

Generalized Relation Learning with Semantic Correlation Awareness for Link Prediction

Yao Zhang¹, Xu Zhang¹, Jun Wang², Hongru Liang¹,
Wenqiang Lei^{3*}, Zhe Sun⁴, Adam Jatowt⁵, Zhenglu Yang^{1*}

¹ TKLNDST, CS, Nankai University, China

² Ludong University, China ³ National University of Singapore, Singapore

⁴ Computational Engineering Applications Unit, RIKEN, Japan ⁵ University of Innsbruck, Austria
{yaozhang, xuzhang, lianghr, junwang}@mail.nankai.edu.cn, wenqianglei@gmail.com,
zhe.sun.vk@riken.jp, adam.jatowt@uibk.ac.at, yangzl@nankai.edu.cn

Abstract

Developing link prediction models to automatically complete knowledge graphs has recently been the focus of significant research interest. The current methods for the link prediction task have two natural problems: 1) the relation distributions in KGs are usually unbalanced, and 2) there are many unseen relations that occur in practical situations. These two problems limit the training effectiveness and practical applications of the existing link prediction models. We advocate a holistic understanding of KGs and we propose in this work a unified Generalized Relation Learning framework GRL to address the above two problems, which can be plugged into existing link prediction models. GRL conducts a generalized relation learning, which is aware of semantic correlations between relations that serve as a bridge to connect semantically similar relations. After training with GRL, the closeness of semantically similar relations in vector space and the discrimination of dissimilar relations are improved. We perform comprehensive experiments on six benchmarks to demonstrate the superior capability of GRL in the link prediction task. In particular, GRL is found to enhance the existing link prediction models making them insensitive to unbalanced relation distributions and capable of learning unseen relations.

Introduction

Knowledge graphs (KGs), representing facts in semantic graph structures, have been applied to multiple artificial intelligence tasks, e.g., recommendation (Lei et al. 2020a; Wang et al. 2021), dialogue generation (Moon et al. 2019; Lei et al. 2020b), and question answering (Christmann et al. 2019; Zhu et al. 2021). In KGs, facts are formed as triples, (head entity, relation, tail entity), where the head entity is linked to the tail entity via the relation. New knowledge emerges continuously, and hence the issue of incompleteness of KGs has triggered wide research interests in link prediction task, which requires predicting the missing links in KGs (Seyed and David 2018). The mainstream link prediction models (Bordes et al. 2013; Dettmers et al. 2018) learn the embeddings of entities and relations, and then use a score func-

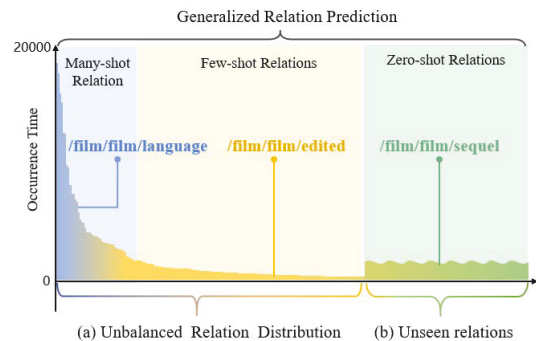


Figure 1: (a) The unbalanced relation distribution in the FB15K-237 dataset where relations are sorted according to their frequency. (b) Lots of unseen relations. Three *film*-related relations are respectively categorized into the many-shot class, few-shot class, and zero-shot class as marked.

tion to estimate the validity of triples. However, we believe using the embedding learning for mainstream link prediction models results in two key problems:

- *Unbalanced relation distribution.* As shown in Figure 1, the relation distribution in an off-the-shelf KG learning resource (i.e., FB15K-237 (Toutanova and Chen 2015)) is quite unbalanced. For example, the frequencies of the two relations */film/film/language* and */film/film/edited* differ greatly. Mainstream link prediction models assume enough training instances for all relations and pay less attention to few-shot relations, disregarding the fact that few-shot relation learning may influence the learning performance to a high degree.
- *Existence of unseen relations.* Real-world KGs tend to be open and evolve quickly, and accordingly there is a large number of zero-shot relations unseen in the off-the-shelf learning resources, for example, the relation */film/film/sequel* in Figure 1. The unseen relations are beyond the capacity of mainstream link prediction models, as there are no training instances to learn their embeddings. This problem may restrict the use of these models in downstream tasks.

*Corresponding authors.

Recently, some efforts have been conducted on addressing the above problems. Xiong et al. (2018), Shi and Weninger (2018), and Chen et al. (2019a) adopted the meta-learning or metric-based approaches to train on limited training samples and perform fast learning on new few-shot relations. These studies show promise in few-shot relation learning, however they have difficulty in tackling unbalanced relation distributions, which is mainly attributed to the excessive time cost required for training numerous relations. More recently, Chen et al. (2019b), Qin et al. (2020) predicted the unseen relations by extracting information from textual descriptions. They were able to successfully complete the unseen relation prediction task. However, these models are not appropriate for link prediction task, since textual descriptions tend to be noisy and also cannot build a bridge between seen and unseen relations. In general, an ideal link prediction model should be able to jointly learn many-, few-, and zero-shot relations.

Regarding the joint relation learning, we noticed that semantic correlations, which denote the similarities of relations in semantics, can serve as a bridge to connect the learning of many-, few-, and zero-shot relations. Take Figure 1 as an instance. The many-shot relation “/film/film/language”, few-shot relation “/film/film/edited”, and zero-shot relation “/film/film/sequel” are all related to “film”. Based on the assumption that semantically similar relations should be located near each other in embedding space (Yang et al. 2015), it makes sense to exploit semantic correlations, such as the one in the above-mentioned example, to accomplish the joint relation learning. Inspired by this, we propose a **Generalized Relation Learning** framework (abbreviated to GRL) based on learning semantic correlations. GRL can be plugged into a mainstream link prediction model to make it (1) insensitive to unbalanced relation distributions and (2) capable of learning zero-shot relations.

Specifically, GRL is plugged into a link prediction model after the embedding learning stage. To optimize the relation embedding, GRL extracts rich semantic correlations through an attention mechanism, fuses different relations, and minimizes the classification-aware loss to enable these implicitly embedded semantic correlations in the relation embeddings. Then, the closeness of semantically similar relations in vector space and the discrimination of dissimilar relations can be improved. In this way, few-shot relations can learn knowledge from the semantically similar many-shot relations; for zero-shot relations, their most semantically similar relation can also be predicted. In our experiments, we improve two base models (DistMult (Yang et al. 2015) and ConvE (Dettmers et al. 2018)) by incorporating the proposed GRL framework on all relation classes, i.e., many, few, and zero-shot relations. Our work is an important step towards a holistic understanding of KGs and a generalized solution of relation learning for the link prediction task.

Our contributions are as follows:

- We carefully consider two key problems of the embedding learning used by mainstream link prediction models and we highlight the necessity of jointly learning many-, few-,

and zero-shot relations.

- We propose GRL framework by leveraging the rich semantic correlations between relations to make the link prediction models insensitive to unbalanced relation distributions and capable of learning zero-shot relations.
- We perform experiments on six benchmarks to evaluate the link prediction capability of GRL, and show that GRL lets the base link prediction models perform well across many-, few-, and zero-shot relations.

Related Work

Since KGs are populated by automatic text processing they are often incomplete, and it is usually infeasible to manually add to them all the relevant facts. Hence, many researches approached the task of predicting missing links in KGs.

Mainstream link prediction models widely use embedding-based methods to map entities and relations into continuous low-dimensional vector space and use a score function to predict whether the triples are valid. They can be broadly classified as translational based (Bordes et al. 2013; Wang et al. 2014; Lin et al. 2015; Ji et al. 2016), multiplicative based (Nickel, Tresp, and Kriegel 2011; Yang et al. 2015; Trouillon et al. 2016), and neural network based (Dettmers et al. 2018; Schlichtkrull et al. 2018). These models are based on the implicit assumption that all relations are distributed within the dataset in a balanced way. Hence, they perform poorly in few-shot relation learning scenarios because these models neglect the imbalanced distributions, as well as they cannot properly handle zero-shot relations due to keeping only the knowledge of existing relations and not learning information on unseen relations.

Few-shot relation learning models attempt to adopt the meta-learning (Chen et al. 2019a; Lv et al. 2019) and metric-based (Xiong et al. 2018; Wang et al. 2019) methods to learn knowledge from only a few samples. However, the few-shot learning models are computationally expensive because they need to spend extra time retraining on each few-shot relation (meta-learning), or need to compare the few-shot relations one by one (metric-based). In practice, the many-shot and few-shot scenarios are not explicitly distinguished. **Zero-shot relation learning models** aim to learn relations that are unseen in the training set. Researchers have proposed several models to deal with zero-shot relations by leveraging information from textual descriptions (Chen et al. 2019b; Qin et al. 2020). They perform well on predicting the zero-shot relations, but are not appropriate in the link prediction task because textual descriptions could be noisy and a bridge connecting seen and unseen relations could be missing.

In this work, we focus on jointly learning many-, few-, and zero-shot relations without requiring extra textual knowledge. Recently, some computer vision works (Ye et al. 2019; Shi et al. 2019) have attempted to approach the generalized image classification. Nonetheless, they are not designed for coping with graph structures, e.g., KGs. We leverage in this work the rich semantic correlations between relations as a bridge to connect the learning of many-, few-, and zero-shot relations. Zhang et al. (2019) integrated the rich

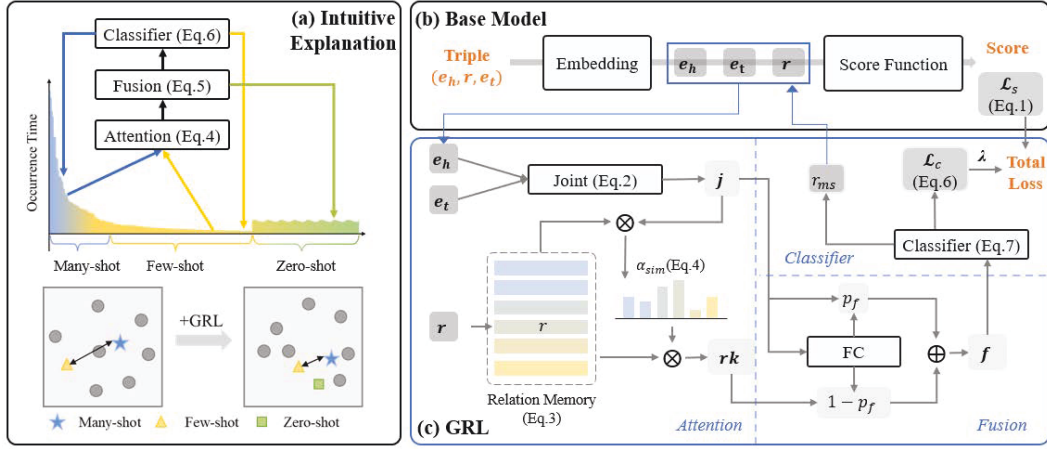


Figure 2: The illustration of GRL, which consists of the intuitive explanation (a), the base model (b) and the detailed architecture (c). The base model denotes the mainstream link prediction model. GRL is plugged after the embedding component of the base model and contains three components: attention, fusion, and classifier.

semantic correlations between specific hierarchical relations into relation extraction. That method performs well only on hierarchical relations, as well as it predicts relations from text, hence it does not cope with the link prediction task.

Method

Figure 2 provides the illustration of the proposed framework GRL. The figure consists of three parts: the intuitive explanation of GRL in Figure 2 (a), base model shown in Figure 2 (b), and the detailed architecture in Figure 2 (c).

The intuitive explanation of GRL is shown to utilize the semantic correlations between many-shot and few-shot relations so that the relation embedding learning can benefit from semantically similar relations. We devise three modules, i.e., *Attention*, *Fusion* and *Classifier*, to embed and fuse the rich semantic correlations among many-shot and few-shot relations in the training phase; and to select the most similar relation embedding for zero-shot relations in the testing phase. In this way, GRL can improve the performance on all relation classes, i.e., many, few, and zero-shot relations. The base model denotes the existing mainstream link prediction model consisting of an embedding component and a score function component. GRL can be plugged between the embedding and the score function components to make it (1) insensitive to imbalanced relation distributions and (2) capable of detecting zero-shot relations.

Before delving into the model description, we first formally represent a KG as a collection of triples $\mathcal{T} = \{(e_h, r, e_t) | e_h \in \mathcal{E}, e_t \in \mathcal{E}, r \in \mathcal{R}\}$, where \mathcal{E} and \mathcal{R} are the entity and relation sets, respectively. Each directed link in KG represents a triple (i.e., e_h and e_t are represented as the nodes and r as the labeled edge between them). The link prediction task is to predict whether a given triple (e_h, r, e_t) is valid or not. In particular, for the zero-shot relations, we need to emphasize that we mainly focus on predicting the validity of the triple with a zero-shot relation, rather than predicting the zero-shot relations, i.e., the relation prediction

task (Chen et al. 2019b; Qin et al. 2020). However, GRL has also the ability to predict the most semantically similar relation of a given zero-shot relation through learning from the many- and few-shot relations, not from the text description.

Base Model

We select a mainstream link prediction model as the base model and apply GRL to it. The base model can be seen as multi-layer neural network consisting of an embedding component and a score function component. For the base link prediction model, given an input triple (e_h, r, e_t) , the embedding component maps the head and tail entities (e_h, e_t) and the relation r to their distributed embedding representations (e_h, r, e_t) through the entity and relation embedding layers, respectively. After the embedding representations are obtained, the score function component is adopted to calculate the likelihood of (e_h, r, e_t) being a valid fact. The following binary cross entropy loss is used to train model parameters:

$$\mathcal{L}_s = -\frac{1}{N} \sum_{i=1}^N (t_i \log p(s_i) + (1 - t_i) \log(1 - p(s_i))), \quad (1)$$

where s_i is the score of the i -th input triple, t_i is the ground truth label, t_i is 1 if the input relation is valid and 0 otherwise, and N is the number of input triples.

GRL Framework

The loss used by mainstream link prediction models is score-oriented and lacks an in-depth exploration of rich semantic correlations in KGs. We propose the GRL framework to learn appropriate representations for relations by embedding semantic correlations into classification-aware optimization. GRL contains three specific modules:

1) *Attention Module*, which builds the knowledge-aware attention distribution and the relational knowledge vector. The aim of this module is to extract the semantic correlations and the degree of these correlations.

2) *Fusion Module*, which fuses the relational knowledge vector with the joint vector obtained from the attention module. This module realizes the fusion of different relations, according to semantic correlations.

3) *Classifier Module*, which calculates the classification-aware loss to implicitly enable the rich semantic correlations embedded in the embeddings. Thanks to it, both the compactness of semantically similar relations and discrimination of dissimilar relations can be enhanced.

The following is a detailed introduction to each module.

Attention Module

Joint Block. The classification-aware loss is calculated by the relation classification results based on the head and tail entities from the given triple (e_h, r, e_t) . Inspired by (Qin et al. 2020), the joint vector of the head and tail entities has the ability to represent the potential relation between them. The head and tail entities representations (i.e., e_h and e_t) are jointed together at the joint block for which we adopt three different alternatives:

$$\mathbf{j} = \begin{cases} e_h - e_t, & \textit{sub} \\ e_h \otimes e_t, & \textit{multiply} \\ W_1[e_h; e_t] + b_1, & \textit{concat} \end{cases}, \quad (2)$$

where \otimes denotes the element-wise multiplication operator, and W_j and b_j are the learnable parameters.

Relation Memory Block. Using a memory block to store class information is widely used in image classification (Snell, Swersky, and Zemel 2017; Karlinsky et al. 2019; Liu et al. 2019). Following these studies, we design a relation memory block to store all relation information by sharing parameters with the relation embedding layer as

$$\mathbf{M} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_{K-1}, \mathbf{r}_K\}, \quad (3)$$

where $M \in \mathbb{R}^{K \times \text{dim}}$, K is the number of relation classes. As the training progresses, the relation embedding layer and relation memory block are updated synchronously.

Relational Knowledge. To realize the classification-aware optimization objective, we extract useful relational knowledge from the relation memory block to enrich the joint vector. The semantic correlation degree between different relations may vary; thus, we adopt the attention mechanism to customize specific relational knowledge for each joint vector. Concretely, the relational knowledge vector \mathbf{rk} is computed as a weighted sum of each relation representation in the relation memory block M , i.e., $\mathbf{rk} = \alpha_{sim} \mathbf{M}$, where $\alpha_{sim} \in \mathbb{R}^K$ represents the knowledge-aware attention distribution.

Attention Distribution The knowledge-aware attention distribution α_{sim} describes the similarity between the joint vector and each relation representation in the relation memory block. We estimate α_{sim} as

$$\alpha_{sim} = \textit{softmax}(\mathbf{j} \mathbf{M}^\top), \quad (4)$$

where *softmax* is the activation function, and \mathbf{M}^\top represents the transposed matrix of \mathbf{M} . Note that the attention value of the ground-truth relation is masked with 0.

Fusion Module

In this module, the joint vector and relational knowledge vector are fused. Intuitively, the proportion of fusion is different for each joint vector. Inspired by the pointer-generator

	$ \mathcal{E} $	$ \mathcal{R} $	Train	Valid	Test
YAGO3-10	123k	37	1M	5k	5k
FB15K-237	15k	237	273k	18k	20k
NELL-995	75k	200	150k	543	4k
Kinship	104	25	9k	1k	1k
WN18	41k	18	141k	5k	5k
NELL-ONE	69k	358	190k	1k	2k

Table 1: Statistics of datasets. $|\mathcal{E}|$ and $|\mathcal{R}|$ represent the cardinalities of the entity and relation sets.

network (See, Liu, and Manning 2017) that facilitates copying words from the source text during new words generation, we propose a soft switch, that is, the fusion probability $p_f \in [0, 1]$, to adaptively adjust the fusion proportion between the joint vector and relational knowledge vector. The fusion probability p_f is estimated according to the joint vector as $p_f = \textit{sigmoid}(FC(\mathbf{j}))$, where FC is the fully connected neural network, and *sigmoid* is the activation function. Finally, we obtain the following fusion vector \mathbf{f} over the joint vector \mathbf{j} and relational knowledge vector \mathbf{rk} as

$$\mathbf{f} = (1 - p_f)\mathbf{j} + p_f \mathbf{rk}. \quad (5)$$

Classifier Module

Classification-aware Loss. The fusion vector \mathbf{f} is mapped to a class probability through the classifier block as

$$D \sim \textit{softmax}(\mathbf{f}^\top W_c), \quad (6)$$

where $W_c \in \mathbb{R}^{\text{dim} \times K}$ is the classification weight matrix, and *softmax* is the activation function.

Given the ground truth relation r_i from the i -th input (e_{h_i}, r_i, e_{t_i}) , we adopt cross entropy to assess the classification-aware loss as

$$\mathcal{L}_c = -\frac{1}{N} \sum_{i=1}^N \log p(r_i | (e_{h_i}, e_{t_i})), \quad (7)$$

where $p(r_i | (e_{h_i}, e_{t_i})) \in D_i$ is the probability of the ground truth relation r_i .

Most Similar Relation. Existing mainstream link prediction models have achieved impressive performance, yet they can only learn the patterns observed in the closed datasets, thereby limiting their scalability for handling the rapidly evolving KGs. Specifically, when a zero-shot relation r_z (i.e., one not existing in the training set) occurs between an entity pair (e_h, e_t) , it is almost impossible for the existing models to distinguish whether this new triple (e_h, r_z, e_t) is valid or not. All r_z will be then identified as an ‘unknown’ vector \mathbf{u} by the embedding component, and the newly constructed triple representation (e_h, \mathbf{u}, e_t) will receive a low score. To alleviate this defect, GRL selects the most semantically similar relation for replacing to enhance the learning ability of base model on zero-shot relations. We argue that the relation which corresponds to the maximum similarity in α_{sim} reflects the semantic relation of two entities in the best way. Therefore, we use the vector of the most similar relation \mathbf{r}_{ms} to replace the vector \mathbf{u} and evaluate the newly constructed triple representation $(e_h, \mathbf{r}_{ms}, e_t)$.

	YAGO3-10			FB15K-237			NELL-995			Kinship			WN18		
	MRR	HITS@N		MRR	HITS@N		MRR	HITS@N		MRR	HITS@N		MRR	HITS@N	
		@10	@1		@10	@1		@10	@1		@10	@1		@10	@1
ComplEx	36.0	55.0	26.0	24.7	42.8	15.8	48.2	60.6	39.9	82.3	97.1	73.3	94.1	94.7	93.6
R-GCN	-	-	-	24.8	41.7	15.3	12.0	18.8	8.2	10.9	23.9	3.0	81.4	96.4	69.7
ConvKB	-	-	-	28.9	47.1	19.8	43.0	54.5	37.0	61.4	95.3	43.6	-	-	-
D4-STE	47.2	64.3	38.1	32.0	50.2	23.0	-	-	-	-	-	-	94.6	95.2	94.2
D4-Gumbel	38.8	57.3	29.4	30.0	49.6	20.4	-	-	-	-	-	-	94.6	95.4	94.2
DistMult	34.0	54.0	24.0	24.1	41.9	15.5	48.5	61.0	40.1	51.6	86.7	36.7	82.2	93.6	72.8
+GRL	41.2	59.9	31.1	25.8	43.9	16.9	54.3	64.6	47.6	52.2	86.4	37.3	86.1	95.2	79.2
(± sd)	(0.3)	(1.0)	(0.1)	(0.2)	(0.3)	(0.1)	(0.2)	(0.3)	(0.3)	(0.2)	(0.8)	(0.2)	(1.0)	(0.4)	(1.1)
ConvE	52.0	66.0	45.0	31.6	49.1	23.9	49.1	61.3	40.3	83.3	98.1	73.8	94.2	95.5	93.5
+GRL	55.4	69.0	47.4	32.6	50.2	24.8	49.4	60.6	41.5	83.4	97.8	74.5	94.8	95.7	94.4
(± sd)	(1.0)	(0.1)	(0.1)	(0.3)	(0.2)	(0.3)	(0.2)	(0.3)	(0.3)	(0.2)	(0.5)	(0.5)	(0.1)	(0.4)	(0.0)

Table 2: Link prediction results (mean ± sd) of the compared models (%): the best results are marked in bold (pairwise t-test at 5% significance level).

	YAGO3-10			NELL-995		
	Many	Few	All	Many	Few	All
DistMult	38.1	26.7	34.0	52.6	41.9	48.5
DistMult+GRL	44.8	34.2	41.2	57.3	48.8	54.3
(Increment)	(↑6.7)	(↑7.5)	(↑7.2)	(↑4.7)	(↑6.9)	(↑5.8)
ConvE	57.9	20.0	52.4	52.0	42.2	49.1
ConvE+GRL	59.4	24.6	55.4	52.4	43.9	49.4
(Increment)	(↑1.5)	(↑4.6)	(↑3.0)	(↑0.4)	(↑1.7)	(↑0.3)

Table 3: Link prediction results with the increment (%) on many-shot and few-shot sub-groups, and entire test set.

Learning Scheme

We follow the definition of score-aware loss in existing base models and propose a classification-aware loss to train the model. The overall optimization follows the joint learning paradigm that is defined as a weighted combination of constituent losses as $\mathcal{L} = \mathcal{L}_s + \lambda \mathcal{L}_c$, where λ is a hyperparameter to balance the importance between the score-aware loss and classification-aware loss for optimization.

Experiments and Results

Datasets

We select two categories of datasets to comprehensively evaluate GRL as follows, whose statistical descriptions are shown in Table 1:

- *Imbalanced datasets*: YAGO3-10 (Mahdisoltani, Biega, and Suchanek 2015), FB15K-237 (Toutanova and Chen 2015), NELL-995 (Xiong, Hoang, and Wang 2017), Kinship (Lin, Socher, and Xiong 2018), and WN18 (Bordes et al. 2013). These datasets contain both many-shot and few-shot relations.
- *Few-shot dataset*: NELL-ONE (Xiong et al. 2018), which is specially constructed for the few-shot learning task in

KG. The relations with less than 500 but more than 50 training triples are selected as testing data.

Baselines

We adopt two embedding-based models, DistMult (Yang et al. 2015) and ConvE (Dettmers et al. 2018), as the base models of our proposed modules, and compare the two enhanced models with the following popular relation prediction models: RESCAL (Nickel, Tresp, and Krieger 2011), TransE (Bordes et al. 2013), DistMult (Yang et al. 2015), ComplEx (Trouillon et al. 2016), ConvE (Dettmers et al. 2018), ConvKB (Nguyen et al. 2018), D4-STE, D4-Bumbel (Xu and Li 2019), and TuckER (Balazevic, Allen, and Hospedales 2019). Besides the above general models, we test two additional models, GMatching (Xiong et al. 2018) and CogKR (Du et al. 2019), which are designed specifically for the few-shot relation learning.

Experimental Configuration

We implement the base models and our proposed two modules in PyTorch (Paszke et al. 2017). Throughout the experiments, we optimize the hyperparameters in a grid search setting for the best mean reciprocal rank (MRR) on the validation set. We use Adam to optimize all the parameters with initial learning rate at 0.003. The dimensions of entity and relation embeddings are both set to 200. The loss weight λ is set to 0.1. According to the frequency of relations, we take the top 20% and bottom 80% of relations as many-shot and few-shot relation classes, respectively. The experimental results of our model are averaged across three training repetitions, and standard deviations (sd) are also reported.

Experiment I: Link Prediction

Setting We follow the evaluation protocol of (Dettmers et al. 2018): each input (e_h, r, e_t) is converted to two queries, that is, tail query $(e_h, r, ?)$ and head query $(?, r, e_t)$; then, the ranks of correct entities are recorded among all entities for each query, excluding other correct entities that

	MRR	HITS@N		
		@10	@5	@1
TransE [†]	9.3	19.2	14.1	4.3
GMatching [†]	18.8	30.5	24.3	13.3
CogKR [*]	25.6	35.3	31.4	20.5
DistMult [†]	10.2	17.7	12.6	6.6
DistMult+GRL	14.4	23.0	18.2	9.8
(\pm sd)	(2.0)	(2.1)	(1.9)	(2.3)
ConvE [*]	17.0	30.6	23.0	10.5
ConvE+GRL	25.6	38.9	33.6	18.8
(\pm sd)	(2.3)	(3.7)	(3.1)	(2.1)

Table 4: Few-shot relation learning results (mean \pm sd) on NELL-ONE dataset (%): the results marked by ‘[†]’ or ‘^{*}’ are taken from (Xiong et al. 2018; Du et al. 2019).

were observed in any of the train/valid/test sets for the same query. We use the filtered HITS@1, 5, 10, and MRR as evaluation metrics.

Results Table 2 records the results on five imbalanced datasets, which reflect the general performance of the compared models in solving the link prediction task. It shows that two base models (DistMult and ConvE) are generally improved by incorporating the proposed GRL framework. That is, GRL improves DistMult by an average of 3.84% and improves ConvE by an average of 1.08% under the MRR evaluation. Especially, the enhanced model ConvE+GRL generally outperforms the compared models on YAGO3-10, FB15K-237, Kinship, and WN18, and the enhanced model DistMult+GRL also performs well on NELL-995. We also evaluate the performance of GRL in learning many-shot and few-shot relations and show the MRR results of DistMult, DistMult+GRL, ConvE, and ConvE+GRL on YAGO3-10 and NELL-995 (c.f. Table 3). The results indicate that GRL achieves consistent improvements on both “many-shot” and “few-shot” sub-groups. We assume this may be because handling many-shot relations can be improved thanks to useful implicit information from few-shot relations, even though there are already numerous training samples for many-shot relations. From this aspect, it is sensible for the mainstream link prediction models to rely on GRL regarding the imbalanced relation issue.

Experiment II: Few-shot Relation Learning

Setting To further evaluate the performance of GRL in the few-shot relation learning case, which is tricky for a link prediction model, especially, when relations are very insufficient, we test approaches on the NELL-ONE dataset wherein each test relation has only one instance in the training set. We follow the evaluation protocol and metrics of (Xiong et al. 2018).

Results Table 4 shows that GRL consistently improves both of the base models by average 4.2% and 8.6% MRR scores. Especially for ConvE, incorporating GRL helps it outperform the other approaches on three metrics. CogKR,

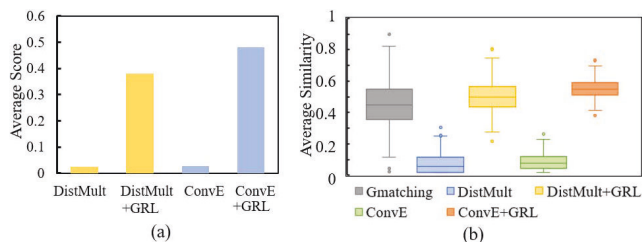


Figure 3: Zero-shot relation learning results: (a) the average score of the testing triples, and (b) the average similarity between the zero-shot relation with the predicted relation.

a path learning based model, performs best under HITS@1. The reason might be that the testing query is easy to be completed by finding KG paths on the few-shot relation datasets, such as NELL-ONE. Although there is only one training instance for each testing query, GRL can effectively embed the few-shot relations by learning from the semantically similar relations in the many-shot class.

Experiment III: Zero-shot Relation Learning

Setting To evaluate the performance on zero-shot relations of GRL, we construct a testing set containing 500 triples whose relations are unseen in the training phase. The testing triples are randomly sampled from the FB15K dataset (Bordes et al. 2013), and the training set is FB15K-237 to ensure the authenticity of the triples. We adopt the fundamental testing protocol that quantitatively determines the scores of triples with zero-shot relations.

Most of existing zero-shot relation studies have to depend on textual descriptions, while the zero-shot learning addressed in this work does not require this information. Therefore, we select the GMatching model (Xiong et al. 2018) for comparison, which can predict similar relations by learning a matching metric without any additional information. We use the classical method TransE (Bordes et al. 2013) to learn the relation embeddings in the FB15K dataset and calculate the similarity between the zero-shot relation and the predicted relation.

Results Figure 3 (a) demonstrates results of the average score of the testing triples with zero-shot relations. Note that we use the fusion vector as the zero-shot relations embedding. We can see that two base models (DistMult and ConvE) cannot get a good average score because all zero-shot relations will be identified as an ‘unknown’ relation. When GRL is plugged, two enhanced models (DistMult+GRL and ConvE+GRL) are both boosted in learning positive relations, proving that the GRL framework can effectively improve the validation capabilities on triples with zero-shot relations of the base models. Figure 3 (b) shows the performance on predicting zero-shot relations. We can see that the base models perform worse due to their superficial way of embedding zero-shot relations as mentioned before. When equipping with GRL, the enhanced models perform better than GMatching, indicating that learning from the semantic correlations between unseen relations and seen

		YAGO3-10	NELL-ONE
(1)	ConvE	52.0	17.0
(2)	ConvE+GRL($p_f = 0$)	52.6	23.3
(3)	ConvE+GRL($p_f = 0.5$)	53.9	24.7
(4)	ConvE+GRL($p_f = 1$)	52.2	20.3
(5)	ConvE+Direct	51.2	10.5
(6)	ConvE+GRL	55.4	25.6

Table 5: Ablation Study.

relations provides a comparably good way as learning from neighbor information.

Further Analysis of GRL

Ablation Study

Study of Fusion Probability To assess the effect of the fusion vector, we make a comparison on three variants from the fusion probability perspective based on ConvE, see Table 5 (2)-(4). The three variants are the followings: only using the joint vector (i.e., $p_f = 0$), only using the relational knowledge vector (i.e., $p_f = 1$), and using the joint and relational knowledge vectors with an equal weight (i.e., $p_f = 0.5$). Compared with three variants, fusing the joint and relational knowledge vectors (i.e., ConvE+GRL) performs best, which suggests the semantic correlations in the relational knowledge vectors can help the base model learn more appropriate representations of relations and thus boost the general performance. Moreover, the adaptive fusion probability can improve the flexibility of the fusion operator.

Direct Fusion vs. GRL We test now a direct fusion method that fuses the relational knowledge vector with the relation representation as the updated relation representation without considering the classification-aware loss. Table 5 (5) shows the MRR performance of ConvE when enhanced by the direct method. Rich semantic correlations in KGs cannot be adequately learned by the direct method because it simply leverage the superficial semantic correlations, rather than embedding them into relation vectors. Moreover, the direct method will make embedding learning more confusing especially for the few-shot relation data such as NELL-ONE.

Case Study

Visualization of Knowledge-aware Attention GRL is able to make the base model fully learn semantic correlations between relations. To verify this, we display the attention distribution for the base model (ConvE) and enhanced model (ConvE+GRL) on FB15K-237 in Figure 4, and show the average attention distribution of 237 relation classes where each row represents a type of relation. The base model learns little about semantic correlations between relations, while the enhanced model (ConvE+GRL) can perfectly capture the semantic correlations. The attention distribution of few-shot relations is more discrete than many-shot relations due to insufficient training data.

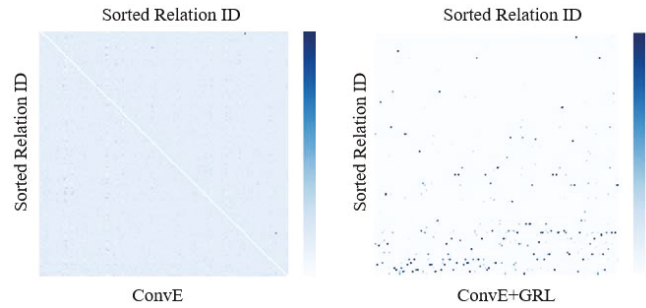


Figure 4: Case study: knowledge-aware attention cases with a heat map.

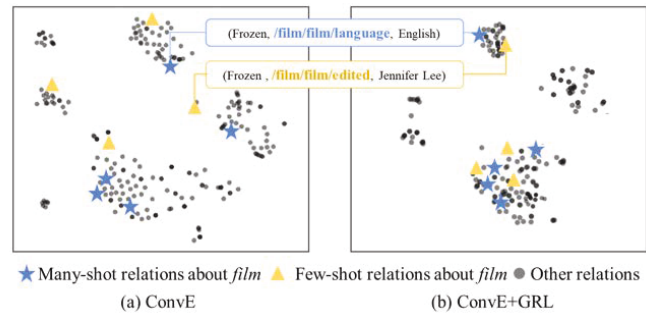


Figure 5: Case study: t-SNE visualization of relation embeddings in FB15K-237 (better view in color). The semantically similar relations get closer after plugging GRL.

Visualization of Relation Embedding In addition, we also show in Figure 5 the t-SNE (Maaten and Hinton 2008) plot of all relations on FB15K-237 in embedding space. To provide more insights we highlight the relations associated with “film”. The Stars and Triangles represent the many-shot and few-shot relations, respectively. We can see that the many-shot and few-shot relations are more compact in the case of the enhanced model than the base model .

Conclusion and Future Work

In this work, we study two natural problems in the link prediction task: 1) unbalanced relation distribution, and 2) unseen relations. To address them, we focus on generalized relation learning and propose a framework, GRL, that uses semantic correlations among relations as a bridge to connect semantically similar relations. Through extensive experiments on six datasets, we demonstrate the effectiveness of GRL, providing a comprehensive insight into the generalized relation learning of KGs. There are a few loose ends for further investigation. We will consider combining the external text information and the semantic knowledge of KGs to facilitate the relation learning. We will also try to deploy GRL to downstream applications that involve generalized relation learning scenarios to gain more insights.

Acknowledgments

This work was supported in part by the Ministry of Education of Humanities and Social Science project under grant 16YJC790123 and the Natural Science Foundation of Shandong Province under grant ZR2019MA049.

References

- Balazevic, I.; Allen, C.; and Hospedales, T. 2019. TuckER: Tensor Factorization for Knowledge Graph Completion. In *Proceedings of EMNLP-IJCNLP*, 5184–5193.
- Bordes, A.; Usunier, N.; Garcia-Durán, A.; Weston, J.; and Yakhnenko, O. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Proceedings of NeurIPS*, 2787–2795.
- Chen, M.; Zhang, W.; Zhang, W.; Chen, Q.; and Chen, H. 2019a. Meta Relational Learning for Few-Shot Link Prediction in Knowledge Graphs. In *Proceedings of EMNLP-IJCNLP*, 4216–4225.
- Chen, W.; Zhu, H.; Han, X.; Liu, Z.; and Sun, M. 2019b. Quantifying Similarity between Relations with Fact Distribution. In *Proceedings of ACL*, 2882–2894.
- Christmann, P.; Saha Roy, R.; Abujabal, A.; Singh, J.; and Weikum, G. 2019. Look before you Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion. In *Proceedings of CIKM*, 729–738.
- Dettmers, T.; Minervini, P.; Stenetorp, P.; and Riedel, S. 2018. Convolutional 2D Knowledge Graph Embeddings. In *Proceedings of AAAI*, 1811–1818.
- Du, Z.; Zhou, C.; Ding, M.; Yang, H.; and Tang, J. 2019. Cognitive Knowledge Graph Reasoning for One-shot Relational Learning. *arXiv preprint arXiv:1906.05489*.
- Ji, G.; Liu, K.; He, S.; and Zhao, J. 2016. Knowledge Graph Completion with Adaptive Sparse Transfer Matrix. In *Proceedings of AAAI*, 985–991.
- Karlinsky, L.; Shtok, J.; Harary, S.; Schwartz, E.; Aides, A.; Feris, R.; Giryes, R.; and Bronstein, A. M. 2019. Rep-Met: Representative-Based Metric Learning for Classification and Few-Shot Object Detection. In *Proceedings of CVPR*, 5197–5206.
- Lei, W.; He, X.; de Rijke, M.; and Chua, T.-S. 2020a. Conversational Recommendation: Formulation, Methods, and Evaluation. In *Proceedings of SIGIR*, 2425–2428.
- Lei, W.; Zhang, G.; He, X.; Miao, Y.; Wang, X.; Chen, L.; and Chua, T.-S. 2020b. Interactive Path Reasoning on Graph for Conversational Recommendation. In *Proceedings of KDD*, 2073–2083.
- Lin, X. V.; Socher, R.; and Xiong, C. 2018. Multi-Hop Knowledge Graph Reasoning with Reward Shaping. In *Proceedings of EMNLP*, 3243–3253.
- Lin, Y.; Liu, Z.; Sun, M.; Liu, Y.; and Zhu, X. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. In *Proceedings of AAAI*, 2181–2187.
- Liu, Z.; Miao, Z.; Zhan, X.; Wang, J.; Gong, B.; and Yu, S. X. 2019. Large-Scale Long-Tailed Recognition in an Open World. In *Proceedings of CVPR*, 2537–2546.
- Lv, X.; Gu, Y.; Han, X.; Hou, L.; Li, J.; and Liu, Z. 2019. Adapting Meta Knowledge Graph Information for Multi-Hop Reasoning over Few-Shot Relations. In *Proceedings of EMNLP-IJCNLP*, 3374–3379.
- Maaten, L. v. d.; and Hinton, G. 2008. Visualizing Data Using T-SNE. *Journal of Machine Learning Research* 2579–2605.
- Mahdisoltani, F.; Biega, J.; and Suchanek, F. 2015. YAGO3: A Knowledge Base from Multilingual Wikipedias. In *Proceedings of CIDR*.
- Moon, S.; Shah, P.; Kumar, A.; and Subba, R. 2019. OpenDialKG: Explainable Conversational Reasoning with Attention-based Walks over Knowledge Graphs. In *Proceedings of ACL*, 845–854.
- Nguyen, D. Q.; Nguyen, T. D.; Nguyen, D. Q.; and Phung, D. 2018. A Novel Embedding Model for Knowledge Base Completion Based on Convolutional Neural Network. In *Proceedings of NAACL*, 327–333.
- Nickel, M.; Tresp, V.; and Kriegel, H.-P. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data. In *Proceedings of ICML*, 809–816.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic Differentiation in PyTorch. In *Proceedings of NeurIPS Workshop*.
- Qin, P.; Wang, X.; Chen, W.; Zhang, C.; Xu, W.; and Wang, W. Y. 2020. Generative Adversarial Zero-Shot Relational Learning for Knowledge Graphs. In *Proceedings of AAAI*, 8673–8680.
- Schlichtkrull, M.; Kipf, T. N.; Bloem, P.; Van Den Berg, R.; Titov, I.; and Welling, M. 2018. Modeling Relational Data with Graph Convolutional Networks. In *Proceedings of ESWC*, 593–607.
- See, A.; Liu, P. J.; and Manning, C. D. 2017. Get To The Point: Summarization with Pointer-Generator Networks. In *Proceedings of ACL*, 1073–1083.
- Seyed, M. K.; and David, P. 2018. Simple Embedding for Link Prediction in Knowledge Graphs. In *Proceedings of NeurIPS*, 4284–4295.
- Shi, B.; and Weninger, T. 2018. Open-World Knowledge Graph Completion. In *Proceedings of AAAI*, 1957–1964.
- Shi, X.; Salewski, L.; Schiegg, M.; Akata, Z.; and Welling, M. 2019. Relational Generalized Few-Shot Learning. *arXiv preprint arXiv:1907.09557*.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical Networks for Few-shot Learning. In *Proceedings of NeurIPS*, 4077–4087.
- Toutanova, K.; and Chen, D. 2015. Observed Versus Latent Features for Knowledge Base and Text Inference. In *Proceedings of CVSC*, 57–66.

- Trouillon, T.; Welbl, J.; Riedel, S.; Gaussier, É.; and Bouchard, G. 2016. Complex Embeddings for Simple Link Prediction. In *Proceedings of ICML*, 2071–2080.
- Wang, X.; Huang, T.; Wang, D.; Yuan, Y.; and Chua, T. S. 2021. Learning Intents behind Interactions with Knowledge Graph for Recommendation .
- Wang, Z.; Lai, K.; Li, P.; Bing, L.; and Lam, W. 2019. Tackling Long-Tailed Relations and Uncommon Entities in Knowledge Graph Completion. In *Proceedings of EMNLP-IJCNLP*, 250–260.
- Wang, Z.; Zhang, J.; Feng, J.; and Chen, Z. 2014. Knowledge Graph Embedding by Translating on Hyperplanes. In *Proceedings of AAAI*, 1112–1119.
- Xiong, W.; Hoang, T.; and Wang, W. Y. 2017. DeepPath: A Reinforcement Learning Method for Knowledge Graph Reasoning. In *Proceedings of EMNLP*, 564–573.
- Xiong, W.; Yu, M.; Chang, S.; Guo, X.; and Wang, W. Y. 2018. One-Shot Relational Learning for Knowledge Graphs. In *Proceedings of EMNLP*, 1980–1990.
- Xu, C.; and Li, R. 2019. Relation Embedding with Dihedral Group in Knowledge Graph. In *Proceedings of ACL*, 263–272.
- Yang, B.; Yih, W.-t.; He, X.; Gao, J.; and Deng, L. 2015. Embedding Entities and Relations for Learning and Inference in Knowledge Bases. In *Proceedings of ICLR*.
- Ye, H.-J.; Hu, H.; Zhan, D.-C.; and Sha, F. 2019. Learning Classifier Synthesis for Generalized Few-Shot Learning. *arXiv preprint arXiv:1906.02944* .
- Zhang, N.; Deng, S.; Sun, Z.; Wang, G.; Chen, X.; Zhang, W.; and Chen, H. 2019. Long-tail Relation Extraction via Knowledge Graph Embeddings and Graph Convolution Networks. In *Proceedings of NAACL*, 3016–3025.
- Zhu, F.; Lei, W.; Wang, C.; Zheng, J.; Poria, S.; and Chua, T.-S. 2021. Retrieving and Reading: A Comprehensive Survey on Open-domain Question Answering.