust to noise in KG. Note that, all these works consider noise as it is inherently present in KG. Therefore, they proposed approaches to make the KGE models robust against such noise. However, none of them evaluated the performance changes of the KGE models when such noise is added as a form of non-adversarial attacks.

3. Preliminaries and Notation

Let \mathcal{E} be the set of entities that are of interest and \mathcal{R} the set of relations that exist between these entities. We express assertions about the entities using triples. A triple (h, r, t) comprises a head and a tail entity ($h, t \in \mathcal{E}$) and a relation $r \in \mathcal{R}$ that holds between them. We define a knowledge graph \mathcal{G} as a collection of triples:

$$\mathcal{G} := \{(h, r, t) \in \mathcal{E} \times \mathcal{R} \times \mathcal{E}\}. \quad (1)$$

KGs are representations of information in a discrete space. However, many modern algorithms cannot process such a graph. Hence, KGE algorithms have been suggested to represent the knowledge of a KG in a continuous, low-dimensional embedding space.

Let \mathbb{V} denote a normed-division algebra, e.g. $\mathbb{R}, \mathbb{C}, \mathbb{H}$, or \mathbb{O} (Balažević et al., 2019a; Demir et al., 2021; Yang et al., 2014; Trouillon et al., 2016; Zhang et al., 2019b). A KGE model of a KG comprises entity embeddings $\mathbf{E} \in \mathbb{V}^{|\mathcal{E}| \times d_e}$ and relation embeddings $\mathbf{R} \in \mathbb{V}^{|\mathcal{R}| \times d_r}$, where d_e and d_r are the size of the embedding vectors. In the following, we use d as size for all embedding vectors, as it has been shown that $d_e = d_r$ holds for many types of models (Nickel et al., 2015). Throughout this paper, we will denote embedding vectors with bold fonts, for instance, the embedding of h , r , and t will be denoted as \mathbf{h} , \mathbf{r} , and \mathbf{t} , respectively.

Given a KG, a KGE algorithm has the goal to find a KGE model that optimizes its scoring function. Most of these algorithms are tailored towards link prediction (Chami et al., 2020; Hogan et al., 2021), i.e., their scoring function is $\phi_{\Theta} : \mathcal{E} \times \mathcal{R} \times \mathcal{E} \mapsto \mathbb{R}$, where Θ denotes parameters and often comprise \mathbf{E} , \mathbf{R} , and additional parameters (e.g., affine transformations, batch normalizations, convolutions). Given an assertion in the form of a triple $(h, r, t) \in \mathcal{E} \times \mathcal{R} \times \mathcal{E}$, a prediction $\hat{y} := \phi_{\Theta}(h, r, t)$ signals the likelihood of (h, r, t) being true (Dettmers et al., 2018). Since \mathcal{G} contains only assertions that are assumed to be true, assertions assumed to be false have to be generated. While different generation methods exist, we will focus on KvsAll (Dettmers et al., 2018), since recent KGE approaches are commonly trained with this strategy (Balažević et al., 2019a,b; Nguyen et al., 2018; Demir and Ngomo, 2021; Ruffinelli et al., 2020).

Let \mathcal{D} = denote the training dataset for the KvsAll training strategy. It comprises training data points $(\mathbf{x}, \mathbf{y}) \in \mathcal{D}$ that correspond to unique head entity and relation pairs ($\mathbf{x}=(\mathbf{h}, \mathbf{r})$) that occur in \mathcal{G} with a binary label vector $\mathbf{y} \in \{0, 1\}^{|\mathcal{E}|}$, where $\mathbf{y}_i = 1$ for the i -th entity $e \in \{e | (h, r, e) \in \mathcal{G}\}$, otherwise 0. Consequently, $|\mathcal{D}|$ equals to the number of unique head entity relation pairs in the graph $|\{(h, r) \in \mathcal{E} \times \mathcal{R} \mid x \in \mathcal{E} \wedge (h, r, x) \in \mathcal{G}\}|$. During the training process, most KGE algorithms divide the training data into mini-batches. A mini-batch \mathcal{B} consists of m data points with $m \times |\mathcal{E}|$ binary labels. The data points are used to update the entity and relation embeddings \mathbf{E} and \mathbf{R} . Hence, during training, a mini-batch is typically represented using the embedding vectors, i.e., \mathcal{B} is expressed as $\mathcal{B} = \{(\mathbf{x}, \mathbf{y})\}$, where $\mathbf{x} = (\mathbf{h}, \mathbf{r})$ comprises the embedding vectors of h and r . The training is typically

performed in several epochs. Within each epoch, all mini-batches are used to update the model’s parameters.

4. Methodology

In this work, we perform non-adversarial attacks on KGE algorithms, i.e., our goal is to reduce the performance of KGE models in a downstream task. To this end, we use three different attack surfaces: 1. the input knowledge graph, 2. the target labels, and 3. the model parameters. Every attack incorporates a parameter k which allows us to regulate the extent to which data is perturbed during the attack.

4.1. Graph Perturbation

The first attack surface that we look at is the training data that is gathered from the KG. During this attack, the attacker perturbs $k\%$ of the input data within each mini-batch by changing the head or relation information of the data points. Let \mathcal{B} be a mini-batch and let $\mathcal{B}^* \subset \mathcal{B}$ be a randomly sampled subset comprising $k\%$ of the training examples of \mathcal{B} . The attacker replaces the original mini-batch with a perturbed version of the batch by replacing \mathcal{B}^* with $\mathcal{B}^{*'}$:

$$\mathcal{B}' = \mathcal{B}^{*'} \cup (\mathcal{B} \setminus \mathcal{B}^*), \quad (2)$$

with $|\mathcal{B}^{*'}| = |\mathcal{B}^*|$. Hence, the graph Perturbation (GP) attack is defined as generating $\mathcal{B}^{*'}$ based on \mathcal{B}^* by perturbing the head or relation information that has been gathered from the KG. Let $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{B}^*$ be the i -th data point in \mathcal{B}^* . Let ξ_i be the i -th random value sampled from a Bernoulli distribution with a probability of 0.5 being either 0 or 1. Let h'_i and r'_i be randomly sampled elements from \mathcal{E} and \mathcal{R} , respectively. Within this attack, an attacker perturbs the data point $(\mathbf{x}_i, \mathbf{y}_i)$ by creating \mathbf{x}'_i as follows:

$$\mathbf{x}'_i = \begin{cases} (h'_i, r_i) & \text{if } \xi_i = 0, \\ (h_i, r'_i) & \text{else.} \end{cases} \quad (3)$$

This perturbation is applied to all data points in \mathcal{B}^* to form $\mathcal{B}^{*'}$:

$$\mathcal{B}^{*' } = \{(\mathbf{x}'_i, \mathbf{y}_i) | (\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{B}^*\}. \quad (4)$$

To give an example of such a perturbation let us assume one of the data points from the set \mathcal{B}^* as $\mathbf{x} = (:Einstein, :bornIn)$. If $\xi_i = 0$ then the perturbed point could be $\mathbf{x}' = (:Laplace, :bornIn)$, whereas if $\xi_i \neq 0$ then $\mathbf{x}' = (:Einstein, :capitalOf)$.

Note that, the idea of manipulating the training data by adding some specific triples in order to make the trained model giving specific predictions is termed as *data poisoning* or adversarial attacks (Bhardwaj et al., 2021a; Zhang et al., 2019a; Pezeshkpour et al., 2019). In the context of knowledge graph embeddings, different attack strategies are gaining popularity due to the critical downstream applications of KGE models. The goal of an attacker is to introduce *malicious* facts in terms of adding triples in the training data, leading to the generation of *poisoned* KGE models. In this work, however, we do not aim to perform such adversarial attacks on the KGE models. More specifically, we do not have a specific target of introducing a fact, rather, we perform non-adversarial attacks by

perturbing either the head or the relations of specific triples in the graph and then finally measure how the perturbation would affect the robustness of the KGE models.

4.2. Label Perturbation

The Label Perturbation (LP) is a similar attack as GP. However, within this attack, the attacker perturbs a data point $(\mathbf{x}_i, \mathbf{y}_i)$ by inverting the label vector as follows:

$$\mathbf{y}'_i = \{-y_{i,j} | y_{i,j} \in \mathbf{y}_i\}. \quad (5)$$

This perturbation is applied to all data points in \mathcal{B}^* to form $\mathcal{B}^{*'}:$

$$\mathcal{B}^{*' } = \{(\mathbf{x}_i, \mathbf{y}'_i) | (\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{B}^*\}. \quad (6)$$

For a data point with $\mathbf{x}_i = (:Einstein, :bornIn)$ and a vector \mathbf{y}_i filled with zeros except for a single 1 at the id of the entity `:Ulm`, the attack would perturb the label vector by inverting all its values. Hence, the new label vector \mathbf{y}'_i would express that the embedding model is expected to predict that the entity `:Einstein` has the relation `:bornIn` to all entities, except `:Ulm`.

In ML, label perturbation of training data is commonly employed to mitigate overfitting and noise. For instance, Szegedy et al. (2016) used a variant of label perturbation called label smoothing to improve the generalization performance. There exist other such methods such as bootstrapping loss by Reed et al. (2015) and label correction by Patrini et al. (2017), which are different types of label perturbation techniques, introduced to generate robust models. Some studies have further explored strategies for adversarial attacks on deep learning models via label perturbations (Song et al., 2018; Zhang et al., 2022). However, none of these prior works considered knowledge graphs in this context, which we look into. The attack defined above is a direct application of the method proposed by Song et al. (2018); Zhang et al. (2022).

4.3. Parameter Perturbation

The third attack surface does not focus on the training data but on the learned parameters. This Parameter Perturbation (PP) changes $k\%$ of the learned vectors, before each of the training epochs. More formally, let \mathcal{B} be the vector-based representation of a mini-batch \mathcal{B} between two epochs and let $\mathcal{B}^* \subset \mathcal{B}$ be a randomly sampled subset comprising $k\%$ of the training examples of \mathcal{B} . The attacker replaces the original mini-batch with a perturbed version of the batch by replacing \mathcal{B}^* with $\mathcal{B}^{*'}$:

$$\mathcal{B}' = \mathcal{B}^{*' } \cup (\mathcal{B} \setminus \mathcal{B}^*). \quad (7)$$

The attack is defined as generating $\mathcal{B}^{*'}$ by perturbing the head or relation vectors in \mathcal{B}^* . Let $(\mathbf{x}_i, \mathbf{y}_i) \in \mathcal{B}^*$ be the i -th data point in \mathcal{B}^* . Let ξ_i be the i -th random value sampled from a Bernoulli distribution with a probability of 0.5 being either 0 or 1. Let \mathbf{q} be a d -dimensional vector with randomly sampled values. Within this attack, an attacker perturbs the data point $(\mathbf{x}_i, \mathbf{y}_i)$ by creating \mathbf{x}'_i as follows:

$$\mathbf{x}'_i = \begin{cases} (\mathbf{h} + \mathbf{q}, \mathbf{r}) & \text{if } \xi_i = 0, \\ (\mathbf{h}, \mathbf{r} + \mathbf{q}) & \text{else.} \end{cases} \quad (8)$$

Some existing works showed that such perturbations could be used by attackers to attack the learned models. For instance, Kurita et al. (2020) proposed an optimization algorithm to perturb the weights of a DNN model in such a way so that whenever specific feature values are present in the input, the output will be predicted to a specific class.¹ There is a different line of work by Bai et al. (2021) that performs such kind of perturbation on the models’ parameters, however on the memory level by flipping the bits of the parameters. However, similar to the perturbation performed on the graph and labels, the works mentioned above belong to a line of works that perform adversarial attacks with a specific goal in mind. For KGE models no such works have considered evaluating the models’ performance when non-adversarial attacks are performed on the learned parameters.

5. Evaluation

In this section, we first describe the experimental setup of our evaluation describing the datasets and the models we consider in this work. Next, we report the results of our evaluation based on the three attack surfaces we look into.

5.1. Datasets

Table 1 lists the datasets and their features that we use for the evaluation of the impact of the non-adversarial attacks on KGE algorithms. The UMLS dataset by McCray (2003) contains 135 medical entities and their connections using 46 distinct relations. The KINSHIP dataset by Denham (2014) describes the Alyawarra tribe’s kinship dynamics with 25 unique relationship types. Apart from these two smaller datasets, we also use three larger datasets. WN18RR by Dettmers et al. (2018) is a version of WordNet optimized for the link prediction task proposed by Bordes et al. (2013a). NELL-995-h100 is a subset of the Never-Ending Language Learning dataset by Xiong et al. (2017). FB15k-237 by Toutanova and Chen (2015) is a subset of the Freebase knowledge graph.

Table 1: Datasets used throughout the evaluation and their features (number of entities, relations, and triples in each split).

Dataset	$ \mathcal{E} $	$ \mathcal{R} $	$ \mathcal{G}^{\text{Train}} $	$ \mathcal{G}^{\text{Validation}} $	$ \mathcal{G}^{\text{Test}} $
UMLS (McCray, 2003)	135	46	5,216	652	661
KINSHIP (Denham, 2014)	104	25	8,544	1,068	1,074
WN18RR (Dettmers et al., 2018)	40,943	22	86,835	3,034	3,134
NELL-995-h100 (Xiong et al., 2017)	22,411	43	50,314	3,763	3,746
FB15K-237 (Toutanova and Chen, 2015)	14,541	237	272,115	17,535	20,466

5.2. Experimental Setup

Throughout our evaluation, we use 5 KGE algorithms with different embedding spaces: DistMult (\mathbb{R}) (Yang et al., 2014), ComplEx (\mathbb{C}) (Trouillon et al., 2016), QMult (\mathbb{H}) (Zhang

1. Note that, in the literature, such attacks are also called *trojan attacks* where the attacker aims that the model predicts a specific class when some specific feature values are present (Liu et al., 2018).

et al., 2019b), MuRE (\mathbb{R}) (Balažević et al., 2019), and Keci (Demir and Ngomo, 2023). With our experiments, we compare the performance of these KGE algorithms on the aforementioned datasets with and without non-adversarial attacks. For each attack described in Section 4, we evaluate the performance of the algorithms using an increasing perturbation ratio $k \in \{0, 0.01, 0.02, 0.04, 0.08, 0.16, 0.32, 0.64\}$. For attack surfaces, that rely on probability distributions, we use an even distribution. We repeat each experiment 5 times with different seed values for random number generators. We measure the KGE performance in terms of Hits@N and Mean Reciprocal Rank (MRR). However, we only report the MRR values on the test data within this paper due to the brevity of this work.² For each KGE approach, we choose the size of embedding vectors d so that all vectors can be represented as 32-dimensional real-valued vectors. Furthermore, we apply a consistent set of hyperparameters across all experiments. We use a learning rate of 0.1, a training duration of 100 epochs, a mini-batch size of 1024, and the KvsAll scoring technique. Additionally, for the Keci algorithm, we set its two additional parameters $p = 0$ and $q = 1$ to the default values suggested by Demir and Ngomo (2022).

5.3. Results

5.3.1. GRAPH PERTURBATION.

Figure 1 reports the average test MRR performances of the aforementioned KGE algorithms with different ratios of graph perturbation on the aforementioned datasets. The results show that on 4 out of 5 datasets, nearly all KGE algorithms show a clear decline in MRR at higher perturbation levels, specifically at 32% and 64% perturbation ratios. On the *FB15k-237* and *WN18RR* datasets, the decline already starts earlier, e.g., MuRE and Keci show significant performance reductions starting from 8% perturbation on *FB15k-237*. Only on the *NELL-995-h100* dataset, our results suggest that the perturbations have little to no effect on the test MRR. In addition, the results on *WN18RR* and *NELL-995-h100* indicate that the QMult KGE model is sensitive to the randomness induced which is demonstrated by the varied performance of QMult at same ratios but for different random seeds.

5.3.2. LABEL PERTURBATION.

The impact of the Label Perturbation on the KGE algorithm performance is larger compared to the Graph Perturbation. Figure 2 shows that the MRR on the test set drops dramatically with higher perturbation rates on the small *UMLS* and *KINSHIP* datasets. On the other three datasets, the effect is even more severe. Even a small perturbation rate of 0.1% already causes the MRR of all KGE approaches to drop close to 0.

5.3.3. PARAMETER PERTURBATION.

Figure 3 summarizes the results of the Parameter Perturbation experiments. On all datasets, the performance of the KGE decreases with perturbation rates of 16% or higher. On the small datasets, the effect already starts with a perturbation rate of 1 or 2%. In contrast, small perturbation ratios do not seem to have a big influence on the KGE algorithm performance when tested with larger datasets. As in the Graph Perturbation experiment, QMult

2. <https://figshare.com/s/4367528fa5c6af381a5a>

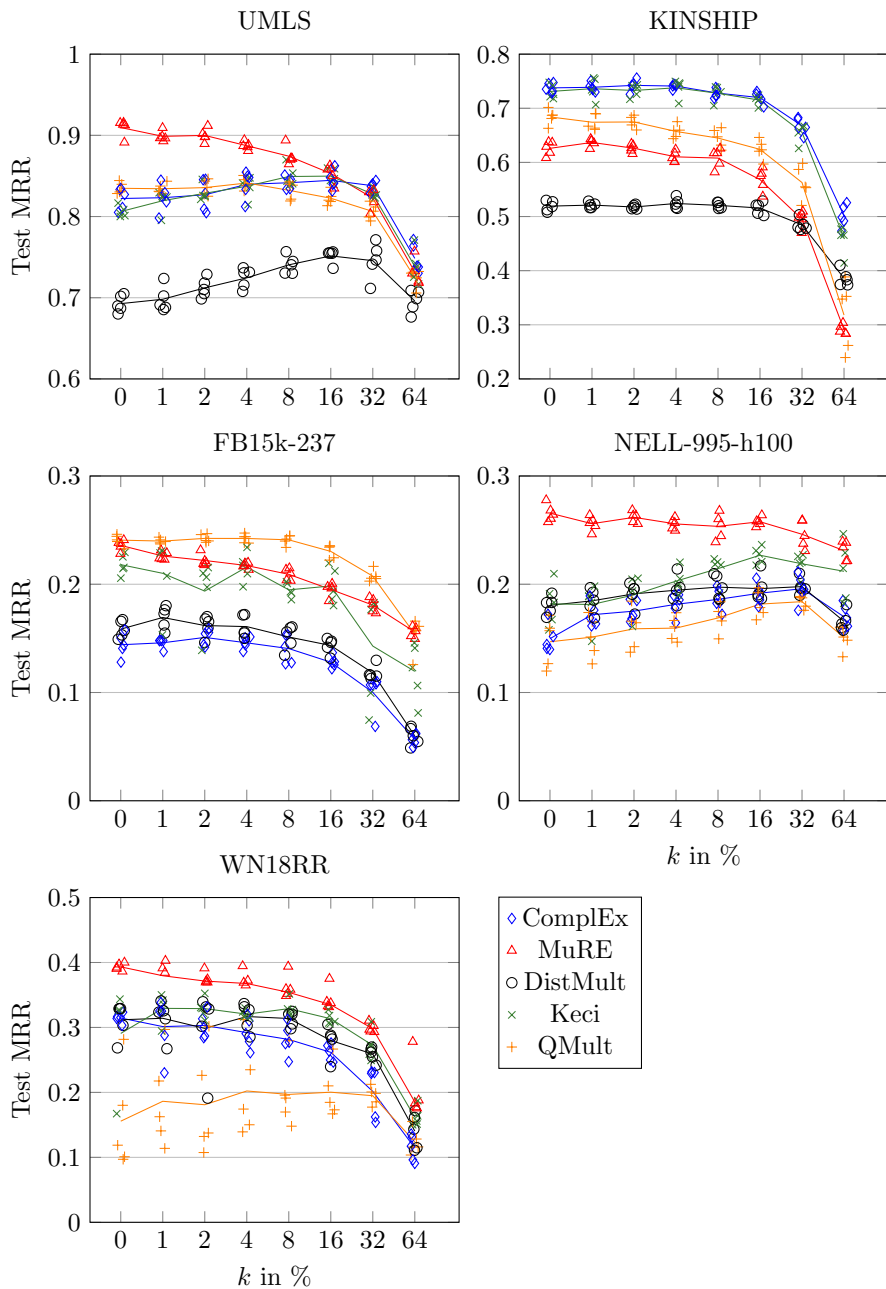


Figure 1: Test MRR performance of the KGE approaches on different datasets with Graph Perturbation and varying perturbation ratios.

shows to be sensitive to the seed value of an internal random number generator, which is again demonstrated by the varied performance of QMult on WN18RR.




Figure 2: Test MRR performance of the KGE approaches on the UMLS and KINSHIP datasets with Label Perturbation and varying perturbation ratios.

5.4. Discussion

The results reported above indicate that all tested KGE algorithms can be influenced by perturbing the data on nearly all datasets. However, the results vary depending on the KGE algorithm, attack surface, and dataset size. Below we further discuss them.

5.4.1. GRAPH PERTURBATION.

The results of the Graph Perturbation experiments allow two conclusions. First, although a non-adversarial attack on the graph input, i.e., on \mathbf{x} , has a negative effect on the performance, the perturbation rate has to be higher than for the other two attack surfaces to achieve a similar reduction. Second, in some cases, a small perturbation showed the opposite effect, i.e., the performance of some KGE algorithms increases, such as DistMult on UMLS and the results of all approaches except MuRE on NELL-995-h100. A similar result can be seen for the MRR measured on the validation split of the datasets. Hence, we conclude that the perturbed data acts like a regularizer in the training process of some KGE algorithms, making them less vulnerable to overfitting and, hence, boosting their performance on the test and validation data. Such behavior is not surprising and there are works by [Orvieto et al. \(2023\)](#) who mention the effect of explicit regularization in the ML model by performing perturbation on the data (e.g., via injecting noise). In this work, we observe a similar effect.

5.4.2. LABEL PERTURBATION.

The attacks on the label vector \mathbf{y} showed the highest impact in our experiments. The reason for the high impact can be explained by comparing an attack on a single label vector




Figure 3: Test MRR performance of the KGE approaches on different datasets with Parameter Perturbation and varying perturbation ratios.

with the number of edges that would have to be added by a Graph Perturbation attack to achieve a similar effect. Consider a training example (\mathbf{x}, \mathbf{y}) in which the label vector has only a single 1 and all other values are 0. If this vector is inverted, the attack has an effect of adding $|\mathcal{E}| - 1$ edges to the graph. For example, for the WN18RR dataset, changing a

single label vector would add more than 40k false edges to a graph that contains 86k edges. After changing a single vector, nearly 1/3 of the training data that the KGE algorithms rely on becomes faulty. This effect is bigger, the larger and the more sparse the graphs are. The two smaller graphs UMLS and KINSHIP have a small number of entities and a high node degree with 38.6 and 82.1 edges per node, respectively. The impact of a small perturbation rate is not as big as on the large datasets. WN18RR, NELL-995-h100 and FB15k-237 have a node degree of 2.1, 2.2, and 18.7, respectively. Changing only 0.1% of the label vectors already adds 3.5M, 1.1M, and 3.9M faulty triples to the training data of these datasets, which are many more triples than the size of the training split. Therefore, label perturbation has been shown to be much more effective in degrading the performance as it can entirely change the structure of the underlying KG. After considering the results obtained through the experiments we can conclude that the typical label-flipping attacks suggested by [Song et al. \(2018\)](#); [Zhang et al. \(2022\)](#) for ML algorithms cannot simply be used for KGEs. Future works proposing new techniques pertaining to label-flipping attacks for KGE models need to be studied. Since this would require a significant extension of the current paper we consider it as a possible future work.

5.4.3. PARAMETER PERTURBATION.

Attacks on the parameter surface show a stronger negative impact on the performance of KGE algorithms when compared to the Graph Perturbation attack. When compared to the Label Perturbation attack, the dataset size seems to have a large influence. On larger datasets, such as in WN18RR, NELL-995-h100, and FB15K-237, an attacker has to reach a higher perturbation ratio to achieve an effect consequently making this attack weaker than the Label Perturbation. However, on the two small datasets, i.e., on UMLS and KINSHIP, the opposite is the case. The Parameter Perturbation leads to a larger performance drop with smaller perturbation rates on both datasets. This shows that KGE models learned on the larger datasets are less susceptible to perturbations compared to the models learned on the smaller datasets. One possible explanation for such an outcome can be that the embedding space of the model learned from the larger datasets is broader and hence, a higher perturbation ratio is needed to cause significant changes on the model. On the contrary, the models learned on smaller datasets have less widened embedding spaces and even a little perturbation can lead to a large impact on the performance.

6. Conclusion

In this work, we have introduced non-adversarial attacks considering three attack surfaces of KGE models. We have performed such non-adversarial attacks on 5 state-of-the-art KGE algorithms considering 5 datasets across 3 attack surfaces, considering 8 different perturbation ratios. Our results suggest that non-adversarial attacks on different surfaces have different rates of performance degradation changes. While attacking the graph by considering lower perturbation ratios can lead to performance improvements, the same ratio can completely degrade the performance when considering the label perturbation.

Therefore, the findings emphasize the importance of evaluating KGE models against different types of perturbations to ensure their robustness, especially if they are to be deployed in dynamic environments where the input data or the model parameters might be

subject to variations. The goal would be to develop KGE models that not only perform well under ideal conditions but can also withstand and adapt to unexpected changes in their operational parameters. Potential approaches could include the development of models that inherently account for parameter variability, the use of robust optimization techniques, or the implementation of adaptive learning rates that could mitigate the impact of high perturbation ratios. Moreover, we envision future research exploring how perturbations can be leveraged to improve KGE model performance effectively. This study serves as an initial step toward a broader investigation into enhancing the robustness of KGE models.

References

- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary G. Ives. Dbpedia: A nucleus for a web of open data. In *The Semantic Web, 6th International Semantic Web Conference ISWC*, 2007.
- Jiawang Bai, Baoyuan Wu, Yong Zhang, Yiming Li, Zhifeng Li, and Shu-Tao Xia. Targeted attack against deep neural networks via flipping limited weight bits. In *9th International Conference on Learning Representations, ICLR*, 2021.
- Ivana Balažević, Carl Allen, and Timothy M Hospedales. Hypernetwork knowledge graph embeddings. In *Artificial Neural Networks and Machine Learning–ICANN, 2019*, pages 553–565, 2019a.
- Ivana Balažević, Carl Allen, and Timothy M Hospedales. Tucker: Tensor factorization for knowledge graph completion. *arXiv preprint arXiv:1901.09590*, 2019b.
- Ivana Balažević, Carl Allen, and Timothy Hospedales. Multi-relational poincaré graph embeddings, 2019.
- Peru Bhardwaj, John D. Kelleher, Luca Costabello, and Declan O’Sullivan. Poisoning knowledge graph embeddings via relation inference patterns. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP*, 2021a.
- Peru Bhardwaj, John D. Kelleher, Luca Costabello, and Declan O’Sullivan. Adversarial attacks on knowledge graph embeddings via instance attribution methods. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing, EMNLP*, 2021b.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Durán, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS’13*, page 2787–2795, Red Hook, NY, USA, 2013a. Curran Associates Inc.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems*, 26, 2013b.

- Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. Low-dimensional hyperbolic knowledge graph embeddings. *arXiv preprint arXiv:2005.00545*, 2020.
- Kewei Cheng, Yikai Zhu, Ming Zhang, and Yizhou Sun. Noigan: Noise aware knowledge graph embedding with adversarial learning. In *ICLR 2020 Conference*, 2020. URL <https://api.semanticscholar.org/CorpusID:226951634>.
- Jeffrey Dalton, Laura Dietz, and James Allan. Entity query feature expansion using knowledge base links. In *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2014.
- Caglar Demir and Axel-Cyrille Ngomo. Clifford embeddings—a generalized approach for embedding in normed algebras. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, 2023.
- Caglar Demir and Axel-Cyrille Ngonga Ngomo. Convolutional complex knowledge graph embeddings. In *The Semantic Web: 18th International Conference, ESWC 2021, Virtual Event, June 6–10, 2021, Proceedings 18*, pages 409–424. Springer, 2021.
- Caglar Demir and Axel-Cyrille Ngonga Ngomo. Hardware-agnostic computation for large-scale knowledge graph embeddings. *Software Impacts*, 13:100377, 2022.
- Caglar Demir, Diego Moussallem, Stefan Heindorf, and Axel-Cyrille Ngonga Ngomo. Convolutional hypercomplex embeddings for link prediction. In *Asian Conference on Machine Learning*, pages 656–671. PMLR, 2021.
- Woodrow W. Denham. The detection of patterns in alyawarra nonverbal behavior. In *The Detection of Patterns in Alyawarra Nonverbal Behavior, Semantic Scholar*, 2014. URL <https://api.semanticscholar.org/CorpusID:140416458>.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- David A. Ferrucci, Eric W. Brown, Jennifer Chu-Carroll, James Fan, David Gondek, Aditya Kalyanpur, Adam Lally, J. William Murdock, Eric Nyberg, John M. Prager, Nico Schlaefler, and Christopher A. Welty. Building watson: An overview of the deepqa project. *AI Mag.*, 31(3):59–79, 2010.
- Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2018.
- Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d’Amato, Gerard de Melo, Claudio Gutierrez, Sabrina Kirrane, José Emilio Labra Gayo, Roberto Navigli, Sebastian Neumaier, et al. Knowledge graphs. *ACM Computing Surveys (CSUR)*, 54(4):1–37, 2021.

- Keita Kurita, Paul Michel, and Graham Neubig. Weight poisoning attacks on pretrained models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL*, 2020.
- Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. Trojaning attack on neural networks. In *25th Annual Network and Distributed System Security Symposium, NDSS*. The Internet Society, 2018.
- A. T. McCray. An upper-level ontology for the biomedical domain. *Comp Funct Genomics*, 4(1):80–84, 2003. doi: 10.1002/cfg.255.
- Mojtaba Nayyeri, Sahar Vahdati, Emanuel Sallinger, Mirza Mohtashim Alam, Hamed Shariat Yazdi, and Jens Lehmann. Pattern-aware and noise-resilient embedding models. In *Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR*, 2021.
- Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. A novel embedding model for knowledge base completion based on convolutional neural network. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 2018.
- Maximilian Nickel, Kevin Murphy, Volker Tresp, and Evgeniy Gabrilovich. A review of relational machine learning for knowledge graphs. *Proceedings of the IEEE*, 104(1):11–33, 2015.
- Antonio Orvieto, Anant Raj, Hans Kersting, and Francis R. Bach. Explicit regularization in overparametrized models via noise injection. In *International Conference on Artificial Intelligence and Statistics*, 2023.
- Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017.
- Pouya Pezeshkpour, Yifan Tian, and Sameer Singh. Investigating robustness and interpretability of link prediction via adversarial modifications. In *1st Conference on Automated Knowledge Base Construction, AKBC*, 2019.
- Scott E. Reed, Honglak Lee, Dragomir Anguelov, Christian Szegedy, Dumitru Erhan, and Andrew Rabinovich. Training deep neural networks on noisy labels with bootstrapping. In *3rd International Conference on Learning Representations, ICLR*, 2015.
- Daniel Ruffinelli, Samuel Broscheit, and Rainer Gemulla. You CAN teach an old dog new tricks! on training knowledge graph embeddings. In *8th International Conference on Learning Representations, ICLR*. OpenReview.net, 2020.
- Yingchun Shan, Chenyang Bu, Xiaojian Liu, Shengwei Ji, and Lei Li. Confidence-aware negative sampling method for noisy knowledge graph embedding. In *2018 IEEE International Conference on Big Knowledge (ICBK)*, pages 33–40, 2018. doi: 10.1109/ICBK.2018.00013.

- Q. Song, H. Jin, X. Huang, and X. Hu. Multi-label adversarial perturbations. In *2018 IEEE International Conference on Data Mining (ICDM)*, 2018.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2016.
- Kristina Toutanova and Danqi Chen. Observed versus latent features for knowledge base and text inference. In *Proceedings of the 3rd Workshop on Continuous Vector Space Models and their Compositionality*. Association for Computational Linguistics, 2015.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *International conference on machine learning*, pages 2071–2080. PMLR, 2016.
- Ruobing Xie, Zhiyuan Liu, Fen Lin, and Leyu Lin. Does william shakespeare REALLY write hamlet? knowledge representation learning with confidence. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence*, 2018.
- Wenhan Xiong, Thien Hoang, and William Yang Wang. DeepPath: A reinforcement learning method for knowledge graph reasoning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, September 2017.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575*, 2014.
- Xiaoyu You, Beina Sheng, Daizong Ding, Mi Zhang, Xudong Pan, Min Yang, and Fuli Feng. Mass: Model-agnostic, semantic and stealthy data poisoning attack on knowledge graph embedding. In *Proceedings of the ACM Web Conference, WWW*, 2023.
- Hengtong Zhang, Tianhang Zheng, Jing Gao, Chenglin Miao, Lu Su, Yaliang Li, and Kui Ren. Data poisoning attack against knowledge graph embedding. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI*, 2019a.
- Peng-Fei Zhang, Zi Huang, Xin Luo, and Pengfei Zhao. Robust learning with adversarial perturbations and label noise: A two-pronged defense approach. In *Association for Computing Machinery*, 2022. ISBN 9781450394789.
- Shuai Zhang, Yi Tay, Lina Yao, and Qi Liu. Quaternion knowledge graph embeddings. *Advances in neural information processing systems*, 32, 2019b.
- Zhao Zhang, Fuzhen Zhuang, Hengshu Zhu, Chao Li, Hui Xiong, Qing He, and Yongjun Xu. Towards robust knowledge graph embedding via multi-task reinforcement learning. *IEEE Trans. Knowl. Data Eng.*, 35(4):4321–4334, 2023.