



μ KG: A Library for Multi-source Knowledge Graph Embeddings and Applications

Xindi Luo¹, Zequn Sun¹ , and Wei Hu^{1,2}

¹ State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

xdluo.nju@gmail.com, zqsun.nju@gmail.com, whu@nju.edu.cn

² National Institute of Healthcare Data Science, Nanjing University, Nanjing, China

Abstract. This paper presents μ KG, an open-source Python library for representation learning over knowledge graphs. μ KG supports joint representation learning over multi-source knowledge graphs (and also a single knowledge graph), multiple deep learning libraries (PyTorch and TensorFlow2), multiple embedding tasks (link prediction, entity alignment, entity typing, and multi-source link prediction), and multiple parallel computing modes (multi-process and multi-GPU computing). It currently implements 26 popular knowledge graph embedding models and supports 16 benchmark datasets. μ KG provides advanced implementations of embedding techniques with simplified pipelines of different tasks. It also comes with high-quality documentation for ease of use. μ KG is more comprehensive than existing knowledge graph embedding libraries. It is useful for a thorough comparison and analysis of various embedding models and tasks. We show that the jointly learned embeddings can greatly help knowledge-powered downstream tasks, such as multi-hop knowledge graph question answering. We will stay abreast of the latest developments in the related fields and incorporate them into μ KG.

Resource Type: Software

License: GPL-3.0 License

GitHub Repository: <https://github.com/nju-websoft/muKG>

Keywords: Multi-source knowledge graphs · Representation learning · Link prediction · Entity alignment · Entity typing

1 Introduction

Knowledge graphs (KGs), such as Freebase [4], DBpedia [21], Wikidata [45], and YAGO [25], store rich structured knowledge about the real world. They have been widely used in a variety of knowledge-driven applications, including semantic search, question answering, and logic reasoning [19]. Learning vector

X. Luo and Z. Sun—Equal contributors.

© The Author(s), under exclusive license to Springer Nature Switzerland AG 2022
U. Sattler et al. (Eds.): ISWC 2022, LNCS 13489, pp. 610–627, 2022.
https://doi.org/10.1007/978-3-031-19433-7_35

representations (a.k.a. embeddings) of KGs has become critical to support these intelligent applications. In the past ten years, various KG embedding models such as TransE [5], ConvE [12], RotatE [41] and TuckER [3] were proposed and achieved promising performance. Please refer to the recent surveys [31, 46] for an overview. With these applications becoming more and more popular and diverse, they put forward higher demands to KGs in terms of coverage, richness and multilingualism. Oftentimes, a single KG cannot meet all these demands. This difficulty calls for the integration of multiple KGs. Learning from multi-source KGs with entity alignment has drawn a lot of attention in recent years [10, 36]. The joint KG embeddings have demonstrated useful for a variety of downstream tasks such as entity typing and multi-source KG completion [11, 35].

Table 1. Comparison of existing KG embedding libraries and ours.

Libraries	Multi-KG support	Deep learning libraries		KG tasks			
		PyTorch	TensorFlow	LP	EA	ET	Multi-LP
OpenKE [16] ^a	✗	✓	TF1	✓	✗	✗	✗
DGL-KE [57] ^b	✗	✓	✗	✓	✗	✗	✗
Pykg2vec [53] ^c	✗	✓	TF2	✓	✗	✗	✗
PyKEEN [1, 2] ^d	✗	✓	✗	✓	✗	✗	✗
TorchKGE [6] ^e	✗	✓	✗	✓	✗	✗	✗
LibKGE [7] ^f	✗	✓	✗	✓	✗	✗	✗
OpenEA [40] ^g	✓	✗	TF1	✗	✓	✗	✗
EAKit [54] ^h	✓	✓	✗	✗	✓	✗	✗
NeuralKG [56] ⁱ	✗	✓	✗	✓	✗	✗	✗
μ KG (Ours)	✓	✓	TF2	✓	✓	✓	✓

^a<https://github.com/thunlp/OpenKE>.

^b<https://github.com/awsmlabs/dgl-ke>.

^c<https://github.com/Sujit-O/pykg2vec>.

^d<https://github.com/pykeen/pykeen>.

^e<https://github.com/torchkge-team/torchkge>.

^f<https://github.com/uma-pil/kge>.

^g<https://github.com/nju-websoft/OpenEA>.

^h<https://github.com/THU-KEG/EAKit>.

ⁱ<https://github.com/zjukg/NeuralKG>.

To support the easy use of KG embeddings and foster reproducible research into KG embedding techniques, much effort has been dedicated to developing KG embedding libraries, including OpenKE [16], DGL-KE [57], Pykg2vec [53], PyKEEN [1, 2], TorchKGE [6], LibKGE [7], OpenEA [40], EAKit [54] and NeuralKG [56]. The majority of these libraries concentrates on the typical KG embedding task of link prediction. Only OpenEA and EAKit support multi-source KG embedding and the corresponding task entity alignment. Besides, most of them only support one deep learning library, especially PyTorch. No one supports another prominent deep learning library TensorFlow2 (TF2 for short). This limits the contexts in which these libraries can be used. Facing these limitations of

existing work and being aware of the effectiveness of multi-source KG embeddings, we develop a new scalable library, namely μ KG, for multi-source KG embeddings and applications. Table 1 compares μ KG with the existing popular KG embedding libraries. In summary, μ KG has the following features:

- **Comprehensive.** μ KG is a full-featured Python library for representation learning over a single KG or multi-source KGs. It is compatible with the two widely-used deep learning libraries PyTorch and TF2, and can therefore be easily integrated into downstream applications. It integrates a variety of KG embedding models and supports four KG tasks including link prediction, entity alignment, entity typing, and multi-source link prediction.
- **Fast and Scalable.** μ KG provides advanced implementations of KG embedding techniques with the support of multi-process and multi-GPU parallel computing, making it fast and scalable to large KGs.
- **Easy-to-Use.** μ KG provides simplified pipelines of KG embedding tasks for easy use. Users can interact with μ KG through both methods APIs and command line. It also has high-quality documentation.
- **Open-Source and Continuously Updated.** The source code of μ KG is publicly available. Our team will keep up-to-date on new related techniques and integrate new (multi-source) KG embedding models, tasks, and datasets into μ KG. We will also keep improving existing implementations.

Our experiments on several benchmark datasets demonstrate the effectiveness and efficiency of our library μ KG. Moreover, we carefully design two new tasks, multi-source link prediction and multi-source knowledge graph question answering (KGQA), with experiments to demonstrate the potential of multi-source KG embeddings:

- For **Multi-source Link Prediction**, we can convert the multiple KGs into a joint graph by merging their aligned entities, on which we learn joint KG embeddings for link prediction over each KG. This differs from the traditional link prediction, which first trains the model on a single KG and then predicts links for the same KG. In our joint learning setting, to avoid the test set of a KG’s link prediction task having overlap with other KGs’ training set, we do not consider relation alignment in multi-source KGs, and also remove these overlapping triples from the training set if they exist. Our experiment on DBP15K [36] shows that the joint trained TransE [5] outperforms its separately trained variant by 122% on *Hits@1*.
- For **multi-source KGQA**, as a downstream application of KGs, we have attempted to use multi-source KG embeddings to aid in the task of multi-hop question answering over a KG. The typical pipeline of using KG embeddings to answer natural language questions [32] is learning to align the question representation (encoded by a pre-trained language model like BERT [13]) with the answer entity’s embedding (encoded by a KG embedding model like ComplEx [44]). Conventional methods and datasets only consider QA over a single KG. We introduce an additional KG for joint embedding with the

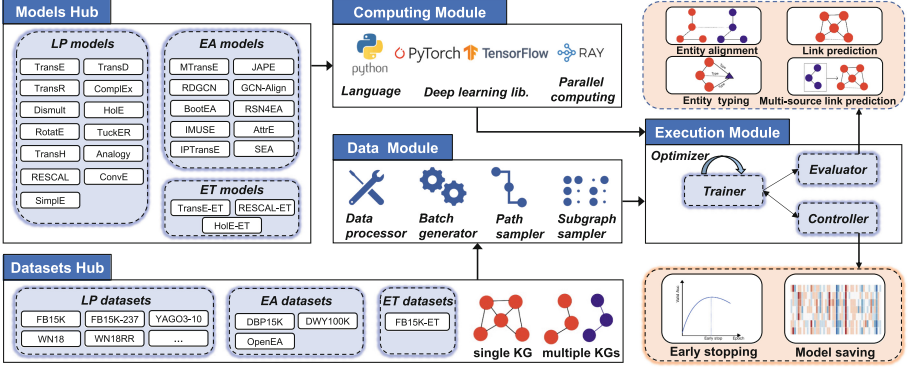


Fig. 1. Framework overview of μ KG.

target KG using μ KG. The embeddings of the target KG from the joint space are used for QA. Our results on WebQuestionsSP [52] show that joint KG embeddings can improve the accuracy by 8.6% over independently trained embeddings on a single KG.

Overall, these experiments show that the multi-source KG embeddings are able to promote knowledge fusion and transfer, and therefore benefit downstream tasks. We hope that our μ KG library can encourage the use of multi-source KG embeddings and promote their applications.

2 μ KG

μ KG is a scalable library for multi-source KG embeddings and applications. It also supports representation learning over a single KG. Its architecture is shown in Fig. 1. μ KG supports a variety of link prediction, entity alignment, and entity typing models, as well as the datasets that go with them. It consists of three modules. The *data module* converts the input single KG or multi-source KGs into the training data format (e.g., triples, paths or subgraphs) used by the embedding model. The *computing module* supports the execution module with neural computation and parallel training solutions. As a result, the *execution module* can be used for large-scale KGs and is compatible with the widely-used deep learning libraries PyTorch and TensorFlow. The execution module trains a KG embedding model with the training data produced by the data module. The *controller* keeps track of and records the training process. The *evaluator* employs the pre-trained embedding model to perform KG tasks, such as link prediction, entity alignment, entity typing, and multi-source link prediction.

2.1 Data Module

We hereby introduce the data module of μ KG.

Data Processor. The goal of the *data processor* is to generate numerical IDs for entities, relations, and attributes from the input single KG or multi-source KGs in the Datasets Hub. The numerical ID is the identifier of a resource in an embedding model. The data processor first reads the original relation triples and attribute triples from the `txt` or `ttl` files. Then, it assigns an ID to each entity, relation, and attribute. It currently provides two ID generation algorithms. The *unique-ID algorithm* generates a unique ID for each resource in KGs. It can be used for both single KG and multi-source KGs. The *shared-ID algorithm* generates the same ID for aligned entities in different KGs. In this way, the multiple KGs are merged as a “single” joint graph.

Batch Generator. The *batch generator* takes as input KG triples and divides the complete data into multiple fixed-size batches for model training. If the model requires relational paths or subgraphs for training, the batch generator would first call the path or subgraph sampler to convert triples. The batch generator also includes several negative sampling methods to randomly generate negative examples (e.g., negative alignment pairs or negative triples) for each positive example. The positive and negative examples are used in the embedding model for contrastive embedding learning. The *uniform negative sampling* method replaces an entity in a triple or an alignment pair with another randomly-sampled entity to generate a negative example. It gives each entity the same replacement probability. Such uniform negative sampling has the problem of inefficiency since many sampled negative samples are obviously false as training goes on, which does not provide any meaningful information. μ KG also supplies the self-adversarial negative sampling method [41] and the truncated negative sampling method [37] that seek to generate hard negative examples.

Path Sampler. The *path sampler* is to support some embedding models that are built by modeling the paths of KGs, such as IPTransE [58] and RSN [15]. It can generate three types of paths based on random walks. The first is the relational path like $(e_1, r_1, e_2, r_2, e_3)$, where e_i stands for an entity and r_j denotes a relation. It is an entity-relation chain. The second is the entity path like (e_1, e_2, e_3) , and the third is the relation path like (r_1, r_2) .

Subgraph Sampler. The *subgraph sampler* is to support GNN-based embedding models like GCN-Align [49] and AliNet [39]. It can generate both first-order (i.e., one-hop) and high-order (i.e., multi-hop) neighborhood subgraphs of entities. The GNN-based models represent an entity by aggregating the embeddings of its neighbors in the subgraphs.

2.2 Execution Module

This module carries out the training task of embedding models.

Trainer. The *trainer* directs the model training and evaluation based on the detailed configurations of users. It manages the model’s training progress. μ KG configures trainers for entity alignment models, link prediction models, and entity typing models, respectively. The *trainer* provides three optimizers, including the standard stochastic gradient descent, Adagrad, and Adam. It implements four loss functions, including the mean-squared loss, marginal ranking loss, limit-based loss, and noise-contrastive estimation loss.

Evaluator. The *evaluator* is to assess the performance of the trained model on specific test data. For (joint) link prediction, it uses the energy function to compute the plausibility of a candidate triple. For entity alignment or typing, it provides several metrics to measure entity embedding similarities, such as the cosine, inner, Euclidean distance, and cross-domain similarity local scaling. The evaluation process can be accelerated using multi-processing. The implemented metrics for assessing the performance of embedding tasks include *Hits@K*, mean rank (*MR*) and mean reciprocal rank (*MRR*). *Hits@K* measures the percentage of the test cases in which the correct counterpart is ranked in the top *k*. *MR* calculates the mean of these ranks. *MRR* is the average of the reciprocal ranks of results. Higher *Hits@K* and *MRR* or lower *MR* values indicate better performance.

Controller. The *controller* is in charge of the trainer. During the training process, the controller calls the evaluator to assess the model performance on validation data. If the performance begins to drop continuously, the controller would terminate the training (i.e., early stopping). After that, the controller saves the model and embeddings for further use.

2.3 Computing Module

In this section, we introduce the computing module.

Support of PyTorch and TF2. The computing module uses PyTorch and TF2 as the backbone for neural computing. Users can choose one of the backbones to run μ KG or carry on secondary development based on μ KG. If no backbone is specified by the user, μ KG can automatically detect which backbone has already been installed in the Python environment.

Multi-GPU and Multi-processing Computation. Scalability is a key consideration when we develop μ KG, because KGs in real-world applications are typically very large. Although PyTorch and TensorFlow both provide interfaces for parallel computing, they differ greatly in implementation and are difficult for users to use. Hence, we use Ray¹ to provide a uniform and easy-to-use interface

¹ <https://www.ray.io/>.

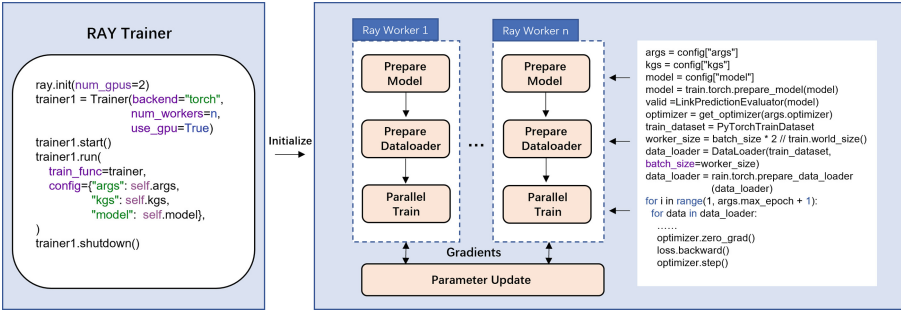


Fig. 2. Code snippet for training KG embedding models in the parallel mode.

```
PS D:\Knober\src\py> python main_args.py -t lp -m transe -o train -d data/FB15K
PS D:\Knober\src\py> python main_args.py -t lp -m transe -o train -d data/FB15K -r gpu:2 -w 2
```

Fig. 3. Command line for using μ KG.

for multi-GPU and multi-processing computation. Fig. 2 shows our Ray-based implementation for parallel computing and the code snippet to use it. Users can set the number of CPUs or GPUs used for model training.

2.4 User Interface

μ KG gives users two options for running KG embedding models. For users that are unfamiliar with μ KG, they can run a model on a dataset with the command line, as shown in Fig. 3. For advanced users, they can modify the configurations of a model and call the model’s running function in their Python code.

3 Experiments

In this section, we report our experiments to evaluate the effectiveness and efficiency of μ KG. The source code is available at our GitHub repository.²

3.1 Experiments on Effectiveness

To evaluate the effectiveness, we compare the results produced by our library with the corresponding official results reported in the models’ papers. We consider one single-KG task *link prediction*, and three multi-KG tasks *entity alignment*, *entity typing* and *multi-source link prediction*.

² <https://github.com/nju-websoft/muKG>.

Table 2. LP results on FB15K.

Models		<i>Hits@1</i>	<i>Hits@10</i>	<i>MRR</i>
RESCAL	Original	—	0.284	—
	Ours	0.129	0.342	0.202
TransE	Original	—	0.471	—
	Ours	0.194	0.647	0.353
TransH	Original	—	0.585	—
	Ours	0.188	0.604	0.332
TransD	Original	—	0.742	—
	Ours	0.214	0.595	0.345

Table 3. LP results on FB15K-237.

Models		<i>Hits@1</i>	<i>Hits@10</i>	<i>MRR</i>
TransE	Original	—	0.465	0.294
	Ours	0.174	0.463	0.270
ConvE	Original	0.237	0.501	0.325
	Ours	0.237	0.514	0.327
RotatE	Original	0.205	0.480	0.297
	Ours	0.172	0.456	0.260
TuckER	Original	0.266	0.544	0.358
	Ours	0.254	0.535	0.346

Link Prediction. We choose two benchmark datasets, FB15K [5] and FB15K-237 [42], for link prediction evaluation. FB15K-237 was created from FB15K to ensure that the testing and evaluation datasets do not have inverse relation test leakage. Recent link prediction models use FB15K-237 for evaluation. On FB15K, we compare the *Hits@1*, *Hits@10* and *MRR* results of four old but popular models in Table 2, including RESCAL [29], TransE [5], TransH [48], and TransD [18]. “—” denotes the unreported results. We can see that our implemented RESCAL, TransE and TransH can achieve better results than the original code due to our modern implementations. We also notice that the implemented TransD shows lower *Hits@10* performance than its original version. The reason lies in the different evaluation settings. The original TransD removes the corrupted triplets in the training, validation and test sets before ranking. But our implementation only removes those in the training set following other methods because this is more reasonable. On FB15K-237, we compare TransE and other three recent models including ConvE [12], RotatE [41],³ and TuckER [3] in Table 3. The results of TransE on FB15K-237 are taken from [41] because TransE was not evaluated on this dataset. As we can see, our implementations of TransE and ConvE in μ KG perform very similarly to their original code. As for RotatE and TuckER, the performance of our implementations is slightly lower than the original results, but also in the range of acceptance. This is due to different hyperparameter settings. In consideration of GPU resources, we do not set the embedding dimension to 1,000, which is used in their original papers but would cost too much GPU memory. Generally, a large dimension leads to good performance. In summary, our implementations of link prediction models can basically reproduce the reported results.

Entity Alignment. We use the recent benchmark dataset OpenEA [40] for entity alignment evaluation. OpenEA also provides the implementations of several entity alignment models using TensorFlow 1.12. We choose three structure-based entity alignment models GCN-Align [49], SEA [30] and BootEA [37], as

³ We use uniform negative sampling for a fair comparison with other models.

Table 4. Entity alignment results on EN-DE and EN-FR 15K.

Models	Backends	EN-DE			EN-FR		
		<i>Hits@1</i>	<i>Hits@10</i>	<i>MRR</i>	<i>Hits@1</i>	<i>Hits@10</i>	<i>MRR</i>
GCN-Align	OpenEA	0.481	0.753	0.571	0.338	0.680	0.451
	TF2 (ours)	0.480	0.754	0.571	0.335	0.670	0.446
	PyTorch (ours)	0.460	0.747	0.560	0.337	0.671	0.453
SEA	OpenEA	0.530	0.796	0.617	0.280	0.642	0.328
	TF2 (ours)	0.536	0.806	0.624	0.281	0.630	0.304
	PyTorch (ours)	0.561	0.834	0.650	0.321	0.679	0.439
BootEA	OpenEA	0.675	0.865	0.740	0.507	0.794	0.603
	TF2 (ours)	0.671	0.866	0.737	0.503	0.786	0.597
	PyTorch (ours)	0.662	0.884	0.738	0.493	0.811	0.599
IMUSE	OpenEA	0.580	0.778	0.647	0.569	0.777	0.638
	TF2 (ours)	0.567	0.672	0.636	0.571	0.777	0.640
	PyTorch (ours)	0.596	0.804	0.670	0.564	0.776	0.640

well as an attribute-enhanced model IMUSE [17], as baselines. We compare their *Hits@1*, *Hits@10* and *MRR* results on OpenEA’s EN-DE and EN-FR 15K settings with our PyTorch-based implementations and TF2-based implementations in Table 4. We can see that the two implementations of a model in μ KG achieve similar performance. For SEA and IMUSE, PyTorch-based implementations perform better than TF2-based implementations. We think this is caused by the difference between the two backbones. When compared to the results of OpenEA, μ KG achieves comparable results. This demonstrates the efficacy of our implementations for entity alignment models.

Entity Typing. Entity typing can be seen as a special link prediction task across an instance KG and an ontological KG. For example, given (“Michael Jackson”, “rdf:type”, $_$), the task is to predict the target type “/music/artist”. We use the FB15K-ET dataset for evaluation [26]. FB15K-ET is an expansion of FB15K with entity types. We follow [26] to implement two baselines, RESCAL-ET and HolE-ET, for entity typing. The two models are built based on the link prediction models RESCAL [29] and HolE [28], respectively. We compare our results with those in [26] in Table 5. We can observe that our implementations achieve similar or even better performance than those in [26], demonstrating the effectiveness of μ KG in entity typing.

Multi-source Link Prediction. This is a new task that we propose, which is inspired by both link prediction in a single KG and entity alignment between two KGs. We believe that training embeddings solely on a KG for link prediction is ineffective because the KG may be very incomplete. We introduce another background KG with entity alignment to the target KG for joint KG embedding learning. We use the shared-ID generation method in μ KG to merge the two

Table 5. Entity typing results on FB15K-ET.

Models		<i>Hits@1</i>	<i>Hits@10</i>	<i>MRR</i>
RESCAL-ET	Original	0.097	0.376	0.190
	Ours	0.128	0.456	0.236
HolE-ET	Original	0.133	0.382	0.220
	Ours	0.129	0.522	0.252

KGs and learn embeddings of the joint KG with a KG embedding model such as TransE [5]. When the learning progress is completed, only the embeddings of the target KG are used to participate in link prediction. For evaluation, we choose DBP15K_{ZH-EN} [36]. It is an entity alignment dataset, and we denote the two KGs in DBP_{ZH-EN} by DBP_{ZH} and DBP_{EN}, respectively. Following TransE [5], we divide triples into training, validation and test sets. Specifically, DBP_{ZH} has 63,372 training triples, 3,522 validation triples and 3,520 test triples, while DBP_{EN} has 85,627, 4,758 and 4,757, respectively. Conventional link prediction is usually carried out on a single KG. However, for multi-source link prediction with entity alignment, it would be interesting to see the performance of link prediction based on the jointly-trained KG embeddings. Based on μ KG, we train a TransE model over the joint graph of DBP_{ZH} and DBP_{EN}. We choose three translational models TransE [5], TransH [48] and TransD [18]; four semantic matching models DistMult [51], HolE [28], ComplEx [44] and Analogy [23]; as well as two neural models ProjE [34] and ConvE [12], as baselines. From Table 6, we can see that μ KG (TransE) outperforms the translational and semantic matching models. ConvE achieves better results than our method, but its model complexity is also much higher than TransE. By encoding alignment information, μ KG (TransE) greatly outperforms TransE. The results demonstrate the joint training is effective to improve the separately-trained models on link prediction. We think that this is because the alignment information between two KGs can complement the incomplete relational structures of each other.

3.2 Experiments on Efficiency

In this section, we evaluate the efficiency of the proposed library μ KG. The experiments were conducted on a server with an Intel Xeon Gold 6240 2.6GHz CPU, 512GB of memory and four NVIDIA Tesla V100 GPUs.

Efficiency of Multi-GPU Training. Figure 4 compares the training time of RotatE and ConvE on FB15K-237 when using different numbers of GPUs. As we can see, using multiple GPUs for parallel computing can significantly accelerate training. The final link prediction results are not affected by parallel computing.

Table 6. Link prediction results with joint KG embeddings.

Models	DBP _{ZH}			DBP _{EN}		
	<i>Hits</i> @1	<i>Hits</i> @10	<i>MRR</i>	<i>Hits</i> @1	<i>Hits</i> @10	<i>MRR</i>
TransE	0.100	0.529	0.248	0.099	0.512	0.241
TransH	0.103	0.519	0.274	0.125	0.535	0.263
TransD	0.097	0.506	0.237	0.114	0.517	0.251
DistMult	0.095	0.375	0.188	0.100	0.385	0.195
HolE	0.114	0.327	0.186	0.122	0.405	0.221
ComplEx	0.174	0.374	0.245	0.195	0.435	0.279
Analogy	0.145	0.363	0.220	0.169	0.375	0.241
ProjE	0.257	0.613	0.317	0.265	0.629	0.323
ConvE	0.291	0.597	0.398	0.322	0.631	0.429
μ KG (TransE)	0.222	0.549	0.331	0.252	0.585	0.363

For example, the *Hits*@1 scores of ConvE when using 1, 2, and 4 GPUs for computing are 0.241, 0.239 and 0.227, respectively. This experiment shows the efficiency of our multi-GPU training.




Fig. 4. Running time comparison on FB15K-237 with multi-GPU training.

Efficiency Comparison Against LibKGE and PyKEEN. We further compare the training time used by μ KG with LibKGE [7] and PyKEEN [1]. They are both PyTorch-based libraries for efficient training, evaluation, and optimization of KG embeddings. The backbone of μ KG in this experiment is also PyTorch. Table 7 gives the training time of ConvE and RotatE on FB15K-237 with a single GPU for calculation. For a fair comparison, we use the same hyper-parameter settings (e.g., batch size and maximum training epochs) for each model in the three libraries. We discover that μ KG costs less time than LibKGE and PyKEEN to train a KG embedding model, which demonstrates its efficiency.

Table 7. Running time on FB15K-237 with a GPU.

Models	μ KG	LibKGE	PyKEEN
RotatE	639 s	3,260 s	1,085 s
ConvE	824 s	1,801 s	961 s

4 Application to Multi-hop KGQA

We hereby report the experimental results on the downstream task, i.e., multi-hop KGQA, using our proposed joint embeddings.

Settings. We follow EmbedKGQA [32], a recent popular embedding-based KGQA method, to build a QA pipeline with our μ KG. EmbedKGQA consists of three modules. The KG embedding module learns embeddings for the input KG. Existing KG embedding models such as TransE [5] and ComplEx [44] can be chosen. The question embedding module encodes natural language questions with the help of the pre-trained language model RoBERTa [24], which is a new training recipe that improves on BERT and is widely used for encoding natural language text. The answer selection module chooses the final answer based on the question and relation similarity scores. Using KG embeddings to answer natural language questions can make it more effective in handling the relational sparsity in KGs. The KG embedding model used in EmbedKGQA is ComplEx. In our pipeline for QA, we use μ KG (the embedding model is also ComplEx) to learn joint embeddings based on the target KG and another background KG Wikidata5M [47], which is a subset of Wikidata with million-scale entities. For a fair comparison, we keep other modules in our pipeline the same as those in EmbedKGQA.

Dataset. We choose the popular multi-hop QA benchmark WebQuestionsSP [52] as the dataset. There are 4,737 questions in total. This dataset contains 1-hop and 2-hop questions that may be answered using Freebase entities [4]. Following EmbedKGQA, we limit the KG to a subset of Freebase that contains all relational triples within 2-hops of any entity specified in the WebQuestionsSP questions. We refine it further to include only those relations that are stated in the dataset. There are a total of 1.8 million entities and 5.7 million triples in this selected KG (denoted as FB4QA in this paper) to support these questions. The number of entity links between Wikidata5M and FB4QA is 493,987.

Results. Table 8 presents the QA accuracy. To study the effect of KG sparsity on QA performance, following EmbedKGQA, the FB4QA is used for two settings: Half-FB4QA and Full-FB4QA. The former randomly drops half of the triples in FB4QA to simulate an incomplete KG. The latter uses the full FB4QA to learn entity embeddings. Besides, in the Full-FB4QA w/ rel. pruning setting,

Table 8. QA accuracy on WebQuestionsSP.

	EmbedKGQA [32]	EmbedKGQA + Wikidata5M
Half-FB4QA	0.485	0.547
Full-FB4QA	0.587	0.646
Full-FB4QA w/rel. pruning	0.666	0.723

a relation pruning strategy is employed to reduce the candidate answer space by filtering out the dissimilar relations with the key entity in the question. We can see from the table that EmbedKGQA + Wikidata outperforms the baseline EmbedKGQA in all three settings. This is because our learned embeddings of FB4QA can benefit from the background KG, and thus are more expressive than those in EmbedKGQA. Both our method and EmbedKGQA in the Full-KG setting achieve better accuracy than the corresponding result in the Half-KG setting. This demonstrates that KG incompleteness degrades the quality of KG embeddings, and thus causes a decrease in performance. Our method can improve the incompleteness issue in KGs through knowledge transfer from other background KGs. We also consider extending LibKGE [7] for this new task. We merge two KGs into a large graph and use LibKGE to learn KG embeddings. The accuracy is 0.718 in the setting of Full-FB4QA w/ rel. pruning, a similar performance compared with our μ KG. This result further shows the effectiveness of our library and the potential of multi-source KG embeddings. In summary, this experiment demonstrates that multi-source KG embeddings are also effective in improving KG-related downstream tasks, and knowledge transfer between multi-source KGs is an alternative for boosting performance in real-world applications.

5 Related Work

In this section, we review the related work on KG embedding models and tasks, as well as existing libraries for KG embedding.

5.1 Knowledge Graph Embedding Tasks and Models

Link Prediction. TransE [5] introduces translational KG embeddings. It defines the score function $f_{\text{TransE}}(\tau) = \|\mathbf{h} + \mathbf{r} - \mathbf{t}\|^4$ to measure the plausibility of relational triple $\tau = (h, r, t)$, where h , r and t denote the head entity, relation and tail entity, respectively. Boldfaced letters denote the corresponding vector representations. Although TransE performs well for modeling one-to-one relations, it encounters issues when dealing with more complex relations. For example, if (h, r, t_1) and (h, r, t_2) hold for a one-to-many relation r , we have $\mathbf{h} + \mathbf{r} \approx \mathbf{t}_1$ and $\mathbf{h} + \mathbf{r} \approx \mathbf{t}_2$, then $\mathbf{t}_1 \approx \mathbf{t}_2$. If (h, r_1, t) and (h, r_2, t) hold for h and t , we have $\mathbf{r}_1 \approx \mathbf{r}_2$. To resolve these problems, several improved translational

⁴ Hereafter, $\|\cdot\|$ denotes the L_2 vector norm.

models, such as TransH [48], TransR [22], and TransD [18], have been proposed. They enable entities to have relation-specific embeddings. For example, TransH interprets a relation as a translation vector on a hyperplane, while TransR and TransD embed entities and relations in distinct vector spaces. RotatE [41] is an improved variant in the complex vector space. Besides, semantic matching-based models exploit similarity-based functions to score relational triples. The scores are computed using bilinear functions in RESCAL [29], DistMult [51], ComplEx [44] and Simple [20], while HolE [28] replaces dot product with circular correlation. Embeddings are given analogical qualities in Analogy [23]. Recently, neural network-based models, including ProjE [34], ConvE [12], R-GCN [33], ConvKB [27], KBGAN [8] and LinkNBed [14], achieve superior link prediction performance. μ KG currently supports TransE, TransR, TransH, TransD, Tucker, DisMult, ComplEx, HolE, Analogy, RESCAL, RotatE, Simple and ConvE.

Entity Alignment. Embedding-based entity alignment models usually consist of two modules, i.e., KG embedding and alignment learning. For KG embedding based on relational facts, many models including MTransE [10], IPTransE [58], JAPE [36], KDCoE [9], BootEA [37], SEA [30], AttrE [43], MultiKE [55] and TransEdge [38] adopt TransE [5] or its improved variants. Most of other models like GCN-Align [49], RDGCN [50] and AliNet [39] adopt GCN due to its powerful representation learning ability. Other models like RSN [15] use recurrent neural networks for KG embedding, respectively. In addition to relational facts, some models such as KDCoE, AttrE, MultiKE, RDGCN and IMUSE [17] also exploit entity attributes for KG embedding and achieve good results. For alignment learning, IPTransE and KDCoE use the pair loss. Besides, JAPE, BootEA, AttrE, RSN and MultiKE let aligned entities in seed alignment share the same or similar embeddings by some tailored data processing skills, which can be also regarded as a special case of the pair loss. GCN-Align and RDGCN use the marginal ranking loss and AliNet uses the limit-based loss. To achieve better performance, some models including IPTransE, BootEA, KDCoE and TransEdge further employ semi-supervised learning. μ KG currently supports MTransE, AttrE, SEA, GCN-Align, RDGCN, IPTransE, JAPE, BootEA, RSN and IMUSE.

Entity Typing. Entity typing seeks to predict the “type entities” of an instance entity. It can be regarded as a special link prediction task across an instance KG and an ontological KG. μ KG currently supports TransE-ET, HolE-ET and RESCAL-ET. Please refer to [26] for more details.

5.2 Knowledge Graph Embedding Libraries

As summarized in Table 1, most of existing libraries for KG embeddings only focus on link prediction, a common KG embedding task. Multi-source KG embedding and entity alignment are only supported by OpenEA [40] and EAKit [54].

Only OpenKE [16], OpenEA and Pykg2vec [53] are developed with TensorFlow, other libraries only support PyTorch. LibKGE [7] is a recent library for link prediction with a high degree of modularity. DGL-KGE [57] is developed based on DGL. It supports PyTorch and XMNet, but not TensorFlow. NeuralKG [56] is a recent Python-based library for diverse representation learning of KGs, but it mainly focuses on rule-based link prediction models. By contrast, our library is more comprehensive than existing work.

6 Conclusion and Future Work

In this paper, we present a new scalable library, μ KG, for multi-source KG embeddings and applications. It facilitates joint representation learning across multi-source KGs. It supports PyTorch and TensorFlow2, and can perform multiple tasks, including link prediction, entity alignment, entity typing, and multi-source link prediction, with advanced implementations of the corresponding embedding models. Extensive experiments validate the effectiveness and efficiency of μ KG. We further demonstrate how jointly learned embeddings can greatly aid KG-powered downstream tasks such as multi-hop KGQA. We show that knowledge transfer in multi-source KGs is an efficient way to improve the performance of KG-powered tasks.

Best Practices of KG Embedding Libraries. The proposed μ KG supports multiple tasks, while few libraries support entity typing and multi-source link prediction. For users who want to carry out these two tasks, μ KG is the best choice. μ KG provides many popular methods in both TensorFlow and PyTorch implementations. If the official code of a model only has one implementation but users need another, μ KG is a good choice. μ KG is still in its early stages, and a few methods do not achieve optimal results. In this case, the original works are more suitable. For the models that μ KG currently does not implement, users can try other libraries, e.g., LibKGE [7] and PyKEEN [1] for link prediction, or OpenEA [40] and EAKit [54] for entity alignment.

Future Work. We plan to integrate more KG embedding models and multi-source KG tasks. We also plan to continually improve our implementations.

Acknowledgments. This work was supported by National Natural Science Foundation of China (No. 61872172), Beijing Academy of Artificial Intelligence (BAAI), and Collaborative Innovation Center of Novel Software Technology & Industrialization. Zequn Sun was also grateful for the support of Program A for Outstanding PhD Candidates of Nanjing University.

References

1. Ali, M., et al.: PyKEEN 1.0: a python library for training and evaluating knowledge graph embeddings. *J. Mach. Learn. Res.* **22**, 82:1–82:6 (2021)

2. Ali, M., Jabeen, H., Hoyt, C.T., Lehmann, J.: The KEEN universe - an ecosystem for knowledge graph embeddings with a focus on reproducibility and transferability. In: ISWC, pp. 3–18 (2019)
3. Balazevic, I., Allen, C., Hospedales, T.M.: TuckER: tensor factorization for knowledge graph completion. In: EMNLP, pp. 5184–5193 (2019)
4. Bollacker, K., Evans, C., Paritosh, P., Sturge, T., Taylor, J.: Freebase: a collaboratively created graph database for structuring human knowledge. In: SIGMOD, pp. 1247–1250 (2008)
5. Bordes, A., Usunier, N., García-Durán, A., Weston, J., Yakhnenko, O.: Translating embeddings for modeling multi-relational data. In: NIPS, pp. 2787–2795 (2013)
6. Boschin, A.: TorchKGE: knowledge graph embedding in python and pytorch. CoRR abs/2009.02963 (2020)
7. Broscheit, S., Ruffinelli, D., Kochsiek, A., Betz, P., Gemulla, R.: LibKGE - a knowledge graph embedding library for reproducible research. In: EMNLP (Demonstration), pp. 165–174 (2020)
8. Cai, L., Wang, W.Y.: KBGAN: adversarial learning for knowledge graph embeddings. In: NAACL, pp. 1470–1480 (2018)
9. Chen, M., Tian, Y., Chang, K., Skiena, S., Zaniolo, C.: Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment. In: IJCAI, pp. 3998–4004 (2018)
10. Chen, M., Tian, Y., Yang, M., Zaniolo, C.: Multilingual knowledge graph embeddings for cross-lingual knowledge alignment. In: IJCAI, pp. 1511–1517 (2017)
11. Chen, X., Chen, M., Fan, C., Uppunda, A., Sun, Y., Zaniolo, C.: Multilingual knowledge graph completion via ensemble knowledge transfer. In: Findings of EMNLP, pp. 3227–3238 (2020)
12. Dettmers, T., Minervini, P., Stenetorp, P., Riedel, S.: Convolutional 2D knowledge graph embeddings. In: AAAI, pp. 1811–1818 (2018)
13. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: NAACL, pp. 4171–4186 (2019)
14. Faloutsos, C., Trivedi, R., Sisman, B., Dong, X.L., Ma, J., Zha, H.: Linknbed: multi-graph representation learning with entity linkage. In: ACL, pp. 252–262 (2018)
15. Guo, L., Sun, Z., Hu, W.: Learning to exploit long-term relational dependencies in knowledge graphs. In: ICML (2019)
16. Han, X., et al.: OpenKE: an open toolkit for knowledge embedding. In: EMNLP (Demonstration), pp. 139–144 (2018)
17. He, F., et al.: Unsupervised entity alignment using attribute triples and relation triples. In: DASFAA, pp. 367–382 (2019)
18. Ji, G., He, S., Xu, L., Liu, K., Zhao, J.: Knowledge graph embedding via dynamic mapping matrix. In: ACL, pp. 687–696 (2015)
19. Ji, S., Pan, S., Cambria, E., Marttinen, P., Yu, P.S.: A survey on knowledge graphs: representation, acquisition, and applications. *EEE Trans. Neural Netw. Learn. Syst.* **33**(2), 494–514 (2022)
20. Kazemi, S.M., Poole, D.: Simple embedding for link prediction in knowledge graphs. In: NeurIPS, pp. 4289–4300 (2018)
21. Lehmann, J., et al.: DBpedia - a large-scale, multilingual knowledge base extracted from Wikipedia. *Seman. Web J.* **6**(2), 167–195 (2015)
22. Lin, Y., Liu, Z., Sun, M., Liu, Y., Zhu, X.: Learning entity and relation embeddings for knowledge graph completion. In: AAAI, pp. 2181–2187 (2015)
23. Liu, H., Wu, Y., Yang, Y.: Analogical inference for multi-relational embeddings. In: ICML, pp. 2168–2178 (2017)

24. Liu, Y., et al.: RoBERTa: a robustly optimized BERT pretraining approach. CoRR abs/1907.11692 (2019)
25. Mahdisoltani, F., Biega, J., Suchanek, F.M.: YAGO3: a knowledge base from multilingual wikipe-dias. In: CIDR (2015)
26. Moon, C., Jones, P., Samatova, N.F.: Learning entity type embeddings for knowledge graph completion. In: CIKM, pp. 2215–2218 (2017)
27. Nguyen, D.Q., Nguyen, T.D., Nguyen, D.Q., Phung, D.Q.: A novel embedding model for knowledge base completion based on convolutional neural network. In: NAACL, pp. 327–333 (2018)
28. Nickel, M., Rosasco, L., Poggio, T.A.: Holographic embeddings of knowledge graphs. In: AAAI, pp. 1955–1961 (2016)
29. Nickel, M., Tresp, V., Kriegel, H.: A three-way model for collective learning on multi-relational data. In: ICML, pp. 809–816 (2011)
30. Pei, S., Yu, L., Hoehndorf, R., Zhang, X.: Semi-supervised entity alignment via knowledge graph embedding with awareness of degree difference. In: WWW, pp. 3130–3136 (2019)
31. Rossi, A., Barbosa, D., Firmani, D., Matinata, A., Merialdo, P.: Knowledge graph embedding for link prediction: a comparative analysis. Trans. Knowl. Discov. Data **15**(2), 14:1–14:49 (2021)
32. Saxena, A., Tripathi, A., Talukdar, P.P.: Improving multi-hop question answering over knowledge graphs using knowledge base embeddings. In: ACL, pp. 4498–4507 (2020)
33. Schlichtkrull, M.S., Kipf, T.N., Bloem, P., van den Berg, R., Titov, I., Welling, M.: Modeling relational data with graph convolutional networks. In: ESWC, pp. 593–607 (2018)
34. Shi, B., Weninger, T.: ProjE: embedding projection for knowledge graph completion. In: AAAI, pp. 1236–1242 (2017)
35. Singh, H., Chakrabarti, S., Jain, P., Choudhury, S.R., Mausam: multilingual knowledge graph completion with joint relation and entity alignment. In: AKB (2021)
36. Sun, Z., Hu, W., Li, C.: Cross-lingual entity alignment via joint attribute-preserving embedding. In: ISWC, pp. 628–644 (2017)
37. Sun, Z., Hu, W., Zhang, Q., Qu, Y.: Bootstrapping entity alignment with knowledge graph embedding. In: IJCAI, pp. 4396–4402 (2018)
38. Sun, Z., Huang, J., Hu, W., Chen, M., Guo, L., Qu, Y.: Transedge: translating relation-contextualized embeddings for knowledge graphs. In: ISWC (2019)
39. Sun, Z., et al.: Knowledge graph alignment network with gated multi-hop neighborhood aggregation. In: AAAI (2020)
40. Sun, Z., et al.: A benchmarking study of embedding-based entity alignment for knowledge graphs. In: PVLDB, vol. 13, pp. 2326–2340 (2020)
41. Sun, Z., Deng, Z., Nie, J., Tang, J.: RotatE: knowledge graph embedding by relational rotation in complex space. In: ICLR (2019)
42. Toutanova, K., Chen, D.: Observed versus latent features for knowledge base and text inference. In: CVSC (2015)
43. Trisedya, B.D., Qi, J., Zhang, R.: Entity alignment between knowledge graphs using attribute embeddings. In: AAAI (2019)
44. Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., Bouchard, G.: Complex embeddings for simple link prediction. In: ICML, pp. 2071–2080 (2016)
45. Vrandečić, D., Krötzsch, M.: Wikidata: a free collaborative knowledgebase. Commun. ACM **57**(10), 78–85 (2014)

46. Wang, Q., Mao, Z., Wang, B., Guo, L.: Knowledge graph embedding: a survey of approaches and applications. *IEEE Trans. Knowl. Data Eng.* **29**(12), 2724–2743 (2017)
47. Wang, X., et al.: KEPLER: a unified model for knowledge embedding and pre-trained language representation. *Trans. Assoc. Comput. Linguistics* **9**, 176–194 (2021)
48. Wang, Z., Zhang, J., Feng, J., Chen, Z.: Knowledge graph embedding by translating on hyperplanes. In: *AAAI*, pp. 1112–1119 (2014)
49. Wang, Z., Lv, Q., Lan, X., Zhang, Y.: Cross-lingual knowledge graph alignment via graph convolutional networks. In: *EMNLP*, pp. 349–357 (2018)
50. Wu, Y., Liu, X., Feng, Y., Wang, Z., Yan, R., Zhao, D.: Relation-aware entity alignment for heterogeneous knowledge graphs. In: *IJCAI* (2019)
51. Yang, B., Yih, W., He, X., Gao, J., Deng, L.: Embedding entities and relations for learning and inference in knowledge bases. In: *ICLR* (2015)
52. Yih, W., Richardson, M., Meek, C., Chang, M., Suh, J.: The value of semantic parse labeling for knowledge base question answering. In: *ACL* (2016)
53. Yu, S., Chhetri, S.R., Canedo, A., Goyal, P., Faruque, M.A.A.: Pykg2vec: a python library for knowledge graph embedding. *J. Mach. Learn. Res* **22**, 16:1–16:6 (2021)
54. Zeng, K., Li, C., Hou, L., Li, J., Feng, L.: A comprehensive survey of entity alignment for knowledge graphs. *AI Open* **2**, 1–13 (2021)
55. Zhang, Q., Sun, Z., Hu, W., Chen, M., Guo, L., Qu, Y.: Multi-view knowledge graph embedding for entity alignment. In: *IJCAI* (2019)
56. Zhang, W., et al.: NeuralKG: an open source library for diverse representation learning of knowledge graphs. *CoRR abs/2202.12571* (2022)
57. Zheng, D., et al.: DGL-KE: training knowledge graph embeddings at scale. In: *SIGIR*, pp. 739–748 (2020)
58. Zhu, H., Xie, R., Liu, Z., Sun, M.: Iterative entity alignment via joint knowledge embeddings. In: *IJCAI*, pp. 4258–4264 (2017)