

Simple and Effective Relation-based Embedding Propagation for Knowledge Representation Learning

Huijuan Wang, Siming Dai, Weiyue Su, Hui Zhong, Zeyang Fang, Zhengjie Huang, Shikun Feng, Zeyu Chen, Yu Sun and Dianhai Yu
Baidu, Inc.

{wanghuijuan03, daisiming, suweiyue, zhonghui03, fangzeyang, huangzhengjie, fengshikun01, chenzeyu01, sunyu02, yudianhai}@baidu.com

Abstract

Relational graph neural networks have garnered particular attention to encode graph context in knowledge graphs (KGs). Although they achieved competitive performance on small KGs, how to efficiently and effectively utilize graph context for large KGs remains an open problem. To this end, we propose the *Relation-based Embedding Propagation* (REP) method. It is a post-processing technique to adapt pre-trained KG embeddings with graph context. As relations in KGs are directional, we model the incoming head context and the outgoing tail context separately. Accordingly, we design relational context functions with no external parameters. Besides, we use averaging to aggregate context information, making REP more computation-efficient. We theoretically prove that such designs can avoid information distortion during propagation. Extensive experiments also demonstrate that REP has significant scalability while improving or maintaining prediction quality. Notably, it averagely brings about 10% relative improvement to triplet-based embedding methods on OGBL-WikiKG2 and takes 5%-83% time to achieve comparable results as the state-of-the-art GC-OTE.

1 Introduction

Knowledge graphs (KGs) are potentially valuable for many applications, such as question answering [Saxena *et al.*, 2020; Bian *et al.*, 2021] and recommender systems [Joseph and Jiang, 2019]. Although KGs contain rich facts in the form of triplets (*head entity, relation, tail entity*), they are still far from complete, necessitating a demand for KG completion. A promising approach is to embed KGs into latent space and predict missing facts from existing ones. Triplet-based methods [Bordes *et al.*, 2013; Yang *et al.*, 2015; Sun *et al.*, 2019; Tang *et al.*, 2020] often regard relations as operations over entity embedding space. They can quickly scale to large KGs because of their elegant simplicity and good interpretability. However, triplet-based methods overlook global neighbor information. Recently, considerable literature has grown around graph neural networks (GNNs) to model such graph contexts. Existing studies for multi-relation

KGs [Schlichtkrull *et al.*, 2018; Vashishth *et al.*, 2020; Cai *et al.*, 2019; Wang *et al.*, 2021] usually train relational GNNs from scratch and optimize triplet-based objects. Specifically, these context-based methods update the central entity with the aggregated embeddings of connected entities and relations. Thus, the time complexity is proportional to the number of triplets and the average degree of entities. The high complexity hinders their applications on real-world tasks at a large scale. How to efficiently and effectively utilize graph context in large KGs remains a challenge.

To this end, we propose *Relation-based Embedding Propagation* (REP) to combine triplet-based methods' simplicity and graph context strength. On the one hand, we derive inspiration from simplified GNN methods such as SGC [Wu *et al.*, 2019] and APPNP [Klicpera *et al.*, 2019] and remove unnecessary non-linearity. Instead of parameterizing the aggregation of neighbors with relation-specific transformation, we use triplet-based methods to capture triplet information firstly and then adapt the pre-trained embeddings with graph context. We do not use layer-specific transform matrices during embedding propagation to reduce time and space complexities. Our method incorporates graph structures after parameter training and has no activation function or backward propagation during embedding propagation, making it more computationally efficient.

On the other hand, simple embedding propagation that ignores relations, such as SGC and APPNP, is not suitable for KGs and will cause performance degeneration. Therefore, we propose a novel way of combining neighbor relations and entities. We make opposite assumptions about the head context and the tail context. For example, we regard incoming relations as additions and outgoing relations as subtractions. Then context embeddings are aggregated from neighbor entity-relation pairs. We take an update scalar to balance pre-trained embeddings and relational context embeddings. Besides, we provide theoretical and empirical evidence that our proposed method of incorporating relations can help REP obtain an accuracy gain.

We show that REP-OTE takes 5%-83% time to achieve comparable performance to the state-of-the-art GC-OTE [Tang *et al.*, 2020] on small datasets through a series of experiments on the link prediction task. Our REP also works well on the medium and the large KGs. In particular, it brings about a 10% relative improvement on the medium OGBL-

WikiKG2 [Hu *et al.*, 2020], despite taking less than 20% time for computation. The key advantages of our REP are three-fold:

- By designing the non-parametric embedding propagation without backward propagation, REP becomes computation-efficient and straightforward in utilizing graph context.
- In order to avoid information distortion during embedding propagation in KGs, we incorporate valuable relations using different triplet assumptions so that REP can maintain or mostly improve prediction quality.
- The proposed REP has significant scalability and takes less time to achieve comparable performance on large KGs, promising practical applications.

2 Related Work

2.1 Knowledge Representation Learning

Knowledge representation learning has been used to embed entities and relations in KGs into latent space and then infer missing facts based on existing ones. The literature falls into two major categories: (1) triplet-based methods that model triplet plausibility by score functions; (2) context-based methods that model graph context beyond triplets.

Triplet-based methods often make assumptions about entities and relations in triplets. Some regard relations as operations between connected entities in the latent space, such as addition [Bordes *et al.*, 2013], rotation [Sun *et al.*, 2019], and transformations in higher dimensional space [Zhang *et al.*, 2019; Tang *et al.*, 2020]. Others follow semantic matching and design score functions based on similarities [Yang *et al.*, 2015; Trouillon *et al.*, 2016]. In general, these methods are of elegant simplicity and good interpretability.

Context-based methods usually inherit the message passing framework from GNNs [Gilmer *et al.*, 2017]. As heterogeneous relations always play an essential role in KGs, modeling such information is a promising direction. GC-OTE [Tang *et al.*, 2020] expects the central embeddings close to the aggregation of neighbors with relation-specific orthogonal transformation. R-GCN [Schlichtkrull *et al.*, 2018] extends GCN by introducing relation-specific weight matrices. CompGCN [Vashishth *et al.*, 2020] follows the idea of modeling relations as composition operations. However, it makes the same assumption about entity-relation pairs in both head and tail graph contexts, which requires external parameters to denote the direction of relations.

2.2 Simplified Graph Neural Networks

Recent progress in simplifying GNNs for homogeneous graphs has led to a surge of improvements in scalability. SGC [Wu *et al.*, 2019] empirically observes that local averaging is more critical and removes non-linearity. APPNP [Klicpera *et al.*, 2019] utilizes a propagation schema derived from personalized PageRank and achieves linear computational complexity in the number of edges. Reminiscent of our method, C&S [Huang *et al.*, 2021] uses graph structure as a post-processing mechanism but propagates labels to correct predictions.




Figure 1: Overview of REP. It consists of two phases: (1) relation-based embedding propagation in the dashed box and (2) entity adaptation by weighted aggregation.

3 Relation-based Embedding Propagation

Our approach starts with simple triplet-based methods, which learn embeddings for entities and relations through minimizing objectives under triplet assumptions. After, REP takes the pre-trained embeddings as input and improves them by incorporating neighbor information in graph context. As illustrated in Figure 1, REP consists of (1) relational graph context functions efficiently and effectively aggregate relation-specific neighbor information; (2) the entity adaptation computes new embeddings based on pre-trained and context embeddings.

Notations. We consider knowledge graph as a collection of triplets $\mathcal{T} = \{(h, r, t) \mid h, t \in E, r \in R\}$ with E as the entity set and R as the relation set. Let A^H be the adjacent head list of the knowledge graph and A^T be the adjacent tail list. Specifically, the head adjacent list A_i^H of the entity i is composed of entity-relation pairs (e, r) with $(e_i, r, e) \in \mathcal{T}$. In contrast, the tail adjacent list A_i^T contains pairs (e, r) with $(e, r, e_i) \in \mathcal{T}$. Furthermore, corresponding bold letters denote embeddings.

3.1 Relation-based Context Function

A key motivation behind our relational context function is that heterogeneous relations in KGs play a crucial role in understanding entities' meanings. Triplet-based methods usually take relations as operations between entities such as additions, multiplications, rotations, and orthogonal transformations. Thus, ignoring relations during propagation can bring information distortion. We design relational context functions according to their assumptions to keep consistent with pre-trained embeddings. As shown in Figure 1, there are two kinds of context.

Head Graph Context includes incoming head-relation pairs for a central entity e , where e acts as a tail entity. As triplet-based methods often apply relation operations to head entities, we can directly use their assumptions for the head-relation pairs. We choose four typical triplet-based methods, which regard relations as additions, multiplications, rotations, and orthogonal transformations, respectively, and define corresponding head context functions $g_h(h, r)$ in Table 1.

Tail Graph Context consists of all outgoing tail-relation pairs connected to the central entity. KGs are directed graphs, where a valid (h, r, t) does not mean (t, r, h) is also a real

Method	Score Function $f_r(h, t)$	Head Context Function $g_h(h, r)$	Tail Context Function $g_t(t, r)$	Parameters
TransE	$- \mathbf{h} + \mathbf{r} - \mathbf{t} $	$\mathbf{h} + \mathbf{r}$	$\mathbf{t} - \mathbf{r}$	$\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{R}^n$
DistMult	$\langle \mathbf{r}, \mathbf{h}, \mathbf{t} \rangle$	$\mathbf{h} \odot \mathbf{r}$	$\mathbf{t} \odot \mathbf{r}$	$\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{R}^n$
RotatE	$- \mathbf{h} \odot \mathbf{r} - \mathbf{t} $	$\mathbf{h} \odot \mathbf{r}$	$\mathbf{t} \odot \bar{\mathbf{r}}$	$\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathbb{C}^n$
OTE	$\sum_{i=1}^L \text{diag}(\exp(\mathbf{s}_{r,(i)}))\phi(\mathbf{M}_{r,(i)})\mathbf{h}_{(i)} - \mathbf{t} $	$\sum_{i=1}^L \text{diag}((\exp(\mathbf{s}_{r,(i)}))\phi(\mathbf{M}_{r,(i)})\mathbf{h}_{(i)})$	$\sum_{i=1}^L \text{diag}((\exp(-\mathbf{s}_{r,(i)}))\phi(\mathbf{M}_{r,(i)})^T \mathbf{t}_{(i)})$	$\mathbf{h}, \mathbf{t} \in \mathbb{R}^n$

Table 1: Score functions $f_r(h, t)$ and context functions $g_h(h, r)$ and $g_t(t, r)$ for different triplet-based KG embedding methods, where $\langle \cdot \rangle$ denotes the generalized dot product, \odot denotes the Hadamard product, $\bar{\cdot}$ denotes conjugate for complex vectors, ϕ denotes the Gram Schmidt process, \mathbf{M}_r denotes orthogonal matrix for r , and L denotes the number of parameter groups.

fact. Therefore, we can not use the same assumptions as of the head graph context. Propagating information of tails and relations to heads utilize de facto inverse edges (t, r^{-1}, h) in KGs. These inverse edges are usually modeled through augmenting relations with inverse ones [Dettmers *et al.*, 2018] or using external parameters to represent relations' directions [Schlichtkrull *et al.*, 2018], which requires external computing resources. To make our method concise and informative, we make opposite assumptions about the tail-relation pairs. Specifically, embeddings of tails and relations are combined by the inverse operations of those for the head graph context, namely subtraction, multiplication, inverse rotation, and inverse orthogonal transformation. In this way, the tail context functions $g_t(t, r)$ are defined in Table 1.

Here we introduce how the context functions are derived in detail. For the methods that minimize the distance between \mathbf{t} and the combination of \mathbf{h} and \mathbf{r} , we set their score functions to the desired 0. Then we get $||\mathbf{h} + \mathbf{r} - \mathbf{t}|| = 0$ for addition-based TransE, $||\mathbf{h} \odot \mathbf{r} - \mathbf{t}|| = 0$ for rotation-based RotatE, and $\sum_{i=1}^L ||\text{diag}(\exp(\mathbf{s}_{r,(i)}))\phi(\mathbf{M}_{r,(i)})\mathbf{h}_{(i)} - \mathbf{t}|| = 0$ for orthogonal transformation-based OTE. After, we can deduce \mathbf{h} and \mathbf{t} from these equations. In particular, the relation matrices in OTE are orthogonal so that we can get their inverse by simple transpose. For DistMult that maximizes generalized dot product, we expect two elements as close as possible: either $\mathbf{t} = \mathbf{h} \odot \mathbf{r}$ or $\mathbf{h} = \mathbf{t} \odot \mathbf{r}$.

Context Aggregation aims to combine neighbor information. As this work aims to improve the scalability of graph context-based methods, we use averaging for aggregation, which helps to keep embeddings on the same scale. Then the head context embedding $\tilde{\mathbf{h}}_i^{(k)}$ of e_i is defined as

$$\tilde{\mathbf{h}}_i^{(k)} = \frac{1}{|A_i^H|} C^H(\mathbf{e}_i^{(k)}) = \frac{1}{|A_i^H|} \sum_{e, r \in A_i^H} g_h(\mathbf{e}^{(k)}, r), \quad (1)$$

and the tail context embedding $\tilde{\mathbf{t}}_i^{(k)}$ is

$$\tilde{\mathbf{t}}_i^{(k)} = \frac{1}{|A_i^T|} C^T(\mathbf{e}_i^{(k)}) = \frac{1}{|A_i^T|} \sum_{e, r \in A_i^T} g_t(\mathbf{e}^{(k)}, r). \quad (2)$$

3.2 Entity Adaptation

The propagation schema derived from personalized PageRank [Klicpera *et al.*, 2019] can be formulated as Eq. (3).

$$\mathbf{E}^{(k+1)} = (1 - \alpha)\mathcal{A}\mathbf{E}^{(k)} + \alpha\mathbf{E}^{(0)}, \quad (3)$$

where $\mathbf{E}^{(0)}$ is the input feature matrix, and \mathcal{A} is the symmetrically normalized adjacency matrix with self-loops. Here we denote embeddings after k updates as $\mathbf{E}^{(k)}$.

However, storing both $\mathbf{E}^{(k)}$ and $\mathbf{E}^{(0)}$ is very demanding on memory for large KGs. We do not store the input feature matrix $\mathbf{E}^{(0)}$, i.e., pre-trained embeddings in our method, but use them for embedding initialization. Besides, we separate self-loops from \mathcal{A} and use different weights for entity and context embeddings to retain inherent information in pre-trained embeddings. Especially, Eq. (3) can not utilize relations in KGs, so we design relational context functions in previous Section 3.1.

In summary, we use an update scalar $\alpha \in [0, 1)$ to balance the trade-off between the pre-learned triplet information and the graph context information. Then the final embedding of entity e_i is computed according to Eq. (4).

$$\mathbf{e}_i^{(k+1)} = \alpha\mathbf{e}_i^{(k)} + \frac{1 - \alpha}{|A_i^H| + |A_i^T|} (C^H(\mathbf{e}_i^{(k)}) + C^T(\mathbf{e}_i^{(k)})), \quad (4)$$

where $\mathbf{e}^{(0)}$ is the pre-trained entity embeddings. In particular, instead of updating parameters with gradient descent algorithms, we directly use the computation results for link prediction. As the number of relations is relatively small, we fix the pre-trained embeddings of heterogeneous relations.

3.3 Theoretical Analysis

In this section, we provide theoretical analysis from the parameter update perspective. Generally, triplet-based methods start with their score function $f_r(h, t)$ and learn embeddings through minimizing the margin-based ranking criterion:

$$\mathcal{L} = -\mathbb{E}_{(h,r,t) \in \mathcal{T}} [\gamma + f_r(\mathbf{h}, \mathbf{t}) - f_r(\mathbf{h}', \mathbf{t}')]_+, \quad (5)$$

where γ denotes the parameter margin, $[x]_+$ takes the positive part of x . A negative sample (h', r, t') is constructed by uniformly replacing h or t in (h, r, t) with other entities in E . Parameters of embeddings are updated by gradient descent methods. The typical stochastic gradient descent algorithm can be formulated as Eq. (6).

$$\mathbf{e}'_i = \mathbf{e}_i - \beta \frac{\partial \mathcal{L}}{\partial \mathbf{e}_i}, \quad (6)$$

where β denotes the learning rate. The motivation behind the margin-based ranking criterion is to make valid triplets have relatively higher scores by enlarging the distance to negative samples. Therefore, when embeddings are fully optimized by this criterion, the distances between valid triplets and negative samples reach a local optimum.

Proof. Here we prove that our REP further improves the objective \mathcal{L} by maximizing scores of valid triplets. Without loss of generality, we take TransE as an example. It interprets relations as additions in latent space, i.e., \mathbf{t} is expected to be the

Dataset	# Entities	# Relations	# Triplets
WN18RR	40,943	11	86,835
FB15k-237	14,541	237	272,115
OGBL-WikiKG2	2,500,604	535	16,109,182
WikiKG90M-LSC	87,143,637	1,315	504,220,369

Table 2: Statistics of knowledge graphs on link prediction.

nearest neighbor of $\mathbf{h} + \mathbf{r}$. Accordingly, the score function is defined as $f_r(h, t) = -\|\mathbf{h} + \mathbf{r} - \mathbf{t}\|$. Then the objective of maximizing scores of valid triplets is formulated as Eq. (7).

$$\mathcal{J} = -\mathbb{E}_{(h_i, r_i, t_i) \in \mathcal{T}} \|\mathbf{h}_i + \mathbf{r}_i - \mathbf{t}_i\|. \quad (7)$$

Assuming that parameters are updated by stochastic gradient descent method, we first calculate the partial derivative for h_i as shown in Eq. (8).

$$\begin{aligned} \frac{\partial \mathcal{J}}{\partial h_i} &= 2\mathbb{E}_{(h_j, r_j, t_j) \in \mathcal{T} \wedge (h_i = h_j)} (\mathbf{h}_j + \mathbf{r}_j - \mathbf{t}_j) \\ &= 2\mathbf{h}_i + 2\mathbb{E}_{(h_j, r_j, t_j) \in \mathcal{T} \wedge (h_i = h_j)} (\mathbf{r}_j - \mathbf{t}_j). \end{aligned} \quad (8)$$

In this case, Eq. (6) can be reformulated as Eq. (9) by replacing the partial derivative with Eq. (8).

$$\mathbf{h}'_i = (1 - 2\beta)\mathbf{h}_i + 2\beta\mathbb{E}_{(h_j, r_j, t_j) \in \mathcal{T} \wedge (h_i = h_j)} (\mathbf{t}_j - \mathbf{r}_j). \quad (9)$$

Let $\alpha = 1 - 2\beta$, this update function is equivalent to Eq. (4) when entities act as heads, where the head context function is $\mathbf{t} - \mathbf{r}$. For the tail entity t_i , we have

$$\mathbf{t}'_j = (1 - 2\beta)\mathbf{t}_j + 2\beta\mathbb{E}_{(h_j, r_j, t_j) \in \mathcal{T} \wedge (t_i = t_j)} (\mathbf{h}_j + \mathbf{r}_j), \quad (10)$$

where the addition operation is consistent with the tail context function in Table 1. Above all, REP acts the same as using stochastic gradient descent to maximize valid triplets' scores, which further optimizes the triplet-based objective \mathcal{L} of pre-trained embeddings. \square

4 Experiments

4.1 Datasets

We use four datasets at diverse scales, as reported in Table 2.

Large scale. WikiKG90M-LSC [Hu *et al.*, 2021] of KDD Cup 2021¹ is extracted from Wikidata. Aside from triplets, it provides text features learned by RoBERTa [Liu *et al.*, 2019], which are combined by concatenation when training triplet-based methods. It also provides 1000 randomly sampled entity candidates for each tail in the validation/test set.

Medium scale. OGBL-WikiKG2 [Hu *et al.*, 2020] is a link prediction dataset from Open Graph Benchmark (OGB)². It is constructed from the relational statements in Wikidata after filtering out rare entities. Moreover, 1000 randomly sampled candidates for tail entities are provided likewise.

Small scale. FB15k-237 and WN18RR [Dettmers *et al.*, 2018] are the most widely used datasets for link prediction. FB15k-237 mostly describes people-related facts, while WN18RR consists of hyponym and hypernym relations.

¹<https://ogb.stanford.edu/kddcup2021/>

²<https://ogb.stanford.edu/>

Dataset	REP-OTE	GC-OTE	Speedup
WN18RR	16s	19s	1.2
FB15k-237	96s	1959s	20.4
OGBL-WikiKG2	1.7h	> 100h	> 50

Table 3: Time complexity comparison between REP-OTE and GC-OTE. The values denote the time cost for each epoch.

4.2 Evaluation Metrics

Link prediction aims to predict missing facts based on existing triplets, namely to predict valid heads for (t, r) or valid tails for (h, r) . Specifically, we first corrupt test triplets and construct candidate triplets using entity candidates. As small datasets do not provide candidates, we use all entities e that $(h, r, e) \notin \mathcal{T}$ and $(e, r, t) \notin \mathcal{T}$ as candidates for a test triplet (h, r, t) . Then we compute the plausibility of the test triplet and its candidate triplets and sort them by descending order. Evaluation metrics are based on the rank of test triplets, including Mean Reciprocal Rank (MRR) and Hits@K which denote the proportion of ranks less than K.

4.3 Implementation Details

Experiments were conducted on Intel Xeon Gold 6271C CPUs and Tesla V100 SXM2 GPUs. We use the public code to reproduce triplet-based methods, while the results of context-based methods are from original papers. The code and details of REP are released in Graph4KG³.

4.4 Results and Discussion

Simplicity. For model parameters, REP requires no parameters except for embeddings of entities and relations, while context-based methods inherit from GNNs and require multiple layer-specific weight matrices besides embeddings. For the computation complexity, we conduct speed experiments to compare REP-OTE and GC-OTE. GC-OTE also requires no external parameters. In theory, REP-OTE and GC-OTE have the same time complexity $\mathcal{O}(|\mathcal{T}|\mathcal{D})$, where $|\mathcal{T}|$ is the number of triplets and \mathcal{D} denotes the average degree. However, as REP has no backward propagation, its constant is far smaller than GC-OTE. The speedup numbers in Table 3 empirically prove this. We count the time consumed to traverse all training data. To make a fair comparison, the time of REP-OTE consists of the time REP takes and the time used by OTE in an epoch during pre-training. We can note that REP brings more acceleration when the data scale is larger.

Effectiveness. Results of medium and large datasets are reported in Table 4. We have observed that when there are more than 80K entities and 1.6M triplets, context-based methods like CompGCN will be out of memory on a 32GB GPU to learn embeddings with dimension 200. Thus, we only report their results on small-scale datasets. As we can see, REP-augmented methods perform best on all metrics. In particular, REP-TransE significantly surpasses TransE and is competitive with OTE, the state-of-the-art triplet-based method with

³<https://github.com/PaddlePaddle/PGL/tree/main/apps/Graph4KG>

Base Model	Strategy	OGBL-WikiKG2				WikiKG90M-LSC			
		MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
TransE [Bordes <i>et al.</i> , 2013]	-	0.426	0.397	0.429	0.468	0.881	0.845	0.908	0.947
	REP	0.528	0.481	0.542	0.619	0.884	0.850	0.912	0.947
DistMult [Yang <i>et al.</i> , 2015]	-	0.369	0.322	0.374	0.453	0.863	0.834	0.885	0.918
	REP	0.376	0.329	0.380	0.461	0.868	0.838	0.891	0.925
RotatE [Sun <i>et al.</i> , 2019]	-	0.434	0.399	0.441	0.491	0.886	0.850	0.914	0.952
	REP	0.491	0.438	0.507	0.593	0.889	0.854	0.916	0.953
OTE [Tang <i>et al.</i> , 2020]	-	0.519	0.475	0.532	0.597	0.910	0.881	0.933	0.961
	REP	0.538	0.490	0.554	0.624	0.914	0.884	0.938	0.966

Table 4: Evaluation results on the medium-scale (OGBL-WikiKG2) and large-scale (WikiKG90M-LSC) KGs on link prediction.

Model	FB15k-237				WN18RR			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
TransE [Bordes <i>et al.</i> , 2013]	0.329	0.230	0.368	0.528	0.226	0.016	0.403	0.531
DistMult [Yang <i>et al.</i> , 2015]	0.307	0.219	0.336	0.483	0.440	0.394	0.454	0.532
RotatE [Sun <i>et al.</i> , 2019]	0.336	0.240	0.374	0.531	0.477	0.432	0.494	0.567
OTE [Tang <i>et al.</i> , 2020]	0.352	0.258	0.387	0.539	0.487	0.438	0.502	0.587
PairRE [Chao <i>et al.</i> , 2021]	0.351	0.256	0.387	0.544	0.454	0.411	0.469	0.548
R-GCN [Schlichtkrull <i>et al.</i> , 2018]	0.248	0.151	-	0.417	-	-	-	-
SACN [Shang <i>et al.</i> , 2019]	0.350	0.260	0.390	0.540	0.470	0.430	0.480	0.540
TransGCN [Cai <i>et al.</i> , 2019]	0.356	0.252	0.388	0.555	0.485	0.438	0.510	0.578
CompGCN [Vashishth <i>et al.</i> , 2020]	0.355	0.264	0.390	0.535	0.479	0.443	0.494	0.546
GC-OTE [Tang <i>et al.</i> , 2020]	0.360	0.265	0.398	0.547	0.491	0.442	0.512	0.583
M ² GNN [Wang <i>et al.</i> , 2021]	0.362	0.275	0.398	0.565	0.485	0.444	0.498	0.572
REP-OTE	0.354	0.262	0.388	0.540	0.488	0.439	0.505	0.588

Table 5: Prediction results of different models on FB15k-237 and WN18RR. The best results are bold.

higher complexity. It averagely achieves a 10% relative improvement on OGBL-WikiKG2 compared with four triplet-based baselines. This result may suggest that graph context is of great use for knowledge graphs at a medium scale. The improvement of WikiKG90M-LSC is not as significant as that on OGBL-WikiKG2. We reason that the text features used during pre-training provide external information to entities with few triplets, so graph context is possible to be superfluous. Results on small datasets are reported in Table 5. In this case, REP-OTE outperforms all triplet-based methods and has comparable prediction quality to the state-of-the-art context-based methods. Above all, REP can bring stable improvements to triplet-based methods and achieve competitive performance as context-based methods with less cost.

Practical Effect. We assume in previous experiments that triplet-based methods have fully trained pre-trained embeddings. In this section, we explore how REP performs on not converged embeddings. Specifically, we apply REP to the embeddings trained by triplet-based methods after 0.25N, 0.5N, 0.75N, and N steps separately if the embeddings converge at step N. The MRR results of these embeddings are reported in Figure 2. Experiments on the four triplet-based methods show that REP brings a significant performance boost to not converged embeddings. Particularly, triplet-based methods with REP can achieve comparable results as the converged ones with just half of the training steps. In




Figure 2: The MRR performance of REP on four triplet-based methods at different training stages.

other words, REP allows obtaining high-quality embeddings within less time.

How to select α ? The update scalar α determines how much neighbor information contributes to the central entities. To investigate the impact of α , we set α as values ranging from 0.95 to 0.99 with step 0.01. Then we evaluate embed-




Figure 3: MRR results of triplet-based methods plus REP on FB15k-237 with different hyper-parameters.




Figure 4: Results of triplet-based methods plus REP at different hops on datasets on three different scales.

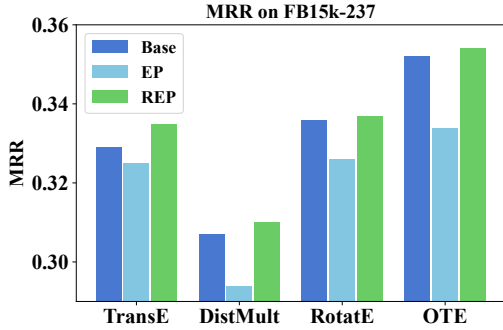


Figure 5: Results of triplet-based methods and the corresponding plus embedding propagation (EP) methods and plus REP methods on the FB15k-237 dataset.

dings after once to 15 times entity adaptations separately and draw heat maps using their MRRs in Figure 3. As illustrated in Section 3.3, $\alpha = 1 - 2\beta$, where β denotes learning rate. Figure 3 also demonstrates this point. When α increases, REP needs more iterations to get the best MRR.

Are multi-hop neighbors always helpful? For REP, the central entity receives neighbor information k -hops away after k times entity adaption. We investigate how multi-hop neighbors influence embedding quality by changing the number of hops from 1 to 20. We set $\alpha = 0.98$ to control variables. Results of TransE and RotatE on datasets at different scales are plotted in Figure 4. As we can see, MRRs of OGBL-WikiKG2 grow almost linearly as the number of hops increases, maybe because of its triplet sparsity. Moreover, the curves of both methods on FB125k-237 and WikiKG90M-LSC are nearly concave. It shows that enlarging the receptive field contributes to embedding quality, but neighbors too far away would hurt performance when triplets are dense, or there is external information.

4.5 Influence of Relations

We have theoretically analyzed the necessity of relations in Section 3.3. In this section, we provide some empirical evidence. *Embedding Propagation* (EP) methods like APPNP have shown strong performance on homogeneous graphs with single relation. However, when applied to KGs, they can not utilize types of relations and regard relations as identity matrices. In order to prove the significance of relations, we conduct ablation studies on the FB15k-237 dataset. Figure 5 plots the evaluation results of four triplet-based methods. Base denotes MRRs of fully trained triplet-based methods. EP propagates entity embeddings to neighbor entities directly and ignores relations during this process. We also apply EP to the pre-trained embeddings as our REP method. As shown in Figure 5, embedding propagation without relations (EP) causes significant performance deterioration. It proves that the relation-based context functions designed in Section 3.1 play a critical role during embedding propagation.

5 Conclusion

We proposed the novel method REP to utilize graph context in KGs during the post-training process. The key idea is incorporating relational graph structure information into pre-trained triplet-based embeddings. For simplicity, we use local averaging to conduct non-parametric embedding propagation without backward propagation. For effectiveness, we design graph context functions for head-relation neighbors and tail-relation neighbors, respectively. As a result, REP can reduce information distortion during embedding propagation in KGs. Furthermore, such design brings REP excellent scalability, which has practical significance in utilizing graph context for large KGs in the real world. Experimental results also show that by enriching pre-trained triplet-based embeddings with graph context, REP improves or maintains prediction quality with less time cost.

Acknowledgments

This work has been supported by National Key Research and Development Program of China (2018AAA0101900).

References

- [Bian *et al.*, 2021] Ning Bian, Xianpei Han, Bo Chen, and Le Sun. Benchmarking knowledge-enhanced common-sense question answering via knowledge-to-text transformation. In *AAAI*, pages 12574–12582. AAAI Press, 2021.
- [Bordes *et al.*, 2013] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. Translating embeddings for modeling multi-relational data. In *NeurIPS*, pages 2787–2795, 2013.
- [Cai *et al.*, 2019] Ling Cai, Bo Yan, Gengchen Mai, Krzysztof Janowicz, and Rui Zhu. Transcgn: Coupling transformation assumptions with graph convolutional networks for link prediction. In *K-CAP*, pages 131–138, 2019.
- [Chao *et al.*, 2021] Linlin Chao, Jianshan He, Taifeng Wang, and Wei Chu. Pairre: Knowledge graph embeddings via paired relation vectors. In *ACL/IJCNLP*, pages 4360–4369, 2021.
- [Dettmers *et al.*, 2018] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. Convolutional 2d knowledge graph embeddings. In *AAAI*, pages 1811–1818, 2018.
- [Gilmer *et al.*, 2017] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *ICML*, volume 70, pages 1263–1272, 2017.
- [Hu *et al.*, 2020] Weihua Hu, Matthias Fey, Marinka Zitnik, Yuxiao Dong, Hongyu Ren, Bowen Liu, Michele Catasta, and Jure Leskovec. Open graph benchmark: Datasets for machine learning on graphs. In *NeurIPS*, 2020.
- [Hu *et al.*, 2021] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. Ogb-lsc: A large-scale challenge for machine learning on graphs. *arXiv preprint arXiv:2103.09430*, 2021.
- [Huang *et al.*, 2021] Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R. Benson. Combining label propagation and simple models out-performs graph neural networks. In *ICLR*, 2021.
- [Joseph and Jiang, 2019] Kevin Joseph and Hui Jiang. Content based news recommendation via shortest entity distance over knowledge graphs. In *WWW*, pages 690–699, 2019.
- [Klicpera *et al.*, 2019] Johannes Klicpera, Aleksandar Bojchevski, and Stephan Günnemann. Predict then propagate: Graph neural networks meet personalized pagerank. In *ICLR*, 2019.
- [Liu *et al.*, 2019] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [Saxena *et al.*, 2020] Apoorv Saxena, Aditay Tripathi, and Partha Talukdar. Improving multi-hop question answering over knowledge graphs using knowledge base embeddings. In *ACL*, pages 4498–4507, 2020.
- [Schlichtkrull *et al.*, 2018] Michael Sejr Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *ESWC*, volume 10843, pages 593–607, 2018.
- [Shang *et al.*, 2019] Chao Shang, Yun Tang, Jing Huang, Jinbo Bi, Xiaodong He, and Bowen Zhou. End-to-end structure-aware convolutional networks for knowledge base completion. In *AAAI*, pages 3060–3067, 2019.
- [Sun *et al.*, 2019] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. Rotate: Knowledge graph embedding by relational rotation in complex space. In *ICLR*, 2019.
- [Tang *et al.*, 2020] Yun Tang, Jing Huang, Guangtao Wang, Xiaodong He, and Bowen Zhou. Orthogonal relation transforms with graph context modeling for knowledge graph embedding. In *ACL*, pages 2713–2722, 2020.
- [Trouillon *et al.*, 2016] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. Complex embeddings for simple link prediction. In *ICML*, pages 2071–2080, 2016.
- [Vashishth *et al.*, 2020] Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha P. Talukdar. Composition-based multi-relational graph convolutional networks. In *ICLR*, 2020.
- [Wang *et al.*, 2021] Shen Wang, Xiaokai Wei, Cicero Nogueira dos Santos, Zhiguo Wang, Ramesh Nallapati, Andrew O. Arnold, Bing Xiang, Philip S. Yu, and Isabel F. Cruz. Mixed-curvature multi-relational graph neural network for knowledge graph completion. In *WWW*, pages 1761–1771, 2021.
- [Wu *et al.*, 2019] Felix Wu, Amauri H. Souza Jr., Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Q. Weinberger. Simplifying graph convolutional networks. In *ICML*, volume 97, pages 6861–6871, 2019.
- [Yang *et al.*, 2015] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. Embedding entities and relations for learning and inference in knowledge bases. In *ICLR*, 2015.
- [Zhang *et al.*, 2019] Shuai Zhang, Yi Tay, Lina Yao, and Qi Liu. Quaternion knowledge graph embeddings. In *NeurIPS*, pages 2731–2741, 2019.