

Tackling Long-Tailed Relations and Uncommon Entities in Knowledge Graph Completion *

Zihao Wang¹, Kwun Ping Lai¹, Piji Li^{1,2}, Lidong Bing³, Wai Lam¹,

¹The Chinese University of Hong Kong

²Tencent AI Lab, Shenzhen, China

³R&D Center Singapore, Machine Intelligence Technology, Alibaba DAMO Academy

zihaoawangbupt@gmail.com

{kplai, wlam}@se.cuhk.edu.hk

pijili@tencent.com

l.bing@alibaba-inc.com

Abstract

For large-scale knowledge graphs (KGs), recent research has been focusing on the large proportion of infrequent relations which have been ignored by previous studies. For example few-shot learning paradigm for relations has been investigated. In this work, we further advocate that handling uncommon entities is inevitable when dealing with infrequent relations. Therefore, we propose a meta-learning framework that aims at handling infrequent relations with few-shot learning and uncommon entities by using textual descriptions. We design a novel model to better extract key information from textual descriptions. Besides, we also develop a novel generative model in our framework to enhance the performance by generating extra triplets during the training stage. Experiments are conducted on two datasets from real-world KGs, and the results show that our framework ¹ outperforms previous methods when dealing with infrequent relations and their accompanying uncommon entities.

1 Introduction

Modern knowledge graphs (KGs)(Bollacker et al., 2008; Lehmann et al., 2015; Vrandečić and Krötzsch, 2014) consist of a large number of facts, where each fact is represented as a triplet consisting of two entities and a binary relation between them. KGs provide rich information and it has been widely adopted in different tasks, such as question answering (Yih et al., 2015), information extraction (Bing et al., 2017, 2015, 2016) and image classification (Marino et al., 2017). However,

The work described in this paper is substantially supported by a grant from the Research Grant Council of the Hong Kong Special Administrative Region, China (Project Code: 14204418).

¹The implementation of our framework can be found in <https://github.com/ZihaoWang/Few-shot-KGC>.

KGs still have the issue of incomplete facts. To deal with the problem, Knowledge Graph Completion (KGC) task is introduced to automatically deduce and fill the missing facts. There exist many previous works focusing on this task and embedding-based methods (Bordes et al., 2013; Wang et al., 2014; Trouillon et al., 2016) achieves the best performance among them. Recent works such as (Xiong et al., 2018) have pointed out that relations in KGs follow a long-tailed distribution. To be more precise, a large proportion of relations have only a few facts in KGs. However, previous works of KGC usually focused on small proportions of frequent relations and ignored the remaining ones. One observation is that they often conducted experiments on small datasets such as FB15k and WN18 (Bordes et al., 2013) where a relation typically possesses thousands of facts. Moreover, after analyzing real-world KGs, we find that the more infrequently a relation appears, the entities within its facts are also more uncommon. Figure 1 depicts the relationship between the relation frequency and the proportion of uncommon entities that appear in the facts of these relations in a KG, where an entity is treated as uncommon when it appears less or equal than 5 times in all triplets of the KG. From Figure 1, it is obvious that less frequent relations involve more uncommon entities than frequent relations. Therefore, when dealing with the problem of infrequent relations, the issue of uncommon entities should also be considered simultaneously, where they are two sides of a coin.

Previous works such as (Xiong et al., 2018) only focused on those infrequent relations and ignored the accompanying problem of uncommon entities. When handling uncommon entities, relying only on the structural information of KGs would lead to inferior performance due to data insufficiency,

and thus additional information is required. Some works (Toutanova et al., 2015; Xie et al., 2016) utilize textual description of entities, but they cannot extract different information from entity description if the entity is involved in more than one relations. A recent work (Shi and Weninger, 2018) tries to tackle this problem by using an attention mechanism considering both entity description and relation, but it adopts a heuristic method that cannot generalize well.

In this paper, we consider performing KGC for infrequent relations and uncommon entities as a few-shot learning problem, and we propose a framework that consists of three main components: description encoder, triplet generator, and meta-learner. In the description encoder, we design a novel structure to handle entities involved with multiple relations by automatically locating and extracting relation-specific information. We also simultaneously learn a triplet generator that is able to generate extra triplets in order to relieve the problem of data sparsity in few-shot learning. Moreover, a meta-learner is further adopted to learn a initial representation of the model that can be easily adapted to unseen relations and entities. As a result, our work has three main contributions as follows:

- We formulate the problem of infrequent relations and uncommon entities as a few-shot learning problem and propose a meta-learning framework to solve it.
- We propose a novel model to extract relation-specific information from entity description for entities with multiple relations.
- We propose a generative model that can enhance the performance of few-shot learning by generating extra triplets during the training stage.

2 Related Work

2.1 Knowledge Graph Completion

The knowledge graph completion (KGC) task focuses on automatically inferring and filling the missing facts in KG. The most successful method of KGC is the embedding-based method that learns a latent embedding in a common space for entities and relations. This method only relies on the structural information of entities and relations in KG. There exist a variety of methods (Bordes

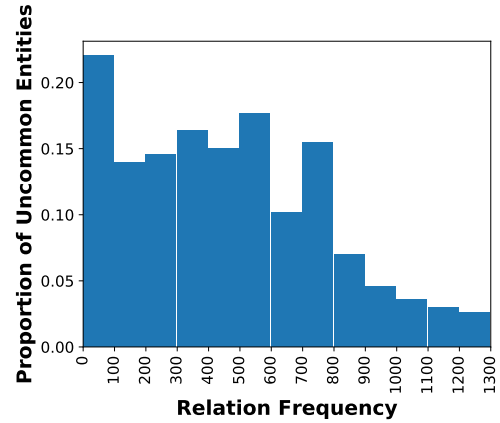


Figure 1: A histogram about relation frequencies and the corresponding proportions of uncommon entities in DBpedia.

et al., 2013; Socher et al., 2013; Wang et al., 2014; Trouillon et al., 2016; Nguyen et al., 2018) that have been proposed to learn good embeddings for entities and relations.

However, embedding of uncommon relation or entities can not learn a good representation due to the data insufficiency. Some research has proposed that additional information can be introduced to enhance the learning performance. Among different types of information, textual descriptions is commonly considered by previous works (Zhong et al., 2015; Toutanova et al., 2015; Xie et al., 2016; Shi and Weninger, 2018). Recently, meta-learning is also proposed by (Xiong et al., 2018) to learn infrequent long-tailed relations in KG.

2.2 Meta-Learning

Meta-learning (Lemke et al., 2015) aims at learning common experiences across different tasks and easy adapting the existing model to new tasks. One interesting application of meta-learning is few-shot learning problem where each task has only a few training data available.

Some research focus on learning a general policy for different tasks using a neural network. An early work (Santoro et al., 2016) proposes that the learning policy can be learned by using a global memory network. Recently, temporal convolution and attention have been considered to learn a common representation and pinpoint common experiences (Mishra et al., 2018).

Another direction is to learn a good initial representation where the learned model can be easily adapted to new data. Prototypical Network (Snell

et al., 2017) is proposed to learn a prototype for each category, and thus new data can be classified by distances between data and prototypes. Model-Agnostic Meta-Learning (MAML) (Finn et al., 2017) focuses on learning a good initial point in parameter space of model, hence a trained model can be quickly adapted to new tasks with several updates. More recently, Reptile (Nichol et al., 2018) proposes to be an approximation of MAML. In the Reptile model parameters are updated after a number of steps of inner iteration that can maximize within-task generalization.

3 Background and Overview

3.1 Problem Setting

Knowledge graph (KG) consists of a set of facts. Each fact has the form of a triplet (h, r, t) where h is a head entity, r is a relation and t is a tail entity. KGs are usually sparse, incomplete, and noisy. Therefore Knowledge Graph Completion (KGC) becomes an important task. Given arbitrary two out of three elements within a triplet, the goal of KGC is to predict the remaining one. We focus on predicting t given h and r in this work since our purpose is to deal with uncommon entities.

Previous works usually considered performing KGC given a set of common relations with lots of triplets. On the contrary, we concentrate on performing KGC on those relations that have only a small number of triplets, which can be viewed as a k -shot learning problem for relations when k is a small number. In the limiting case where k equals to 1, we deal with one-shot learning problem in our framework. Besides, unlike the previous work (Xiong et al., 2018) that focuses on common entities, we also consider uncommon entities during operational or testing phase, which means that some entities could appear only several times or be absent before. Moreover, when dealing with uncommon entities, relying only on the structural information of KG would lead to inferior performance due to data insufficiency, and thus additional information is necessary. Textual descriptions have been widely considered in KGC. Typically, textual descriptions are used to describe an entity or a relation, and each of them can be either a short sentence or a paragraph consisting of several sentences. We utilize textual descriptions of entities and relations in our framework.

3.2 Overview of Learning Method

Meta learning is a popular paradigm for solving the few-shot learning problem, and we adopt it to perform KGC. Given a KG, we treat each relation as a task, and the triplets of each relation can be viewed as specific data of each task. We further divide all tasks into three disjoint sets R_{train} , R_{val} and R_{test} . Hence meta-training, meta-validation and meta-testing can be performed on each set respectively. In each iteration of the meta-training phase, we randomly sample B tasks from R_{train} where B is batch size, and then for each task r in the batch we sample some triplets of r to train the model. After meta-training finishes, we obtain a trained model with model parameters W . Next, we follow the procedure in the previous work (Xiong et al., 2018) to perform meta-validation and meta-testing on R_{val} and R_{test} , where the settings are the same. So we only describe meta-testing for short. In the meta-testing phase, given a new task $r' \in R_{test}$ with $H_{r'}$ triplets, we randomly sample k out of $H_{r'}$ triplets. The trained model is further improved via another training stage with only these k samples and the parameters of model become W' . Then we keep parameters being fixed as W' and evaluate the performance of model on the remaining $H_{r'} - k$ triplets. These procedures are repeated for all tasks in R_{test} .

Given a triplet, the textual descriptions of h , r and t are respectively d_h , d_r and d_t . With textual descriptions, entities and relations can be mapped into a common semantic space. Therefore, uncommon entities can be tackled in this common space as usual.

4 Model Description

In this section we present the architecture and the learning procedure of our proposed framework.

First, given the textual descriptions of a triplet (h, r, t) , the description encoder extracts key information from descriptions and produces corresponding embeddings $O = (o_h, o_r, o_t)$, where o is a u -dimensional vector in the latent semantic space.

Next, the triplet generator participates in the learning procedure. During meta-training phase, it takes O as inputs and learns latent patterns for triplets. However, during the training stage of the meta-testing phase, instead of learning latent patterns, the triplet generator performs triplet aug-

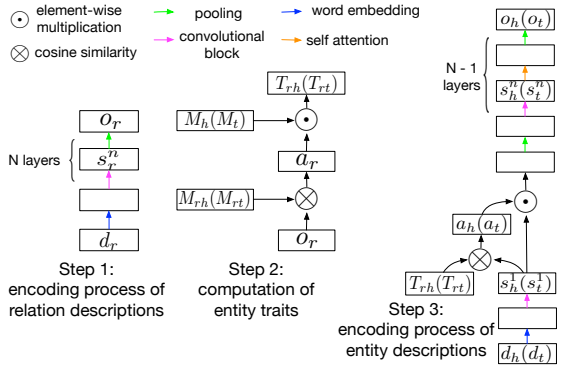


Figure 2: Structure and learning procedures of our description encoder. This figure should be viewed in color.

mentation by generating extra K sets of embeddings $G = \{(g_h, g_r, g_t)\}$. Each set of embeddings (g_h, g_r, g_t) can be viewed as an artificial imitation of O . In the few-shot setting, the size of O is usually too small to learn a good representation for a new task, and thus extra embeddings G is generated for data augmentation.

After previous procedures, we are able to obtain a set of embeddings of triplets $E = \{(e_h, e_r, e_t)\}$, where $E = O$ during meta-training phase and $E = O \cup G$ during training stage of meta-testing phase. With E prepared, a score function F takes E as input and computes the score C for each group of embeddings $(e_h, e_r, e_t) \in E$. Although more sophisticated score functions might be designed, in our framework we adopt a simple formula as follows:

$$C = F(h, r, t) = |e_h + e_r - e_t| \quad (1)$$

from TransE (Bordes et al., 2013), where L_1 -norm is used.

Finally, during the meta-training phase or the training stage of the meta-testing phase, a loss function L related to C is computed, and the meta-learner is adopted to optimize L so that the framework can be easily adapted to new relations and entities. Otherwise, during the testing stage of the meta-testing phase, we collect scores of the correct triplet and other candidates, and then we compute metrics based on the rank of correct triplet within all scores for evaluation.

4.1 Description Encoder

In KG, if an entity is involved in multiple relations, it is natural that different relations are more

relevant to different parts in the description of the entity. However, existing works using textual descriptions have not tackled this issue effectively. In order to deal with this issue, we define a new concept "entity trait" ("trait" for short) that represents the common characteristics of some entities related to a special relation. In another word, an entity owns different traits for different relations it involved. In a sense, a trait is similar to an entity type ("type" for short), but it has more advantages when handling KGC. First, types are not relation-specific but traits are. Besides, a trait may consist of semantics of several different types and hence it is more expressive. Moreover, we cannot easily obtain types in some situations, but traits can always be learned properly since they are latent and data-driven. Formally, we assume that a relation r has two traits T_{rh} and T_{rt} , where the previous one for all the head entities of r and the latter one for all the tail entities of r . In our description encoder, a simple but effective method is adopted to learn and utilize traits to extract relation-specific information from description.

The overall structure and learning process of our description encoder are given in Figure 2. Given the descriptions (d_h, d_r, d_t) of a triplet (h, r, t) , there are three steps to obtain the embeddings of triplet O as depicted in the figure. For entity, we only describe the process of h for simplicity, but everything stays the same for t . In Step 1, the encoding process of relation descriptions takes d_r as input and outputs a relation embedding o_r . Next, o_r is used to learn the trait T_{rh} for all the head entities of the relation r in Step 2. Finally, in Step 3 both d_h and T_{rh} are fed to the encoding process of the entity descriptions, and the output o_h is the embedding of the head entity. Note that the word embedding layer, convolutional blocks and pooling layers in Step 1 share the same parameters and architectures with the corresponding ones in Step 3.

The core of the encoding process is a N -layer convolutional neural network (CNN) (Conneau et al., 2017), which is shown to have excellent ability of extracting information. In our CNN, the basic convolutional block consists of three consecutive operations: two 1-d convolutions, an instance normalization (Ulyanov et al., 2016), and a non-linear mapping. For the pooling strategy, max pooling with a proper stride is used to distill the key information in the previous $N - 1$ layers, and

mean pooling is used to gather the information in the last layer. Moreover, in Step 3, we also apply self-attention mechanism (Vaswani et al., 2017) before each pooling layer in the last $N - 1$ layers. Unlike Step 1, self-attention is necessary here since entity descriptions are often more complex and noisier than relation descriptions according to our observation. Self-attention can assign lower weights to noise, and then those noise would be filtered out in the subsequent pooling layer.

Furthermore, we demonstrate how to compute the trait in Step 2, where the external memories M play an important role. These memories record the global information of relations and entities that can generalize well when encountering new relations and entities. In detail, the o_r computed in Step 1 is transformed to a probability distribution a_r by using m relation memories M_{rh}

$$a_r = \text{softmax}(o_r \otimes M_{rh}), \quad (2)$$

where a_r is m -dimensional, \otimes denotes the cosine similarity, softmax is the commonly used softmax function and M_{rh} is a matrix with shape (m, u) . After that, the u -dimensional trait T_{rh} can be obtained by computing a linear combination of m latent entity memories M_h

$$T_{rh} = \sum_{i=1}^m M_h^i \odot a_r^i, \quad (3)$$

where \odot is the element-wise multiplication between two vectors and M_h^i is a matrix with shape (m, u) . Note that each pair of m latent relation memories M_{rh} and m latent entity memories M_h has a one-to-one correspondence.

Finally, we describe how the trait T_{rh} is used to extract the relation-specific information in Step 3. Given the description d_h , the hidden states s_h^1 can be obtained after the first convolutional block, and then the trait T_{rh} is used to locate important hidden states in s_h^1 that have high relevance to r by assigning them higher weights. The procedure here is the same as before: a probability distribution a_h over s_h^1 is computed by

$$a_h = \text{softmax}(T_{rh} \otimes s_h^1), \quad (4)$$

and then a_h multiplies with s_h^1 element-wise to weight different hidden states. In this way, the hidden states that are not relevant to r are assigned lower weights, and thus they are more likely to be filtered in the subsequent max-pooling layer.

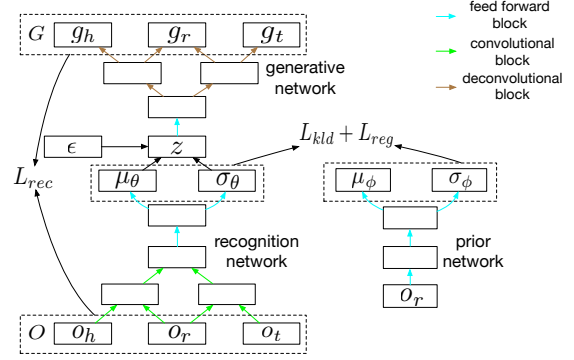


Figure 3: Structure of TCVAE. This figure should be viewed in color.

4.2 Triplet Generator

When handling KGC, learning good representation for infrequent relations and uncommon entities is difficult due to data sparsity. However, recently some research has focused on relieving the data sparsity in few-shot learning by generating extra data with a generative model (Schwartz et al., 2018; Wang et al., 2018). Inspired by these works, we propose a deep generative model that aims at triplet augmentation for k -shot learning. Although the generative adversarial network (Goodfellow et al., 2014) is a popular model that can generate high-quality samples (Frid-Adar et al., 2018), it suffers from an unstable learning process in our framework because of the difficult nature of Nash equilibrium and the influence of meta-learner. On the other hand, VAE is often applied to generate samples and extract latent semantics (Pu et al., 2016; Li et al., 2017) due to its smooth learning procedure. To cope with this issue, we design our triplet generator on the basis of CVAE (Sohn et al., 2015) and we name it triplet CVAE (TCVAE) in this paper. Figure 3 depicts the overall structure of TCVAE that is composed of three important probability distributions:

- Variational posterior distribution $q_\theta(z|O)$ parameterized by θ of the recognition network.
- Conditional prior distribution $p_\phi(z|o_r)$ parameterized by ϕ of the prior network.
- Likelihood distribution $p_\psi(G|z, o_r)$ parameterized by ψ of the generative network.

In the recognition network, there exist two layer of convolutional blocks. Each convolutional block takes two u -dimensional inputs and concatenates

them to form a matrix with shape $(2, u)$, so that 1-d convolution with filter width 2 can be applied to it. Instead of directly concatenate $O = (o_h, o_r, o_t)$ to form a matrix with shape $(3, u)$ and adopt only one layer of convolutional block, such a tree structure of two consecutive layers can better capture the pairwise semantics between any two embeddings in O . Likewise, two layers of the deconvolutional blocks that takes a u -dimensional vector as input and outputs a matrix with shape $(2, u)$ are placed in the generative network. Besides, all feed forward blocks in TCVAE consist of an affine transformation and a non-linear mapping.

During the meta-training phase, the recognition network takes O as input and learns the variational parameters μ_θ and σ_θ of the variational posterior $q_\theta(z|O)$, where the latent semantics z is assumed to follow a Gaussian distribution. Besides, the prior network conditioning on o_r computes the parameters μ_ϕ and σ_ϕ of the prior distribution $p_\phi(z|o_r)$. After that, the generative network samples three u -dimensional embeddings $G = (g_h, g_r, g_t)$ from the likelihood distribution $p_\psi(G|z, o_r)$. In the generative network, firstly, the latent variable z is transformed into a u -dimensional hidden state with the feed forward block after z . Next, the first deconvolutional block receives the hidden state and outputs a matrix with shape $(2, u)$. Finally, the second deconvolutional block receives the matrix before and outputs a matrix with shape $(3, u)$ which is denoted to G . G consists of three u -dimensional embeddings (g_h, g_r, g_t) corresponding to three elements (h, r, t) in the triplet.

Given the procedure, we are able to write down the loss terms of TCVAE: L_{rec} , L_{kld} and L_{reg} . More formally, L_{rec} is the expected log-likelihood that is also the reconstruction loss between the input O and the output G

$$L_{rec} = E_{q_\theta(z|O)} \log p_\psi(G|z, o_r). \quad (5)$$

L_{kld} is the KL-divergence between variational posterior distribution and conditional prior distribution

$$L_{kld} = KL(q_\theta(z|O) || p_\phi(z|o_r)). \quad (6)$$

And L_{reg} is the regularization term for the prior network proposed in (Ivanov et al., 2019)

$$L_{reg} = -\frac{\mu_\phi^2}{2\sigma_\mu^2} + \sigma_\sigma(\log \sigma_\psi - \sigma_\psi), \quad (7)$$

where $\sigma_\mu = 10000$ and $\sigma_\sigma = 0.0001$ are two hyper-parameters. These terms are jointly optimized with the loss function of KGC that we would demonstrate in the following subsection.

During the training stage of meta-testing phase, TCVAE uses prior network to compute μ_ϕ and σ_ϕ given only the relation embedding o'_r within O where $r' \in R_{test}$. Then it obtains K latent variables z with the following transformation

$$z = \mu_\phi + \sigma_\phi \odot \epsilon \quad \epsilon \sim \text{Gaussian}(0, 1). \quad (8)$$

After that, K embeddings of triplet G can be generated from the likelihood distribution $p_\psi(G|z, o'_r)$, and G is merged with O to form the final embeddings E . Please note that E is subsequently used to compute the score C with the score function F in Equation 1.

4.3 Loss Function and Meta-Learner

Following previous works, we adopt a simple strategy to compute the loss function of KGC L_{KGC} in both meta-training phase and the training stage of meta-testing phase. Given a randomly sampled relation r , first a positive triplet (h, r, t) is sampled from all triplets of r . Next, a negative triplet (h, r, t') can be produced by replacing t with another entity t' in KG, where the replacement is based on a uniform negative sampling. Note that if the negative triplet exists in KG, the negative sampling needs to be performed again. With a positive triplet and a negative triplet, embeddings of the positive triplet E^+ and embeddings of the negative triplet E^- can be obtained, and then two scores C^+ and C^- can be computed respectively. Finally, a hinge loss related to both scores is minimized for performing KGC

$$L_{KGC} = \max(0, \gamma + C^+ - C^-), \quad (9)$$

where γ is a margin hyper-parameter greater than 0. Moreover, during meta-training phase, L_{KGC} is also jointly optimized with TCVAE so that the overall loss L is

$$L = L_{KGC} - L_{rec} - \lambda_1 L_{kld} - \lambda_2 L_{reg}, \quad (10)$$

where the negative loss terms of TCVAE are minimized, and λ_1, λ_2 are two hyper-parameters for weighting terms in the overall loss.

In order to ensure that both description encoder and TCVAE have a good generalization ability when handling infrequent relations and uncommon entities, a meta-learner is further used

to optimize the overall loss L . Among different directions of meta-learning, we construct the meta-learner based on Reptile since we find it has the best performance in our task. In the context of KGC, learning with Reptile is different from previous KGC works. During the meta-training phase, Reptile searches for an initial point W within the parameter space of our framework, but such a framework may not perform well when directly used for performing KGC in R_{test} . Instead, during training stage of the meta-testing phase, the framework parameters can be quickly adapted to a new point W' that is suitable for performing KGC given a new relation $r' \in R_{test}$, and such an adaptation only needs a few training triplets of r' available. The procedure of learning with the meta-learner is depicted in Algorithm 1, where Adam (Kingma and Ba, 2015) is used during the S inner-training steps.

Algorithm 1 Learning Procedure of Meta-Learner

- 1: **for** iteration = 1, 2, ... **do**
 - 2: Save parameters of our framework W
 - 3: **for** $i = 1, \dots, B$ **do**
 - 4: Sample a relation $r \in R_{train}$, then sample a positive and a negative triplet of r
 - 5: Train our framework for S steps using Adam with learning rate α_1
 - 6: Save current parameters of our framework W_i and reset them to W
 - 7: **end for**
 - 8: Update $W \leftarrow W + \alpha_2(\frac{1}{B} \sum_{i=1}^B W_i - W)$
 - 9: **end for**
-

5 Experiments

5.1 Datasets

	WDtext	DBPtext
#Entity	60304	51768
#Relation	178	319
#Word	131796	170844
# R_{train}	130	220
# R_{val}	16	30
# R_{test}	32	69
Avg #words	5.3	170.8

Table 1: Statistics of datasets, where "Avg #words" means the average number of words in descriptions.

Existing datasets for KGC usually select triplets consisting of frequently appearing relations and

common entities. Recently two datasets focused on infrequent relations are proposed in (Xiong et al., 2018), but they do not contain textual descriptions for relations and entities. To obtain datasets that fulfill the practical few-shot learning situations as investigated in the problem setting, we manually harvest triplets and their textual descriptions from Wikidata (Vrandečić and Krötzsch, 2014) and DBpedia (Lehmann et al., 2015), and then we construct two datasets called WDtext and DBPtext respectively. The statistics of the two datasets are shown in Table 1. Specifically, the average numbers of words in the textual descriptions have a large variation between two datasets. In WDtext, descriptions are usually short phrases with only several words, on the other hand, descriptions in DBPtext consists of thousands of words, and we use first 200 words so that our model can be processed on GPUs. Such a variation of description length can better reveal the performance of our model in different situations. Besides, in order to collect enough triplets for evaluation, we select those relations whose numbers of triplets are greater than 5 and less than 1000, where the contained entities may exist only several times or even unseen during the meta-training phase. In this way, the problem setting of infrequent relations and uncommon entities is fulfilled in our datasets.

5.2 Experiment Setting

We compare our model with previous KGC models that can make use of textual descriptions, namely, DKRL (Xie et al., 2016) and ConMask (Shi and Weninger, 2018). We adapt their codes with our implementation. Following the experimental protocol of (Shi and Weninger, 2018), we also remove structural features of DKRL so that it can tackle unseen entities. To facilitate fair comparisons, even though these models are designed without using meta-learning, for each relation in R_{val} and R_{test} , we also sample k triplets and put them into R_{train} to ensure that all models make use of similar training data. Besides, the few-shot KGC model GMatching (Xiong et al., 2018) is also used for comparison. We only enable its "neighbor encoder" on WDtext because we cannot collect neighbor information from DBPedia. Moreover, we also design two additional baselines for ablation study. These baselines are constructed by removing specific components and keeping re-

Model	WDtext				DBPtext			
	Hits@10	Hits@5	Hits@1	MRR	Hits@10	Hits@5	Hits@1	MRR
DKRL	0.180	0.143	0.104	0.137	0.100	0.041	0.010	0.054
ConMask	0.279	0.207	0.085	0.156	0.304	0.213	0.059	0.147
GMatching	0.095	0.092	0.090	0.093	0.194	0.141	0.100	0.138
Ours-trait	0.179	0.168	0.081	0.126	0.182	0.124	0.060	0.107
Ours-TCVAE	0.178	0.155	0.138	0.149	0.343	0.258	0.110	0.187
Ours	0.198	0.189	0.148	0.168	0.376	0.321	0.224	0.281

Table 2: One-shot KGC results on WDtext and DBPedia, where **bold** numbers indicate best results over different models on the same metric.

Model	WDtext				DBPtext			
	Hits@10	Hits@5	Hits@1	MRR	Hits@10	Hits@5	Hits@1	MRR
DKRL	0.202	0.173	0.122	0.151	0.099	0.053	0.012	0.052
ConMask	0.303	0.211	0.107	0.173	0.333	0.225	0.077	0.163
GMatching	0.089	0.088	0.085	0.088	0.185	0.146	0.097	0.138
Ours-trait	0.277	0.221	0.113	0.170	0.252	0.174	0.088	0.146
Ours-TCVAE	0.214	0.198	0.159	0.184	0.303	0.227	0.110	0.174
Ours	0.258	0.227	0.180	0.210	0.409	0.319	0.186	0.255

Table 3: Four-shot KGC results on WDtext and DBPedia, where **bold** numbers indicate best results over different models on the same metric.

WDtext		DBPtext	
#Gen	MRR	#Gen	MRR
0	0.149	0	0.187
2	0.166	16	0.247
4	0.172	32	0.274
8	0.179	64	0.304
16	0.146	128	0.330
32	0.144	256	0.256

Table 4: MRR results of our framework when using different number of augmentation in TCVAE, where #Gen means the number of triplets being generated.

maintaining parts in our framework. Specifically, trait is removed in the baseline "Ours - trait" and TCVAE is removed in the baseline "Ours - TCVAE" respectively. The experimental setting of hyperparameters and initialization of our framework and baselines can be found in Appendix A.

To make a fair comparison, we use two categories of common metrics: mean reciprocal rank (MRR) and hits@P which is the percentage of correct tail entities ranked in the top P. Besides, experiments are conducted using four different random seeds and we report the average results of four trials. For each model, we select the epoch that has the best performance when evaluating on R_{val} and report the corresponding results on R_{test} .

5.3 Results of One-Shot Learning

Firstly, we compare our overall framework with baselines by conducting an one-shot KGC experiment, which is the most difficult case in few-shot learning. The results of our overall framework and baselines are shown in Table 2, where we find that our overall framework outperforms the baselines on most metrics. For WDtext, ConMask is a strong baseline and has a better result on Hits@10 and Hits@5, but it performs worse on Hits@1 and MRR compared to our framework. On the other hand, our overall framework outperforms for a large margin compared to all baselines on DBPtext. Since the main difference between the two datasets are the average length of descriptions, we can observe that our framework has a better performance when dealing with long textual descriptions. Besides, the results of the two ablation baselines are significantly worse than our overall framework, and thus we can see both components play an important role in our framework.

5.4 Results of Few-Shot Learning

In real-world KGs, few-shot learning of infrequent relations and uncommon entities are more common than one-shot scenario, so we also conduct four-shot KGC as another experiment. We use the same baselines and metrics as that in the one-shot

KGC experiment. The results are shown in Table 3, where we can see that our overall framework also has the best performance on different comparisons except for Hits@10 on WDtext. By comparing our overall framework with two ablation baselines, the importance of traits and triplet augmentation is demonstrated again. Furthermore, when compared with the previous one-shot KGC results, all baselines in this experiment are able to make use of the extra training data to improve their performances. Our framework also performs better on all metrics in WDtext, but it performs inferior on some metrics in DBPtext when compared with the one-shot scenario. One possible reason is descriptions in DBPtext is longer and more complex than that in WDtext, and thus four descriptions of training triplets are too diverse to learn a good representation.

5.5 Analysis of Triplet Generation

As shown in the previous KGC experiments, the performance of our framework heavily depends on the triplet generation provided by TCVAE. In this subsection, we further explore the effect of triplet generation by comparing the MRR result of our overall framework with different number of generated triplets in one-shot scenario, where the triplet augmentation is particularly helpful. For this analysis, we only conduct experiment and report the result with one trial for simplicity. The results are shown in Table 4, and we can conclude that a proper data augmentation does enhance the performance of our framework when training data available is scarce. Besides, the appropriate number of triplets being generated varies from one dataset to another, and too many or too few generation leads to an inferior performance. In WDtext, generating 8 extra triplets enhances the performance most, but generating 128 triplets is better in DBPtext. One reason for such a difference is that DBPtext has longer descriptions which also increase the variance of the generated triplets, and thus generating more triplets is necessary to learn a stable representation.

6 Conclusions

We consider a new type of KGC where infrequent relations and uncommon entities need to be jointly handled, and we formulate it as a few-shot KGC problem. To tackle the problem, we propose a novel concept "trait" and adopt it to extract

relation-specific information from entity descriptions. Besides, we also design a triplet generator and a meta-learning framework based on Reptile to deal with the issue of few-shot KGC. Moreover, we also conduct two new datasets that focus on this problem setting. The experiments of both one-shot and four-shot scenarios show that our framework has a better performance compared to other baselines.

References

- Lidong Bing, Sneha Chaudhari, Richard C. Wang, and William W. Cohen. 2015. [Improving distant supervision for information extraction using label propagation through lists](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 524–529.
- Lidong Bing, Bhuwan Dhingra, Kathryn Mazaitis, Jong Hyuk Park, and William W. Cohen. 2017. [Bootstrapping distantly supervised IE using joint learning and small well-structured corpora](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3408–3414.
- Lidong Bing, Mingyang Ling, Richard C. Wang, and William W. Cohen. 2016. [Distant IE by bootstrapping using lists and document structure](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, pages 2899–2905.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: A collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data, SIGMOD '08*, pages 1247–1250. ACM.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Durán, Jason Weston, and Oksana Yakhnenko. 2013. [Translating embeddings for modeling multi-relational data](#). In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, pages 2787–2795. Curran Associates Inc.
- Alexis Conneau, Holger Schwenk, Loïc Barrault, and Yann Lecun. 2017. [Very deep convolutional networks for text classification](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 1107–1116. Association for Computational Linguistics.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, pages 1126–1135.
- Maayan Frid-Adar, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan. 2018. Synthetic data augmentation using gan for improved liver lesion classification. *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 289–293.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. [Generative adversarial nets](#). In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, pages 2672–2680. MIT Press.
- Oleg Ivanov, Michael Figurnov, and Dmitry P. Vetrov. 2019. [Variational autoencoder with arbitrary conditioning](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. 2015. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic Web*, 6(2):167–195.
- Christiane Lemke, Marcin Budka, and Bogdan Gabrys. 2015. [Metalearning: a survey of trends and technologies](#). *Artificial Intelligence Review*, 44(1):117–130.
- Piji Li, Zihao Wang, Wai Lam, Zhaochun Ren, and Lidong Bing. 2017. [Salience estimation via variational auto-encoders for multi-document summarization](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3497–3503.
- K. Marino, R. Salakhutdinov, and A. Gupta. 2017. [The more you know: Using knowledge graphs for image classification](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20–28.
- Nikhil Mishra, Mostafa Rohaninejad, Xi Chen, and Pieter Abbeel. 2018. [A simple neural attentive meta-learner](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*.
- Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. 2018. [A novel embedding model for knowledge base completion based on convolutional neural network](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 327–333. Association for Computational Linguistics.
- Alex Nichol, Joshua Achiam, and John Schulman. 2018. [On first-order meta-learning algorithms](#). *CoRR*, abs/1803.02999.

- Yunchen Pu, Zhe Gan, Ricardo Henao, Xin Yuan, Chunyuan Li, Andrew Stevens, and Lawrence Carin. 2016. [Variational autoencoder for deep learning of images, labels and captions](#). In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pages 2352–2360.
- Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. 2016. [Meta-learning with memory-augmented neural networks](#). In *Proceedings of the 33rd International Conference on Machine Learning - Volume 48, ICML'16*, pages 1842–1850. JMLR.org.
- Eli Schwartz, Leonid Karlinsky, Joseph Shtok, Sivan Harary, Mattias Marder, Abhishek Kumar, Rogerio Feris, Raja Giryes, and Alex Bronstein. 2018. [Delta-encoder: an effective sample synthesis method for few-shot object recognition](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 2845–2855. Curran Associates, Inc.
- Baoxu Shi and Tim Weninger. 2018. Open-world knowledge graph completion. In *Proceedings of the Thirty-second AAAI Conference on Artificial Intelligence*, pages 1957–1964. AAAI Press.
- Jake Snell, Kevin Swersky, and Richard S. Zemel. 2017. [Prototypical networks for few-shot learning](#). In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 4077–4087.
- Richard Socher, Danqi Chen, Christopher D. Manning, and Andrew Y. Ng. 2013. [Reasoning with neural tensor networks for knowledge base completion](#). In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1, NIPS'13*, pages 926–934. Curran Associates Inc.
- Kihyuk Sohn, Honglak Lee, and Xinchun Yan. 2015. [Learning structured output representation using deep conditional generative models](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 3483–3491.
- Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoi-fung Poon, Pallavi Choudhury, and Michael Gamon. 2015. [Representing text for joint embedding of text and knowledge bases](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509. Association for Computational Linguistics.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, pages 2071–2080. JMLR.org.
- Dmitry Ulyanov, Andrea Vedaldi, and Victor S. Lempitsky. 2016. [Instance normalization: The missing ingredient for fast stylization](#). *CoRR*, abs/1607.08022.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: A free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.
- Yu-Xiong Wang, Ross B. Girshick, Martial Hebert, and Bharath Hariharan. 2018. Low-shot learning from imaginary data. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7278–7286.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. [Knowledge graph embedding by translating on hyperplanes](#). In *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, AAAI'14*, pages 1112–1119. AAAI Press.
- Ruobing Xie, Zhiyuan Liu, Jia Jia, Huanbo Luan, and Maosong Sun. 2016. Representation learning of knowledge graphs with entity descriptions. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, pages 2659–2665. AAAI Press.
- Wenhan Xiong, Mo Yu, Shiyu Chang, Xiaoxiao Guo, and William Yang Wang. 2018. [One-shot relational learning for knowledge graphs](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 1980–1990.
- Wen-tau Yih, Ming-Wei Chang, Xiaodong He, and Jianfeng Gao. 2015. [Semantic parsing via staged query graph generation: Question answering with knowledge base](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1321–1331.
- Huaping Zhong, Jianwen Zhang, Zhen Wang, Hai Wan, and Zheng Chen. 2015. [Aligning knowledge and text embeddings by entity descriptions](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 267–272.