

KACC: A Multi-task Benchmark for Knowledge Abstraction, Concretization and Completion

Jie Zhou^{1,2*}, Shengding Hu^{1,2*}, Xin Lv^{1,2}, Cheng Yang⁴,
Zhiyuan Liu^{1,2,3†}, Wei Xu⁵, Jie Jiang⁵, Juanzi Li^{1,2}, Maosong Sun^{1,2,3}

¹Department of Computer Science and Technology ²Institute for Artificial Intelligence

³State Key Lab on Intelligent Technology and Systems, Tsinghua University, Beijing, China

⁴School of Computer Science, Beijing University of Posts and Telecommunications, China

⁵Tencent Marketing Solution, Tencent, Shenzhen, China

{zhoujie18, hsd20, lv-x18}@mails.tsinghua.edu.cn

Abstract

A comprehensive knowledge graph (KG) contains an instance-level entity graph and an ontology-level concept graph. The two-view KG provides a testbed for models to “simulate” human’s abilities on knowledge abstraction, concretization, and completion (KACC), which are crucial for human to recognize the world and manage learned knowledge. Existing studies mainly focus on partial aspects of KACC. In order to promote thorough analyses for KACC abilities of models, we propose a unified KG benchmark by improving existing benchmarks in terms of dataset scale, task coverage, and difficulty. Specifically, we collect new datasets that contain larger concept graphs, abundant cross-view links as well as dense entity graphs. Based on the datasets, we propose novel tasks such as multi-hop knowledge abstraction (MKA), multi-hop knowledge concretization (MKC) and then design a comprehensive benchmark. For MKA and MKC tasks, we further annotate multi-hop hierarchical triples as harder samples. The experimental results of existing methods demonstrate the challenges of our benchmark. The resource is available at <https://github.com/thunlp/KACC>.

1 Introduction

Large-scale knowledge graphs (KGs) like Wikidata (Vrandečić and Krötzsch, 2014), DBpedia (Lehmann et al., 2015), and YAGO (Mahdisoltani et al., 2013) usually contain two subgraphs: an instance-level **entity graph** and an ontology-level **concept graph**. The entity graph (a.k.a. the entity view) is composed of entities and relations. It describes factual knowledge such as (*Da Vinci*, *paint*, *Mona Lisa*). The concept graph (a.k.a.

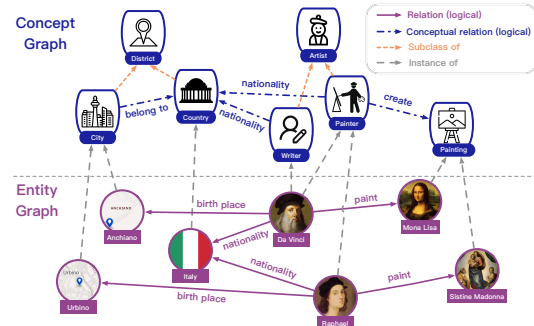


Figure 1: An example of the entity-concept KG.

the concept view) contains concepts and conceptual relations. It provides abstract and common-sense knowledge like (*painter*, *create*, *painting*). In this paper, we name this kind of two-view KG as the entity-concept KG (EC-KG). In a EC-KG, the relations can be grouped into three categories. The “subclassOf” relation forms hierarchical concept structures via triples like (*painter*, *subclassOf*, *artist*). The “instanceOf” relation groups entities into concepts, such as (*Da Vinci*, *instanceOf*, *painter*). These two relations are important for testing models’ abilities on knowledge abstraction and concretization. Other relations are *logical* relations for testing models’ abilities on knowledge completion. An example of the EC-KG is shown in Figure 1.

During the last decade, there are massive works focusing on learning representations for KGs such as TransE (Bordes et al., 2013), DistMult (Yang et al., 2015), ComplEx (Trouillon et al., 2016), and TuckER (Balažević et al., 2019). Though they have achieved promising results on knowledge graph completion, most of them focus on a single graph, especially the entity graph.

Beyond modeling a single graph of KGs, recent studies demonstrate that jointly modeling the two graphs in the EC-KG can improve the understanding of each one (Xie et al., 2016; Moon et al., 2017; Lv et al., 2018; Hao et al., 2019). They also propose

* indicates equal contribution

† Corresponding author: Z.Liu(liuzy@tsinghua.edu.cn)

several tasks on the EC-KG, such as link prediction and entity typing. These tasks focus on partial aspects of knowledge abstraction, concretization, and completion, which are essential abilities for humans to recognize the world and acquire knowledge. For example, in entity typing, a model may link the entity “*Da Vinci*” to the concept “*painter*” which reflects the model’s abstraction ability. However, little work has been devoted to unified benchmarking and studies on KACC.

In this paper, we present a comprehensive benchmark for KACC by improving existing benchmarks in dataset scale, task coverage, and difficulty.

Dataset scale. We have examined the EC-KGs proposed by previous works such as Hao et al. (2019). Due to the data distribution, the concept graphs are small compared to the entity graphs. Furthermore, the cross-links between the two graphs are also sparse (refer to Section 4.3). These may limit the knowledge transfer between the two graphs. To tackle these problems, we construct several different-scale datasets based on Wikidata (Vrandečić and Krötzsch, 2014) with careful filtering, annotation and refinement. As Wikidata contains more fine-grained concepts, our datasets have large concept graphs, abundant cross-view links, as well as dense entity graphs.

Task coverage. Most previous works focus on partial tasks of KACC. In our benchmark, we define the tasks more comprehensively and categorize these tasks into three classes: knowledge abstraction, concretization, and completion.

Difficulty. We propose two new tasks, including multi-hop knowledge abstraction and multi-hop knowledge concretization, which require models to predict multi-hop “instanceOf” and “subclassOf” triples that do not exist in KGs but can be inferred via relation transitivity. These tasks are meaningful and important since correctly modeling these triples is necessary for models to truly understand the concept hierarchy. To ensure the quality of these tasks, we annotate corresponding multi-hop datasets. Our experiments show that these tasks are still challenging for existing models.

Based on our benchmark, we conduct extensive experiments for existing baselines and provide thorough analyses. The experiments show that while the methods specifically designed for modeling hierarchies perform better than general KGE models on abstraction and concretization tasks, they are not competitive to some general KGE models on logi-

cal relations. Moreover, all methods have drastic performance degradation on multi-hop tasks, and the knowledge transfer between the entity graph and the concept graph is still obscure. Finally, we present useful insights for future model design.

2 Related Work

2.1 Knowledge Graph Datasets

Existing datasets for knowledge graph completion are usually subgraphs of large-scale KGs, such as FB15K, FB15K-237, WN18, WN18RR and CoDEX (Bordes et al., 2013; Toutanova et al., 2015; Dettmers et al., 2018; Safavi and Koutra, 2020). These datasets are all single-view KGs, in which FB15K, FB15K-237, and CoDEX focus on the entity view while WN18 and WN18RR can be regarded as concept view KGs. Several datasets try to link the two views in different ways. Firstly, some datasets provide additional type information to the entity graph, such as FB15K+, FB15K-ET and YAGO43K-ET (Xie et al., 2016; Moon et al., 2017). Secondly, some datasets provide concept hierarchies for the entity graph, such as Probase (Wu et al., 2012) and YAGO39K (Lv et al., 2018). However, they do not provide the full concept graphs with logical relations. Thirdly, some datasets provide the full concept graphs (Hao et al., 2019), but both the scale and the depth of the concept hierarchy are limited. For example, the entity numbers of DB111K-174 (Hao et al., 2019) and our dataset KACC-M are similar, but KACC-M has 38 times more concepts than DB111K-174 (see Table 1).

2.2 Knowledge Embedding Methods

Existing knowledge embedding (KE) methods can be categorized as translation models (Bordes et al., 2013; Wang et al., 2014; Lin et al., 2015; Ji et al., 2015; Sun et al., 2019), tensor factorization based models (Yang et al., 2015; Nickel et al., 2016; Trouillon et al., 2016; Balažević et al., 2019), and neural models (Socher et al., 2013; Dettmers et al., 2018; Nguyen et al., 2018). These methods are typically designed for single-view KGs. Although they can be directly applied to EC-KGs by ignoring different characteristics between entity graphs and concept graphs, they cannot take full advantage of the information in EC-KGs.

Several works (Krompaß et al., 2015; Xie et al., 2016; Ma et al., 2017; Moon et al., 2017) incorporate the type information into KE methods to help the completion of entity graphs. ETE (Moon

et al., 2017) further conducts entity typing, which can be seen as a simplified version of our knowledge abstraction task. Though types of entities can be seen as concepts, they omit the concept hierarchy and interactions (conceptual relations) between concepts.

To jointly model the whole EC-KG, TransC (Lv et al., 2018) adopts TransE as the entity graph model and models concepts as spheres that enclosing points of entities. However, it is not flexible enough to model logical relations between concepts. AttH (Chami et al., 2020) further combines hyperbolic embedding methods with KGE methods to simultaneously embed hierarchical and logical relations. JOIE (Hao et al., 2019) uses different training paradigms for training the entity graph, the concept graph, and the cross-view links. It also defines several meaningful tasks on the EC-KG. In this paper, we extend these tasks with several newly proposed tasks, then we categorize and formalize these tasks into a unified benchmark. We also test several KE methods as mentioned above using our benchmarks and analyze their advantages and deficiencies in terms of handling these tasks.

3 Benchmark

In this section, we propose the KACC benchmark with three tasks: knowledge abstraction, knowledge concretization, and knowledge completion.

3.1 Formalizations

We first give formalizations of the EC-KG, then we introduce multi-hop triples used in later tasks.

Formalizations of EC-KG. A EC-KG is a comprehensive KG, which contains two subgraphs and the cross-view links. The entity graph $\mathcal{G}_E = \{\mathcal{E}_E, \mathcal{R}_E, \mathcal{T}_E\}$ is composed of the entity set \mathcal{E}_E , relation set \mathcal{R}_E , and corresponding triple set $\mathcal{T}_E = \{(h^E, r^E, t^E) \mid h^E, t^E \in \mathcal{E}_E, r^E \in \mathcal{R}_E\}$, where h, r, t represent head, relation, and tail of a triple, respectively. The concept graph $\mathcal{G}_C = \{\mathcal{E}_C, \mathcal{R}_C, \mathcal{T}_C\}$ contains the concept set \mathcal{E}_C , conceptual relation set \mathcal{R}_C , and triple set \mathcal{T}_C . In our settings, \mathcal{E}_E and \mathcal{E}_C are disjoint sets, while \mathcal{R}_E and \mathcal{R}_C may contain some relations in common (see Section 4.3). The cross-view links $\mathcal{T}_S = \{(h^S, r_{\text{ins}}, t^S)\}$ connects the two subgraphs, where $h^S \in \mathcal{E}_E, t^S \in \mathcal{E}_C$, and r_{ins} is the “instanceOf” relation. Therefore, the EC-KG is $\mathcal{G} = \{\mathcal{E}_{EC}, \mathcal{R}_{EC}, \mathcal{T}_{EC}\}$, where $\mathcal{E}_{EC} = \mathcal{E}_E \cup \mathcal{E}_C$, $\mathcal{R}_{EC} = \mathcal{R}_E \cup \mathcal{R}_C \cup \{r_{\text{ins}}\}$, and $\mathcal{T}_{EC} = \mathcal{T}_E \cup \mathcal{T}_C \cup \mathcal{T}_S$.

There are two special relations “instanceOf” and “subclassOf” that are crucial for knowledge abstraction and concretization. We use “ins” and “sub” to denote them in the rest of our paper, respectively. The corresponding triples are $\mathcal{T}_{\text{ins}} = \mathcal{T}_S$ and $\mathcal{T}_{\text{sub}} \subset \mathcal{T}_C$. Other relations are *logical* relations. Their corresponding triples in concept graphs are $\mathcal{T}_{C(\text{logic})} = \mathcal{T}_C \setminus \mathcal{T}_{\text{sub}}$ and logical triples in the entity graphs are $\mathcal{T}_{E(\text{logic})} = \mathcal{T}_E$.

Multi-hop Triples. Hierarchical relations like “ins” and “sub” should preserve the multi-hop transitivity, which can be explained by two rules:

$$(e, \text{ins}, c_1) \wedge (c_1, \text{sub}, c_2) \wedge \dots \wedge (c_{N-1}, \text{sub}, c_N) \Rightarrow (e, \text{ins}, c_N), \quad (1)$$

$$(c_0, \text{sub}, c_1) \wedge (c_1, \text{sub}, c_2) \wedge \dots \wedge (c_{N-1}, \text{sub}, c_N) \Rightarrow (c_0, \text{sub}, c_N), \quad (2)$$

in which $\{c_i \mid i \geq 1\}$ are defined as the *high-level concept* for e and c_0 . These two rules indicate that an entity/concept always belongs to its high-level concepts. With these rules, we can collect multi-hop hierarchical triples like (e, ins, c_N) and (c_0, sub, c_N) from the train data and use them as harder samples for knowledge abstraction and concretization testing. Corresponding datasets of multi-hop hierarchical triples are denoted as $\mathcal{T}_{M-\text{Ins}}$ and $\mathcal{T}_{M-\text{Sub}}$.

3.2 Knowledge Abstraction

This task contains tail prediction tasks for one-hop and multi-hop “ins” and “sub” triples. We use KA (knowledge abstraction) and MKA (multi-hop knowledge abstraction) to denote the tasks.

KA-Ins / KA-Sub: KA-Ins and KA-Sub are tail prediction tasks for “ins” triples and “sub” triples respectively. These are all triples in the original datasets and these tasks reflect the direct knowledge abstraction ability of models.

MKA-Ins / MKA-Sub: MKA-Ins and MKA-Sub are tail prediction tasks for multi-hop hierarchical triples $\mathcal{T}_{M-\text{Ins}}$ and $\mathcal{T}_{M-\text{Sub}}$. These tasks reflect models’ abilities on high-level concept abstraction, which aim to predict upper concepts multiple hops away in the concept hierarchy.

3.3 Knowledge Concretization

Similar to KA and MKA tasks, this task contains KC (knowledge concretization) and MKC (multi-hop knowledge concretization) tasks.

KC-Ins / KC-Sub: KC-Ins and KC-Sub are head prediction tasks for “ins” and “sub” triples,

which aim to predict entities for concepts or low-level concepts for high-level ones.

MKC-Ins / MKC-Sub: These tasks are head prediction tasks for multi-hop hierarchical triples. These tasks aim to predict entities for concepts or predict finer concepts for coarser concepts that are multi-hops away.

3.4 Knowledge Completion

The knowledge completion task contains the sub-tasks of entity graph completion (EGC) and concept graph completion (CGC) under two settings. In the “Single” setting, models can only use each single graph to do knowledge graph completion while both the two graphs and the cross-view links are provided in the “Joint” setting.

CGC-Single / EGC-Single: These subtasks are conducted on each single graph \mathcal{G}_C and \mathcal{G}_E . The test phase is conducted on logical triples of each graph \mathcal{T}_E and $\mathcal{T}_{C(\text{logic})}$. The results can be compared with results from CGC/EGC-Joint to test the effectiveness of jointly modeling the two graphs.

CGC-Joint: This subtask requires the model to do link prediction with the full information of the EC-KG \mathcal{G} . The model needs to abstract conceptual knowledge from the entity graph to do link prediction for logical concept graph triples $\mathcal{T}_{C(\text{logic})}$. The results of this subtask can also be used to verify models’ abilities on knowledge abstraction.

EGC-Joint: Models are required to use the guidance from the concept graph to do link prediction for entity graph triples \mathcal{T}_E . For example, a person in the entity graph is more likely to lead some organizations if he is a politician.

4 Dataset Construction

In this section, we first provide the details of our data collection process and annotation process. Then we give a detailed analysis of the statistical characteristics of the datasets.

4.1 Dataset Collection

The dataset construction process has four steps:

Step 1: Entity Filtering. We select entities in FB15K-237 (Toutanova et al., 2015) as our seed entities. We first find out corresponding seed entities in the Wikidata dump via the “FreebaseID” property of each item. Note some entities in Freebase may be labeled as concepts in Wikidata, so we filter out these concepts in our seed entity set. Then we extract one-hop neighbors of the seed entities in

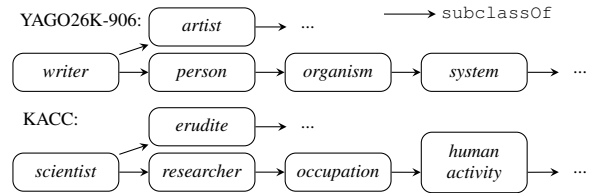


Figure 2: Violation of concept transitivity in two datasets.

the entity graph to form the entity pool, which contains more than 10 million entities. With the entity pool, we can sample an arbitrary size of one-hop neighbors to form the entity graph of our dataset. Our sampling strategy is to select entities with highest degrees and the final entity set consists of all seed entities and the sampled one-hop neighbors. To meet the requirements for different scales, we propose three sizes of datasets: (1) KACC-S, the dataset only contains the seed entities; (2) KACC-M, the expected total entity number is set to 100K; (3) KACC-L, the entity number is set to 1M.

Step 2: Concept Finding. Next, we extract concepts based on selected entities. We use a breadth-first search algorithm to find the concepts. The algorithm starts from entities and search for concepts via “ins” triples and “sub” triples. Since the concept hierarchy follows the structure of a directed acyclic graph, our algorithm ends when all potential concepts are found.

Step 3: Triple Extracting and Filtering. This step firstly extracts cross-view links and all triples in the entity/concept graph. Then we filter all triples by relation statistics and annotation. Relations (1) with less than 10 triples, (2) whose head or tail entity set’s size is smaller than 10, and (3) which are annotated meaningless are dropped. Similar to Toutanova et al. (2015), we further remove reverse relations to prevent valid/test leakage.

Step 4: Concept Filtering. To get more precise concept graphs, we ask human annotators to find out meaningless concepts and we further remove these concepts. These “meaningless” concepts include concepts with no labels or descriptions, concepts used for the self-construction of Wikidata (e.g. “Wikimedia list article”), etc. Details of the annotation step can be found in Appendix A.1.

4.2 Multi-hop Triple Annotation

To support MKA/MKC tasks, we extract multi-hop “ins” and “sub” triples from corresponding train sets according to rule (1) and rule (2). Ideally, hierarchical triples should preserve the multi-hop transitivity. However, when we dive into real-world

Dataset	Entity Graph			Concept Graph			# Cross-links
	# Entities	# Relations	# Triples	# Concepts	# Conceptual Rels	# Triples	
YAGO26K-906	26,078	34	390,738	906	30	8,962	9,962
DB111K-174	111,762	305	863,643	174	20	763	99,748
KACC-S	11,896	82	90,722	2,561	18	6,137	16,061
KACC-M	99,615	209	662,650	6,685	30	15,616	123,342
KACC-L	999,148	377	7,741,272	15,160	44	34,930	1,097,970

Table 1: Statistics of different datasets.

datasets like YAGO26K-906 and ours, we find that the multi-hop transitivity does not always hold true. As illustrated in Figure 2, the transitivity is violated when the transition link goes deep. (*scientist*, *sub*, *occupation*) is meaningful while (*scientist*, *sub*, *human activity*) is not. To make our multi-hop triples meaningful, we further ask human annotators to check the validity of these triples. Details are in Appendix A.2.

4.3 Dataset Analysis

In this subsection, we compare our datasets with existing datasets YAGO26K-906 and DB111K-174 (Hao et al., 2019) in terms of scale, domain coverage, and hierarchical relations. The statistics of these datasets are shown in Table 1.

Scale. From Table 1, we can see that concept graphs in our three datasets have more balanced sizes compared to entity graphs. From the comparison between DB111K-174 and KACC, we can see that entity graphs of these two datasets have similar sizes, but KACC has more concepts, conceptual relations, and triples in the concept graph.

Our datasets also have rich cross-view links. In Table 1, the average numbers of cross-links for each entity are less than 1.0 in YAGO26K-906 (0.38) and DB111K-174 (0.89), which means lots of entities in these datasets are not connected to concepts. In KACC, the ratios are all above 1.09, indicating that one entity may belong to multiple concepts.

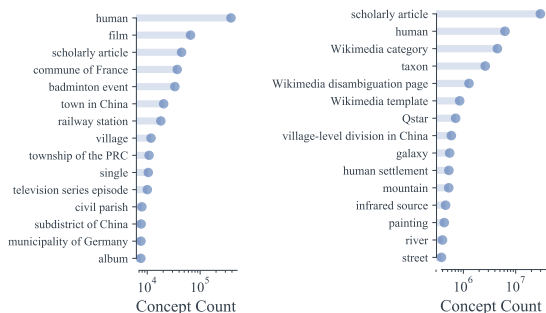


Figure 3: Top 15 most frequent bottom concepts of KACC-L (left) and the original Wikidata dump (right).

Domain Coverage. In Figure 3, we plot 15 most

Dataset	# Duplicate Edges	# Self-loops	# Undetected Concepts
YAGO26K-906	188	44	132 (17.77%)
DB111K-174	27	13	4 (3.48%)
KACC-S	4 → 0	0	26 (1.02%) → 0
KACC-M	13 → 0	0	33 (0.50%) → 0
KACC-L	17 → 0	0	33 (0.22%) → 0

Table 2: Quality check of existing datasets. “→” indicates the filtering process.

frequent bottom concepts, i.e., the concepts that directly connect to entities, in KACC-L and Wikidata dump to illustrate the domains of our datasets. Plots for KACC-S and KACC-M are in Appendix A.4. We find our datasets mainly focus on people, locations, sports, and films, similar to domains of our seed entities extracted from FB15K-237. The comparisons between our datasets and Wikidata dump show that Wikidata dump contains more domains such as scholarly articles, galaxies, and entities related to Wikimedia. Our datasets only focus on partial domains of Wikidata dump, which ensures entities in our datasets are densely connected.

Hierarchical Relations. We present the characteristics of hierarchical relations in our datasets.

We first examine the data quality of *sub* triples in each dataset. We first detect duplicate edges and self-loops. As the global structure of *sub* relations is assumed to be a directed acyclic graph, we use the topological sort algorithm to find loops. We report numbers of concepts that are not detected by the algorithm in each dataset (these concepts are in loops or dangled in loops). The statistics are in Table 2. We can see that our datasets are of high quality with fewer duplicate edges, no self-loops, and less proportion of concepts in loops. Finally, we remove duplicate edges and wrong-labeled triples in loops after a manual check.

Then we examine the depths of the concepts in each dataset. We start from bottom concepts and traverse all concepts via topological sort. We plot numbers of concepts with different depths in Figure 4. From the figure, we can see that our datasets have deeper hierarchical structures than others, which are more informative and useful for

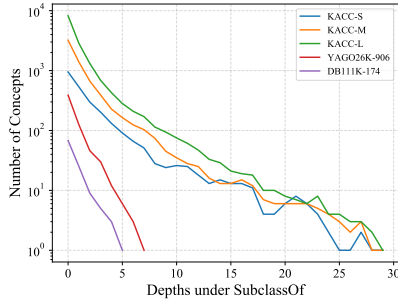


Figure 4: Numbers of concepts with different depths.

Task	Train	Valid	Test
KA-Ins / KC-Ins	$\mathcal{T}_{EC}^{\text{Train}}$	$\mathcal{T}_{\text{ins}}^{\text{Valid}}$	$\mathcal{T}_{\text{ins}}^{\text{Test}}$
KA-Sub / KC-Sub	$\mathcal{T}_{EC}^{\text{Train}}$	$\mathcal{T}_{\text{sub}}^{\text{Valid}}$	$\mathcal{T}_{\text{sub}}^{\text{Test}}$
MKA-Ins / MKC-Ins	$\mathcal{T}_{EC}^{\text{Train}}$	$\mathcal{T}_{\text{M-Ins}}^{\text{Valid}}$	$\mathcal{T}_{\text{M-Ins}}^{\text{Test}}$
MKA-Sub / MKC-Sub	$\mathcal{T}_{EC}^{\text{Train}}$	$\mathcal{T}_{\text{M-Sub}}^{\text{Valid}}$	$\mathcal{T}_{\text{M-Sub}}^{\text{Test}}$
EGC-Joint	$\mathcal{T}_{EC}^{\text{Train}}$	$\mathcal{T}_E^{\text{Valid}}$	$\mathcal{T}_E^{\text{Test}}$
CGC-Joint	$\mathcal{T}_{EC}^{\text{Train}}$	$\mathcal{T}_{C(\text{logic})}^{\text{Valid}}$	$\mathcal{T}_{C(\text{logic})}^{\text{Test}}$
EGC-Single	$\mathcal{T}_E^{\text{Train}}$	$\mathcal{T}_E^{\text{Valid}}$	$\mathcal{T}_E^{\text{Test}}$
CGC-Single	$\mathcal{T}_C^{\text{Train}}$	$\mathcal{T}_{C(\text{logic})}^{\text{Valid}}$	$\mathcal{T}_{C(\text{logic})}^{\text{Test}}$

Table 3: Settings of datasets for different tasks.

models to learn more fine-grained representations.

Finally, we show the characteristic of the “ins” relation in our datasets. Unlike existing datasets where “ins” only connects entities and concepts, concepts in Wikidata also have “ins” connections, which are denoted by $\mathcal{T}_{C(\text{ins})}$. We find these triples are also meaningful as they reflect different level semantics. For example in a triple (*planet*, ins, *astronomical object type*), “planet” is a concept while it can also be regarded as an instance when mentioned in a higher level. The finding is also compatible with human cognition. So we remain these triples in our datasets, and we test them together with other “ins” triples. Therefore, we modify the corresponding definitions in Section 3.1 into $\mathcal{T}_{\text{ins}} = \mathcal{T}_S \cup \mathcal{T}_{C(\text{ins})}$ and $\mathcal{T}_{C(\text{logic})} = \mathcal{T}_C \setminus (\mathcal{T}_{\text{sub}} \cup \mathcal{T}_{C(\text{ins})}$.

Other Characteristics. Our datasets also have some new properties. In existing datasets, relations in the entity graph and conceptual relations in the concept graph are disjoint. However, in our datasets, some relations appear in both the entity graph and the concept graph. For example, the “partOf” relation appears in (*Chile*, partOf, *South America*) in the entity graph and (*hospital*, partOf, *health system*) in the concept graph. Our experiments treat them as different relations while models can also treat these relations as the same, which depends on the hypotheses of the designers.

5 Experimental Settings

5.1 Dataset Partition

Considering the tradeoff between scale and training efficiency, we use the medium-sized dataset KACC-M to conduct the experiments. To generate each task’s train/valid/test data, we firstly split each triple set $\mathcal{T}_S, \mathcal{T}_E$ and \mathcal{T}_C by the proportion 8:1:1. To make it easy for model training and hyperparameter selection, we provide a unified train set $\mathcal{T}_{EC}^{\text{Train}}$ for all tasks defined on the EC-KG, that is $\mathcal{T}_{EC}^{\text{Train}} = \mathcal{T}_E^{\text{Train}} \cup \mathcal{T}_C^{\text{Train}} \cup \mathcal{T}_S^{\text{Train}}$. For the valid and test sets, different tasks have their own valid/test sets for model selection and performance reports.

Train sets are different for EGC-Single and CGC-Single. As they focus on a single graph, we use $\mathcal{T}_E^{\text{Train}}/\mathcal{T}_C^{\text{Train}}$ as train sets respectively. The settings of datasets for different tasks are in Table 3. The statistics are in the Appendix A.3.

5.2 Baselines

To test how existing methods behave in our benchmark, we choose several representative models for single-view KG embedding, as well as JOIE (Hao et al., 2019) and AttH (Chami et al., 2020), which are specially designed for modeling the EC-KG.

Single-view KE Methods. We use TransE (Bordes et al., 2013), DistMult (Yang et al., 2015), ComplEx (Trouillon et al., 2016), and Tucker (Balažević et al., 2019) as our baselines. These baselines treat the EC-KG as a large single-view KG by regarding concepts as entities, conceptual relations and hierarchical relations as ordinary relations defined on a single-view KG.

JOIE. JOIE (Hao et al., 2019) uses traditional KE methods like TransE and DistMult as the backend model to learn logical relations in entity/concept graphs. It further defines specific transformations and loss functions for hierarchical triples. These mechanisms could improve the performance of corresponding backend models.

AttH. AttH (Chami et al., 2020) utilizes the hyperbolic geometry to embed tree-like structures, which is suitable for modeling the concept hierarchy. It also proposes methods to embed logical relations in the hyperbolic space.

5.3 Evaluation Metrics

We test the tasks in the form of link prediction. We use two evaluation metrics in these tasks:

Mean Reciprocal Rank (MRR). The metric computes the mean reciprocal rank of the correct

Method	KA-Ins			MKA-Ins			KA-Sub			MKA-Sub		
	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10
TransE	0.658	0.560	0.832	0.112	0.047	0.242	0.093	0.000	0.288	0.098	0.035	0.225
DistMult	0.712	0.636	0.847	0.131	0.086	0.222	0.135	0.062	0.277	0.122	0.040	0.284
ComplEx	0.737	0.663	0.863	0.121	0.078	0.214	0.226	0.151	0.373	0.135	0.066	0.275
TuckER	0.759	0.681	0.885	0.115	0.077	0.179	0.191	0.107	0.369	0.147	0.068	0.313
JOIE	0.706	0.611	0.873	0.195	0.115	0.370	0.099	0.004	0.289	0.113	0.010	0.351
AttH	0.778	0.693	0.918	0.218	0.116	0.436	0.203	0.089	0.458	0.188	0.081	0.420

Table 4: Results on knowledge abstraction. Best scores are in **bold**.

Method	KC-Ins			MKC-Ins			KC-Sub			MKC-Sub		
	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10	MRR	Hits@1	Hits@10
TransE	0.123	0.083	0.208	0.060	0.033	0.110	0.049	0.000	0.145	0.039	0.000	0.113
DistMult	0.175	0.114	0.282	0.107	0.068	0.183	0.061	0.021	0.136	0.041	0.004	0.128
ComplEx	0.208	0.142	0.321	0.098	0.070	0.153	0.103	0.061	0.179	0.044	0.012	0.107
TuckER	0.169	0.110	0.280	0.074	0.047	0.120	0.087	0.040	0.160	0.051	0.021	0.106
JOIE	0.241	0.200	0.320	0.141	0.093	0.240	0.048	0.010	0.120	0.032	0.001	0.092
AttH	0.172	0.120	0.279	0.112	0.061	0.213	0.081	0.021	0.204	0.051	0.009	0.135

Table 5: Results on knowledge concretization. Best scores are in **bold**.

instances. If the ranks of correct instances are k_i , then the metric computes the average of $\frac{1}{k_i}$.

Hits@N. This metric computes the proportion of the ranks that are no larger than N.

A good model could achieve higher scores on these metrics. We use the ‘‘Filtered’’ setting for all the evaluations, which filters out other true answers from the prediction results to get the final rank for each test case.

5.4 Hyperparameter Settings

According to Ruffinelli et al. (2020), performances of KGE methods are sensitive to hyperparameters. Following them, we run 30 quasi-random trails for all models from predefined hyperparameter spaces. We list the hyperparameter spaces we use in Appendix A.5. We run all trails for 100 epochs.

For all single-view KE methods, we use the implementations from LibKGE (Broscheit et al., 2020), which utilizes the Ax framework to perform quasi-random hyperparameter search.

For AttH, we use the implementation from the authors¹. For JOIE, we use the implementation from the authors². We use TransE as the backend and adopt the suggested hyperparameter space from the paper.

6 Experimental Results

In this section, we provide the experimental results and further propose several future directions.

6.1 Knowledge Abstraction

The results of knowledge abstraction are shown in Table 4. From the results, we can see that AttH has

¹<https://github.com/HazyResearch/KGEmb>

²<https://github.com/JunhengH/joie-kdd19>

a large margin beyond other methods on KA-Ins and also performs well on KA-Sub, which demonstrates the effectiveness of hyperbolic embeddings. JOIE outperforms its backend model TransE.

Comparing results between KA-Ins and MKA-Ins, all the models have performance degradation larger than 0.51 on MRR. We conclude that the composition rule in Equation (1) is hard to learn naturally. Among all the models, AttH performs the best on both tasks and has the least degradation from KA to MKA, showing that hyperbolic space has advantages over Euclidean space in knowledge abstraction. However, the degradation is still drastic, showing the difficulty of the MKA task.

Comparing results between KA-Sub and MKA-Sub, most methods also have performance degradation while TransE-based models (TransE and JOIE) have better performances on MKA-Sub, which is interesting for further investigation. AttH performs best on MKA-Sub, which further confirms the advantage of hyperbolic methods.

6.2 Knowledge Concretization

The results of knowledge concretization are in Table 5. ComplEx and JOIE performs well on KC-Ins and KC-Sub tasks. Similar to tasks in knowledge abstraction, MKC-Ins and MKC-Sub are also harder for existing models. The results of knowledge concretization tasks are lower than corresponding knowledge abstraction tasks, which shows that knowledge concretization is much harder than knowledge abstraction.

6.3 Knowledge Completion

The results of knowledge completion are shown in Table 6. From the table, TuckER performs well

Method	EGC						CGC					
	MRR		Hits@1		Hits@10		MRR		Hits@1		Hits@10	
	Joint	Single	Joint	Single	Joint	Single	Joint	Single	Joint	Single	Joint	Single
TransE	<u>0.305</u>	0.299	<u>0.182</u>	0.175	<u>0.510</u>	0.504	0.242	<u>0.261</u>	0.003	<u>0.095</u>	<u>0.659</u>	0.603
DistMult	<u>0.444</u>	0.440	<u>0.379</u>	0.376	<u>0.566</u>	0.560	<u>0.495</u>	0.481	0.419	<u>0.436</u>	<u>0.631</u>	0.568
ComplEx	<u>0.458</u>	0.453	<u>0.397</u>	0.393	<u>0.572</u>	0.567	0.537	0.501	0.472	0.458	0.684	0.581
TuckER	0.481	0.473	0.415	0.408	0.604	0.595	0.536	0.525	0.468	0.473	<u>0.668</u>	0.615
JOIE	0.171	-	0.094	-	0.308	-	0.218	-	0.018	-	0.622	-
AttH	0.348	<u>0.352</u>	0.235	<u>0.241</u>	<u>0.551</u>	0.545	<u>0.268</u>	0.244	0.100	<u>0.154</u>	<u>0.631</u>	0.418

Table 6: Results on knowledge completion. Best scores among different models in the same task are in **bold**. Best scores for a model between Joint and Single settings are underlined.

Method	KACC	KA	KCon	KCom
TransE	0.374	0.396	0.143	0.585
DistMult	0.396	0.407	0.181	0.599
ComplEx	<u>0.414</u>	<u>0.429</u>	<u>0.184</u>	<u>0.628</u>
TuckER	0.410	0.423	0.165	0.636
JOIE	0.353	0.444	0.177	0.439
AttH	0.452	0.558	0.208	0.591

Table 7: The overall scores. Best scores are in **bold** and second high scores are underlined.

on entity-level logical relations while ComplEx is good at dealing with concept-level logical relations. JOIE does not perform well on logical relations.

From the comparisons between “Joint” and “Single” settings, we find that results on EGC-Joint are usually higher than results on EGC-Single, which shows that incorporating the concept graph and cross-view links helps the understanding of the entity graph. However, the pattern is not obvious on CGC-Joint and CGC-Single, which may due to that entity triples are far more than concept triples, so models tend to focus more on entity triples.

6.4 Overall Results

Finally, we compute an overall KACC score for each method to show their overall performances. Similar to GLUE (Wang et al., 2019a), we average Hits@10 scores of each method on all tasks (except CGC-Single and EGC-Single) to get final scores. We also compute the average scores for knowledge abstraction (KA), knowledge concretization (KCon), and knowledge completion (KCom). In Table 7 we can see that AttH has the best overall score and achieves the highest scores on two tasks. ComplEx also performs well. It is a balanced model since it gets the second place on all tasks. TuckER performs best on knowledge completion. In the future, we plan to test more methods and investigate their abilities.

6.5 Analyses and Future Directions

From the results above, we analyze several problems that existing models cannot handle well and propose several promising future directions.

Multi-hop triple modeling. The prediction scores of multi-hop triples are lower than those of one-hop triples, showing the challenge of multi-hop triple modeling. Besides, how to balance the model to learn from logical and hierarchical relations is also an exciting direction.

Conceptual knowledge completion. Not all models successfully extract conceptual knowledge effectively as their scores of CGC-Joint are lower than those of CGC-Single. The main reason is that KE methods tend to focus more on entity triples due to the losses. They lack the ability to abstract factual knowledge to enrich conceptual knowledge.

Knowledge concretization. The results of concretization tasks are much lower than those of abstraction tasks. It demonstrates that existing models can find proper concepts for entities but cannot find correct entities for concepts. Some solutions may be using contrastive learning to “push” negative entities away from the concepts.

Besides the analyses, there are also several promising future directions of our benchmark.

Contextualized knowledge embedding. Recently, contextualized knowledge embeddings (Wang et al., 2019b) are proposed to capture different semantics of entities and relations in different contexts. These methods only conduct on the entity graph, while incorporating concepts provides more contextual information for entities. For example, an entity of a *painter* is more likely to *paint* than a *politician*. It is a promising direction to model concepts and entities jointly by contextualized embeddings.

Joint modeling EC-KG with text. The EC-KG is a symbolic form of knowledge, and it is interesting to combine it with text. Future directions may include incorporating more commonsense knowledge into the concept graph from language or using the EC-KG to help understand the natural language from the concept level.

7 Conclusion

In this paper, we focus on the problems of knowledge abstraction, concretization, and completion. We propose a benchmark to test the abilities of models on KACC. To conduct the evaluation, we construct large-scale datasets with desired properties, and experiments show that tasks in KACC are challenging. For future work, we plan to test more models and design advanced models to address tasks in KACC.

8 Acknowledgements

This work is funded by the Natural Science Foundation of China (NSFC) and the German Research Foundation (DFG) in Project Crossmodal Learning, NSFC 62061136001 / DFG TRR-169. This work is also supported by Tencent Marketing Solution Rhino-Bird Focused Research Program.

Ethical Considerations

Here we list ethical considerations of our paper:

Intellectual property. All of our datasets are collected from Wikidata and Wikidata offers the data for free with no requirement to attribute under [Creative Commons CC0 License](#).

Privacy. Our datasets are collected from an online resource automatically and the collection process does not involve with participants' privacy rights.

Compensation. For the two annotation processes, the salary for annotating each sample is computed according to the average annotation time and local wage standard. And we ensure that all annotators are well paid.

Potential problems. Though we have manually checked the quality of our datasets and removed meaningless and wrong data, there still may exist false triples. These may lead to wrong predictions in knowledge abstraction, concretization and completion tasks. However, noises are common in human contributed resources such as existing datasets and ours, so the potential risks are low.

References

- Ivana Balažević, Carl Allen, and Timothy M Hospedales. 2019. Tucker: Tensor factorization for knowledge graph completion. In *Proceedings of EMNLP*.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Proceedings of NIPS*.
- Samuel Broscheit, Daniel Ruffinelli, Adrian Kochsiek, Patrick Betz, and Rainer Gemulla. 2020. LibKGE - A knowledge graph embedding library for reproducible research. In *Proceedings of EMNLP: System Demonstrations*.
- Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. 2020. Low-dimensional hyperbolic knowledge graph embeddings. In *Proceedings of ACL*.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of AAAI*.
- Junheng Hao, Muhao Chen, Wenchao Yu, Yizhou Sun, and Wei Wang. 2019. Universal representation learning of knowledge bases by jointly embedding instances and ontological concepts. In *Proceedings of KDD*.
- Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of ACL*.
- Denis Krompaß, Stephan Baier, and Volker Tresp. 2015. Type-constrained representation learning in knowledge graphs. In *Proceedings of ISWC*.
- 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*.
- Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning entity and relation embeddings for knowledge graph completion. In *Proceedings of AAAI*.
- Xin Lv, Lei Hou, Juanzi Li, and Zhiyuan Liu. 2018. Differentiating concepts and instances for knowledge graph embedding. In *Proceedings of EMNLP*.
- Shiheng Ma, Jianhui Ding, Weijia Jia, Kun Wang, and Minyi Guo. 2017. Transt: Type-based multiple embedding representations for knowledge graph completion. In *Proceedings of ECML-PKDD*.
- Farzaneh Mahdisoltani, Joanna Biega, and Fabian M Suchanek. 2013. Yago3: A knowledge base from multilingual wikipedias. In *Proceedings of CIDR*.
- Changsung Moon, Paul Jones, and Nagiza F Samatova. 2017. Learning entity type embeddings for knowledge graph completion. In *Proceedings of CIKM*.
- 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 NAACL*.

- Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016. Holographic embeddings of knowledge graphs. In *Proceedings of AAAI*.
- Daniel Ruffinelli, Samuel Broscheit, and Rainer Gemulla. 2020. You CAN teach an old dog new tricks! on training knowledge graph embeddings. In *Proceedings of ICLR*.
- Tara Safavi and Danai Koutra. 2020. CoDEx: A Comprehensive Knowledge Graph Completion Benchmark. In *Proceedings of EMNLP*.
- Richard Socher, Danqi Chen, Christopher D Manning, and Andrew Ng. 2013. Reasoning with neural tensor networks for knowledge base completion. In *Proceedings of NIPS*.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. Rotate: Knowledge graph embedding by relational rotation in complex space. In *Proceedings of ICLR*.
- 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 EMNLP*.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *Proceedings of ICLR*.
- Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. *Communications of the ACM*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019a. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of ICLR*.
- Quan Wang, Pingping Huang, Haifeng Wang, Song-tai Dai, Wenbin Jiang, Jing Liu, Yajuan Lyu, Yong Zhu, and Hua Wu. 2019b. Coke: Contextualized knowledge graph embedding. *arXiv preprint arXiv:1911.02168*.
- Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of AAAI*.
- Wentao Wu, Hongsong Li, Haixun Wang, and Kenny Q Zhu. 2012. Probbase: A probabilistic taxonomy for text understanding. In *Proceedings of SIGMOD*.
- Ruobing Xie, Zhiyuan Liu, and Maosong Sun. 2016. Representation learning of knowledge graphs with hierarchical types. In *Proceedings of IJCAI*.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2015. Embedding entities and relations for learning and inference in knowledge bases. In *Proceedings of ICLR*.

A Appendices

A.1 Annotations for Meaningless Concepts

In this section, we first present our annotation guidelines for annotators, and then we provide the annotation results.

Task Guidelines. This task aims to find out meaningless “concepts”. For a given instance, you need to check whether it is a “concept”. Here are some definitions in this task:

- **Concept.** A word for a group or a class of things, such as “artist”, “writer”, etc. Humans obtain concepts by abstracting commonalities from things.
- **Entity.** A specific person or thing, such as “Barack Obama”, “Mona Lisa”, etc.

We provide you the Wikidata ID, name, and description of an instance. For more details, you can go to the web “<https://www.wikidata.org/wiki/Qxxxx>” by replacing “Qxxxx” with the specific Wikidata ID. An example is shown in the following:

ID	Name	Description
Q68	computer	general-purpose device for performing arithmetic or logical operations

If an instance is a concept, you should give the correct label. You should give the wrong label in these circumstances:

1. The instance is more like an entity than a concept, such as “Voice over Internet Protocol” (a network protocol).
2. The description and name of the instance are “None”.
3. The instance is used for the website’s construction and is meaningless, such as “Wikimedia list article” and “Wikimedia disambiguation page”.
4. Other cases that are difficult to judge.

Annotation results. We ask human annotators to annotate all concepts in KACC-L. The result of an instance is obtained if two annotators reach the agreement. If not, a third annotator is asked to label the instance. As a result, 482 concepts are removed among 15,642 concepts.

A.2 Annotations for Multi-hop Hierarchical Triples

In this task, we extract multi-hop “instanceOf” and “subclassOf” transitivity links from different train sets and ask annotators to label the position where the hierarchical transitivity holds.

Task Guidelines. This task aims to annotate the transitivity link of concepts. For an example we provided, you need to determine whether semantic drift exists in this link and label the final position that the transitivity holds. Here is some preliminary knowledge:

- **Concept.** A word for a group or a class of things, such as “artist”, “writer”, etc. Humans obtain concepts by abstracting commonalities from things.
- **Entity.** A specific person or thing, such as “Barack Obama”, “Mona Lisa”, etc.
- **Transitivity of concepts.** An example of a transitivity link of concepts is “*scientist* → *researcher* → *occupation* → *human activity*”. The transitivity link starts from an entity or a concept and follows by concepts. The transitivity of concepts assumes that the semantic of the later concept could contain the former concept. For example, the semantic of “*occupation*” contains “*scientist*”, while “*occupation*” is also a more general meaning concept.
- **Semantic drift.** Because of the annotation process of the original data source (Wikidata), we can assume almost all one-hop links are correct, such as “*scientist* → *researcher*” in our example. But semantic drift occurs as the transitivity link goes deep. For example, “*scientist*” can be subclass of “*occupation*” while it cannot belong to “*human activity*”. However, the one-hop link “*occupation* → *human activity*” still holds true.

We provide you the transitivity links of concepts with length 4. These links start from an entity or a concept and is followed by concepts. We provide you the Wikidata ID and the name of the entities and concepts. For more details, you can go to the web “<https://www.wikidata.org/wiki/Qxxxx>” by replacing “Qxxxx” with the specific Wikidata ID. Some examples of this task are shown in Table 8:

Links

- (1) Q4442912 capital of Russia → Q5119 capital → Q515 city → Q702492 urban area
(2) Q3108101 tropical garden → Q1107656 garden → Q386724 work → Q15401930 product

Table 8: Annotation examples for concept transitivity.

You need to label the last position that the concept transitivity holds true starting from the first entity/concept. For example (1) in Table 8, “*capital of Russia*” can be regarded as a sub-concept of “*urban area*”, so the position is 4. In example (2), “*tropical garden*” belongs to “*garden*” while it does not belong to “*work*”, so the position can be labeled as 2.

We can assume that most one-hop links are correct, and you have no need to check the authenticity of them. For example in “*Dewey County* → *county of Oklahoma*”, you do not need to check whether *Dewey* is a county of *Oklahoma*. However, in some specific circumstances, the one-hop link may be wrong, then you can label the case as 1, which means that only the first entity/concept is true.

If you cannot find out meaningful names for entities or concepts, or you meet other cases that are difficult to judge, you can label them as 0.

Annotation Results. We extract 1200, 3000 and 6000 multi-hop “instanceOf” and “subclassOf” triples for KACC-S, KACC-M and KACC-L. These numbers are similar to numbers of “subclassOf” triples in corresponding valid and test sets. We ask two annotators to annotate them, and a third annotator will be added if the two annotators do not reach an agreement. Note that our task requires to label the position, thus there are cases where all these three annotators give different labels. In these cases, we just omit these examples. If a case is labeled as 4, then we can construct both 2-hop and 3-hop triples from the link. If the case is labeled as 3, we can only obtain the 2-hop triple. The statistics of our datasets are in Table 9.

A.3 Statistics of Dataset Split

The statistics of the datasets after partition are shown in Table 10.

A.4 Additional Domain Plot for KACC

We plot the domains of our KACC-S and KACC-M in Figure 5. Domains of KACC-S and KACC-M are similar while KACC-M has more fine-grained

concepts, such as “*town in China*” and “*commune of France*”.

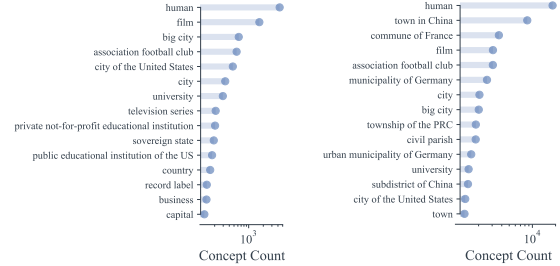


Figure 5: Top 15 most frequent bottom concepts of KACC-S (left) and KACC-M (right).

A.5 Hyperparameter Settings

In this section, we present our hyperparameter selection methods in detail. We run 30 quasi-random hyperparameter search trails on predefined hyperparameter spaces for different baselines (see Table 11 to Table 13). Because we use different implementations, thus hyperparameter spaces are different for different methods.

We run each trail for 100 epochs and save the checkpoint every 20 epochs (150 saved checkpoints for one model in total). Since our benchmark contains multiple tasks, for each task, we use the corresponding valid set to choose the best checkpoint based on the MRR metric, and then we test the selected checkpoint on the test set and compute final metrics.

A.6 Runtime Environment

All experiments are conducted on a server with the following environment.

- Operating System: Ubuntu 18.04.3 LTS
- CPU: Intel(R) Xeon(R) Gold 5218 CPU @ 2.30GHz
- GPU: GeForce RTX 2080 Ti

Dataset	instanceOf			subclassOf		
	# Extracted Links	# 2-hop Links	# 3-hop Links	# Extracted Links	# 2-hop Links	# 3-hop Links
KACC-S	1,200	1,159	1,137	1,200	1,170	1,148
KACC-M	3,000	2,888	2,854	3,000	2,946	2,904
KACC-L	6,000	5,723	5,671	6,000	5,887	5,806

Table 9: Statistics of annotated multi-hop triples.

Data Source	# Train	# Valid	# Test
\mathcal{T}_{EC}	644,332	-	-
\mathcal{T}_E	533,209	-	-
\mathcal{T}_S	98,553	-	-
\mathcal{T}_C	12,570	-	-
\mathcal{T}_E	533,209	64,965	64,476
\mathcal{T}_C	-	1495	1549
\mathcal{T}_{sub}	-	931	965
$\mathcal{T}_{C(\text{logic})}$	-	366	369
$\mathcal{T}_{C(\text{ins})}$	-	198	215
\mathcal{T}_{ins}	-	12,679	12,523
\mathcal{T}_S	-	12,481	12,308
$\mathcal{T}_{C(\text{ins})}$	-	198	215
\mathcal{T}_{sub}	-	931	965
$\mathcal{T}_{C(\text{logic})}$	-	366	369
\mathcal{T}_{M-Ins}	-	2,871	2,871
\mathcal{T}_{M-Sub}	-	2,925	2,925

Table 10: Statistics of data split for KACC-M.

Hyperparameter	Search Range
Training	
Scheme (Complex, DistMult, TuckER)	1vsAll
Scheme (TransE)	Negative Sampling
No. subject samples (TransE)	[1, 1000], log scale
No. object samples (TransE)	[1, 1000], log scale
Batch size	{128, 256, 512}
Loss type	{CE}
Optimizer	
Optimizer	{Adam, Adagrad}
Learning rate	[0.001, 0.1], log scale
Learning rate scheduler's patience	[0, 10]
Embedding	
Dimension	{100, 200, 300}
Initialization	
Std. deviation (Normal)	$[10^{-4}, 1.0]$, log scale
Interval (Uniform)	[-1.0, 1.0]
Regularization	
Entity emb.weight	{None, L_3 , L_2 , L_1 }
Relation emb.weight	$[10^{-20}, 10^{-1}]$, log scale
Frequency weighting	$[10^{-20}, 10^{-1}]$, log scale
Dropout	{True, False}
Entity emb.dropout	[0, 0.5]
Relation emb.dropout	[0, 0.5]

Table 11: Hyperparameter space of quasi-random search for TransE, DistMult, ComplEx, TuckER.

Hyperparameter	Search Range
Training	
Scheme	Negative Sampling
No. negative samples	{1, 10, 20, 50}
Batch size	{128, 256, 512}
Loss type	{F2, N3}
Optimizer	
Optimizer	{Adam, Adagrad}
Learning rate	{0.001, 0.005, 0.01, 0.05, 0.1}
Embedding	
Dimension	{100, 200, 300}

Table 12: Hyperparameter space of quasi-random search for AttH.

Hyperparameter	Search Range
Training	
Backend	{TransE}
Transition method	{CG, CT}
Scheme	Negative Sampling
Batch size	{128, 256, 512}
a_1, a_2	{1.0, 2.5}
m_1, m_2	{0.5, 1.0}
Optimizer	
Learning rate	{0.0005, 0.001, 0.01}
Embedding	
Entity dimension	{50, 100, 200, 300}
Concept dimension	{50, 100, 200, 300}

Table 13: Hyperparameter space of quasi-random search for JOIE.