

Improving Complex Reasoning over Knowledge Graph with Logic-Aware Curriculum Tuning

Tianle Xia¹, Liang Ding², Guojia Wan^{1*}, Yibing Zhan³, Bo Du^{1*}, Dacheng Tao⁴

¹School of Computer Science, National Engineering Research Center for Multimedia Software, Institute of Artificial Intelligence, and Hubei Key Laboratory of Multimedia and Network Communication Engineering, Wuhan University, China

²The University of Sydney

³JD Explore Academy

⁴College of Computing and Data Science at Nanyang Technological University, Singapore 639798
{2018302020065, guojiawan, dubo}@whu.edu.cn, {liangding.liam, dacheng.tao}@gmail.com, zhanyibing@jd.com

Abstract

Answering complex queries over incomplete knowledge graphs (KGs) is a challenging task. Most previous works have focused on learning entity/relation embeddings and simulating first-order logic operators with various neural networks. However, they are bottlenecked by the inability to share world knowledge to improve logical reasoning, thus resulting in suboptimal performance. In this paper, we propose a complex reasoning schema over KG upon large language models (LLMs), containing a curriculum-based logical-aware instruction tuning framework, named LACT. Specifically, we augment the arbitrary first-order logical queries via binary tree decomposition, to stimulate the reasoning capability of LLMs. To address the difficulty gap among different types of complex queries, we design a simple and flexible logic-aware curriculum learning framework. Experiments across widely used datasets demonstrate that LACT has substantial improvements (brings an average +5.5% MRR score) over advanced methods, achieving the new state-of-the-art.

Code — <https://github.com/TianleXia0914/LACT>

Introduction

Industrial-scale knowledge graphs (KGs) like FreeBase store structural knowledge in a collection of fact triplets and are widely adopted by many domains. Unfortunately, KGs are often incomplete, leaving many missing triplets undiscovered. Thus, complex logical reasoning over such incomplete KGs (Hamilton et al. 2018; Bai et al. 2023) is a challenging task and has attracted much attention in the recent years. In the realm of complex logical queries, Existential First Order queries with a single free variable (EFO_1) emerge as a powerful tool that includes logical operators such as the basic projection operation (\exists) and the corresponding negation operation (\neg), two relational operation intersection (\wedge) and union (\vee), etc. A very straightforward approach is to represent the computation graph as a Directed

Acyclic Graph (DAG), which can be addressed by a systematic traversal of a knowledge graph (KG). This process entails the allocation of suitable entities to intermediate variables based on their structural attributes (Dalvi and Suciu 2007).

In the wake of knowledge graph embedding (KGE)’s resounding success, (Bordes et al. 2013; Bai et al. 2021), a series of work seeks to answer complex logical queries by learning query embedding and simulating logical operators with well-designed neural networks (Chen, Hu, and Sun 2022; Zhu et al. 2022; Zhang et al. 2021; Arakelyan et al. 2020). In contemporary research, the exploration of KGE method predominantly centers around the intricate design of various embedding representations, including geometric embedding like Euclidean geometry modeling (Hamilton et al. 2018; Ren, Hu, and Leskovec 2019) and Riemannian geometry modeling (Choudhary et al. 2021), and embedding based on probability modeling (Ren, Hu, and Leskovec 2019), seeking to encapsulate the semantic representation and vector mapping of KGs’ entities and relations.

However, the embedding-based approach described above also has some limitations. (1) **Limited information:** Knowledge graphs often contain incomplete and insufficient information. When only the information from the knowledge graph is available, it becomes challenging to perform complex reasoning that requires additional or missing information. (2) **Insufficient utilization of semantic information:** Current knowledge graph embedding (KGE) methods primarily focus on the structural information of the knowledge graph, paying little attention to the actual semantics of entities and relationships. (3) **High complexity of logical queries:** The intricate nature of real-world knowledge contributes to the complexity of reasoning in practical applications. As a result, it is challenging to accurately model the relationships within real-world knowledge using simple geometric representations. Such simplifications may overlook potentially complex relational information (Choudhary et al. 2021), thereby limiting the effectiveness of advanced logical reasoning. (4) **Generalizability:** A Knowledge Graph Embedding (KGE) model tailored to a specific knowledge graph cannot generalize across different knowledge graphs.

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

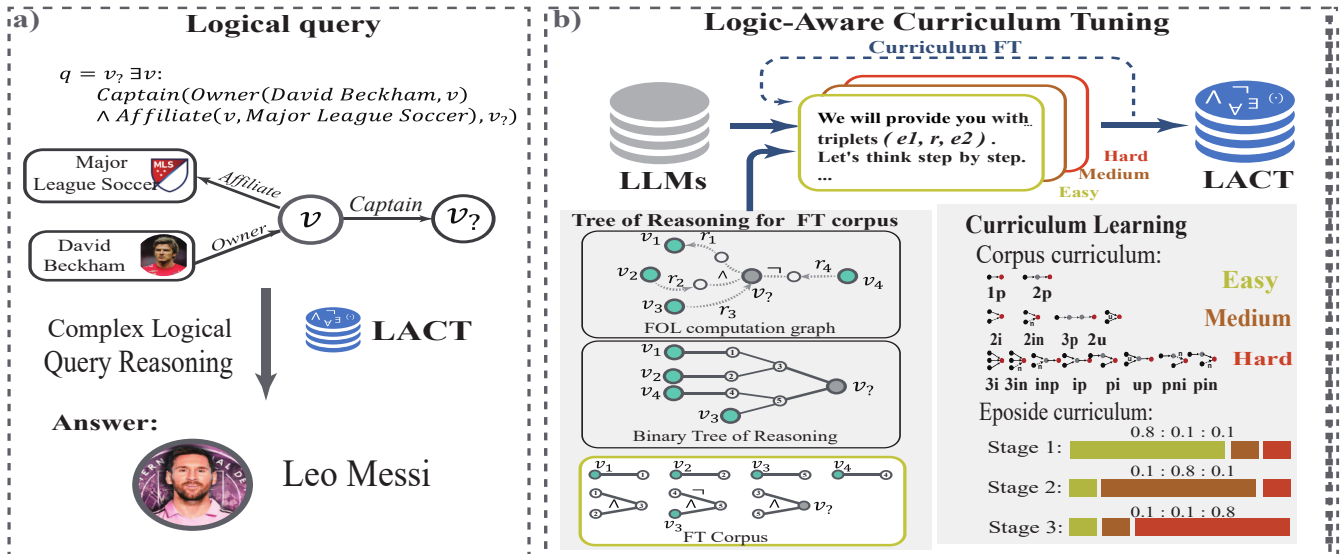


Figure 1: Schematic illustration: a) Answering logical query over KG with LACT. b) The framework of Logic-Aware Curriculum Tuning over Llama. We leverage the binary tree decomposition strategy (Seen in Section Methodology) to construct a logic-rich FT corpus and the Curriculum learning strategy (Seen in Section Methodology) to fine-tune a base LLM. c) Performing reasoning using well-designed prompts.

This limitation hinders the practical application of such approaches in real-world scenarios, where knowledge graphs often vary significantly in both structure and content.

Recently, large language models (LLMs) (Achiam et al. 2023; Touvron et al. 2023a) have showed outperforming capabilities for a wide range of tasks (Ouyang et al. 2022; Zhong et al. 2023b; Peng et al. 2023; Lu et al. 2023; Guo et al. 2025). With the rise of this trend, (Choudhary and Reddy 2023; Liu et al. 2024a) construct prompt templates and apply LLMs as text-generators to answer complex queries. However, LLM without fine-tuning suffers from hallucination problem (Zhang et al. 2023c), especially when faced with such a knowledge-intensive task that generates answers on an incomplete KG rather than simple retrieval. Besides, previous tasks relied on manual classification of queries to improve performance, which is unrealistic in large-scale practical applications and also limits the types of queries. Finally, Previous methods generally decompose the problem into sub-problems, which greatly increases the **cost** of reasoning, especially considering that previous methods are generally based on closed-source models such as GPT. Therefore, compared to pure prompt engineering, we prefer to fine-tune our model to solve the above problems.

In this paper, we propose **Logic-Aware Curriculum Tuning (LACT)**, a novel fine-tune framework for answering complex logical query, which stimulates the ability of LLMs to perform complex reasoning on knowledge graphs. We propose a strategy to incorporate the knowledge contained in the KGs into our training corpus to activate the corresponding knowledge of the LLMs and supplement the missing relevant knowledge of the LLMs during the fine-tuning process. At the same time, we have proven that data

argument by binary tree decomposition can arouse the corresponding capabilities of LLMs and effectively improve their reasoning performance. At last, we show that curriculum learning (Bengio et al. 2009) can effectively smooth the difficulty differences between different types of queries and greatly improve the results of difficult queries. In summary, our contribution manifests in three distinct facets:

- We propose a **Logic-Aware Curriculum tuning (LACT)** paradigm for complex logical reasoning over KGs.
- **LACT** achieves state-of-the-art performance beyond embedding-based and PLM-based methods, using only 7B models.
- Through extensive experimentation, we found that SFT Corpus constructed by queries on incomplete KGs via **Binary Tree Decomposition** and **Curriculum Learning** can significantly enhance LLM logical reasoning ability.

Related Works

Logical Reasoning over Knowledge Graph

Knowledge graphs can be widely used to enhance natural language reasoning (Zhang and Yao 2022; Zhao et al. 2023), understanding (Zhong et al. 2023a; Liu et al. 2023), and generation (Ding, Wang, and Liu 2023; Pan et al. 2024b). Given a EFO_1 query over a KG, complex logical reasoning aims to answer the correct entities, which contains both multi-hop queries and logical operators (Guu, Miller, and Liang 2015; Hamilton et al. 2018) over a incomplete KG. Most of the current approaches have focused on learning meaningful KG embeddings including (Arakelyan et al. 2020; Chen, Hu, and Sun 2022; Zhang et al. 2021; Zhu et al. 2022; Wang et al. 2023b). Neuralizing logical operators through a specific embedding space, thereby embedding EFO_1 queries

into a multi-dimensional vector space (Hamilton et al. 2018; Ren, Hu, and Leskovec 2019), or into a specific probability distribution (Ren, Hu, and Leskovec 2019; Choudhary et al. 2021), then obtain the final answer set by fitting the nearest neighbour representation based on relational operations. Additionally, recent work like CQD (Arakelyan et al. 2020) its optimization work (Arakelyan et al. 2024) improves performance by working to reduce query difficulty by decomposing complex queries into simple one-hop queries, QTO (Bai et al. 2023) introduces query computation tree optimization and LMPNN (Wang et al. 2023f) aggregates logical messages passed on the query graph by training an MLP network. Some work (Yin, Wang, and Song 2023) has begun to expand into the area of cyclic query. Despite their efficacy, embedding-based methods often suffer from a lack of generalization due to their specificity in embedding knowledge graphs. Meanwhile, this limitation on the embedding range hinders their ability to effectively generalize to more complex query structures.

Moreover, PLM-based methods view complex logical reasoning as text-generation tasks by modeling pre-trained language models. EFO_1 queries are organized into input-output sequence pairs after textualization and encoded by PLM (Wang et al. 2023c; Xu et al. 2023; Wang et al. 2023c). However, limited by the performance limitations of the base model, the PLM method has always been deficient in understanding world knowledge and reasoning capabilities, limiting its performance in complex reasoning.

LLMs for KG Reasoning

In recent years, substantial advancements have been witnessed in the domain of LLMs (Achiam et al. 2023; Touvron et al. 2023b; Peng et al. 2023; Zhong et al. 2023b; Zeng et al. 2022). Among these, instruction tuning (IT) (Ouyang et al. 2022) and the alignment of the model (Wang et al. 2023e) with human preferences stand out.

Within the realm of LLM, the integration of LLMs with Knowledge Graphs (KG) (Pan et al. 2024a; Wang et al. 2023a; Luo et al. 2024) constitutes a prominent and consequential research avenue.

Leveraging its potent generative capabilities, LLMs prove invaluable in addressing Knowledge Graph-related tasks, including but not limited to Knowledge Graph Completion (KGC) (Zhu et al. 2023; Zhang et al. 2023b), entity alignment (Zhang et al. 2023a), Knowledge Graph Question Answering (KGQA) (Luo et al. 2024), and others (Luo et al. 2023). Consequently, the synergy between Knowledge Graphs for LLMs (KG4LLM) and LLMs for Knowledge Graphs (LLM4KG) emerges as an essential focal point, bearing significance in advancing the collective capabilities of both entities.

Our research endeavours center around the utilization of Language Models (LLMs) for the Complex Logical Reasoning task, an area that remains relatively unexplored. (Choudhary and Reddy 2023) made the initial attempt by prompt engineer but it lacks in-depth research and simply uses LLMs as text generators.

Curriculum Learning

The concept of progressively training neural networks from simple to complex configurations has its origins dating back to (Elman 1993; Krueger and Dayan 2009). Based on these works, Curriculum Learning is first proposed in (Bengio et al. 2009). A series of illustrative experiments were meticulously crafted to showcase the advantages of employing curriculum strategies in both image classification and language modeling. When focusing on the field of NLP, by experimenting with several heuristics, (Sachan and Xing 2016; Xu et al. 2020) migrated the success of CL to NLU tasks. (Ding et al. 2021; Zhou et al. 2021) improved the machine translation modeling by carefully designing different curricula. Recently, with the rise of LLM, (Liu et al. 2024b) discovered the huge potential of CL in in-context learning, while (Wang et al. 2023d) focus on the improvement of CL for LLM in pre-train. However, we are committed to exploring the huge potential of CL in fine-tuning LLM.

Preliminary

Knowledge Graph

In our work, a knowledge graph is $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ where \mathcal{E}, \mathcal{R} are the set of entity and relation respectively. With regard to generality, KG can be considered as a collection of triplets $\{\mathcal{T} = (v_s, r, v_t) | v_s, v_t \in \mathcal{E}, r \in \mathcal{R}\}$, where v_s/v_t denotes the head/tail entity.

Complex logical query

Complex logical query is used for retrieving and manipulating data stored in knowledge graphs, which is grounded in a subset of EFO_1 query. The process of answering a complex logical query involves trying to match suitable results using the composition of queries:

$$q[v?] = \exists v : q_1 \wedge q_2 \wedge \dots \wedge q_n, \quad (1)$$

or,

$$q[v?] = \exists v : q'_1 \vee q'_2 \vee \dots \vee q'_n, \quad (2)$$

where q denotes a EFO_1 query and q_i denotes a sub-query. Note that Eq. (1) is a conjunctive normal form (CNF) and Eq. (2) is a disjunctive normal form (DNF) (Ren, Hu, and Leskovec 2019). The two can be equivalently converted to each other via De Morgan’s law.

Methodology

Instruction Tuning on LLMs

Within this section, we elucidate the methodology for seamlessly integrating Knowledge Graph (KG) information into the textual prompt.

In the context of complex logical reasoning, we designate a Language Model (LLM) as \mathcal{M} serving as a text decoder to produce the corresponding output. By commencing with the aforementioned definition, we can frame this task as a text generation problem. In contrast to vanilla text generation, triplet generation involves more complexity due to the intricate semantic information associated with the entities and relations in the triplet prompt, as defined by the given knowledge graph (KG). In fact, we want the generated answers to

be entities that exist in KG itself. Without this knowledge, the predicted answers are unreliable and unstable. Thus, to engage LLM in complex logical reasoning, a crucial step involves incorporating knowledge graph (KG) information into the prompt. This integration provides additional auxiliary context, effectively enhancing the performance of LLM.

In particular, when we fine-tune \mathcal{M} , we can treat the training corpus as a set of question-answer pairs $(\mathcal{S}, \mathcal{A})$. When considering complex reasoning over KGs, the input textual sequence \mathcal{S} consists of the description of question \mathcal{D} , knowledge graph neighbourhood information (i.e. related triplets) \mathcal{X} and logical query. In our work, we used a simple but effective method called *greedy depth traversal algorithm* to search for neighbourhood information and we simply discarded all samples that exceeded the token limit. (Detailed algorithm and token distribution can be found in Appendix). The logical query contains the textual information about the query q_τ with the query structure τ and specific query content Q_τ that needs to be processed, which can be denoted as $f_1(q_\tau)$. Likewise, we can denote the output answer \mathcal{A} as $f_2(V_\tau)$, where f_2 indicates textualization of V_τ here. In summary, the fine-tune training corpus \mathcal{C} can be expressed in the following form:

$$\mathcal{C} = (\mathcal{S}, \mathcal{A}) = (\mathcal{D} \oplus \mathcal{X} \oplus f_1(q_\tau), f_2(V_\tau)). \quad (3)$$

The model \mathcal{M} (parameterized by θ) is fine-tuned by the next token prediction task. We fine-tune \mathcal{M} to obtain the final model by maximizing the log-likelihood of the next token. The training loss can be formulated as

$$\mathcal{L} = -\frac{1}{|\mathcal{C}|} \sum_{i=1}^{|\mathcal{C}|} \log P_{\mathcal{M}}(c_i | c_{<i}), \quad (4)$$

where $c_i (i = 1, 2, \dots, |\mathcal{C}|)$ represents the textual tokens of the training corpus \mathcal{C} . For our task, the training objective can be transferred as

$$\mathcal{L} = -\frac{1}{|\mathcal{C}|} \sum_{i=1}^{|\mathcal{C}|} \log P_{\mathcal{M}}(\mathcal{A} | \mathcal{S}). \quad (5)$$

Data Augmentation by Binary Tree Decomposition

This section introduces how to build fine-tuning corpora that make LLMs (Large Language Models) logic-aware based on instruction tuning.

Chain of thought (COT) enables models to decompose multi-step problems into intermediate steps, subsequently improving the reasoning abilities of LLMs (Wei et al. 2022). However, pure-prompt-based reasoning needs more in-context memory to perform complex logical reasoning. Considering complex logical queries, which query structure can be transferred into the form of a DAG and its hierarchical structure becomes a natural fit for decomposition into a series of sub-problems. So we propose a method for data augment based on *Binary Tree Decomposition Mechanism* to stimulate LLMs with the potential to decompose a complex query into a chain of simple queries. **Binary Tree Decomposition Mechanism.** The Binary Tree Decomposition Mechanism is divided into the following three steps:

Query Computation Tree. Existing embedding-based methods rely on DNF to avoid embedding divergence (Ren and Leskovec 2020). Our method bypasses logical operator learning. Consequently, this allows us to convert any query into a logic-equivalent tree form through iterative sub-query recursion. For a complex EFO_1 query, like the example shown in Figure 1, its computation graph that is a directed acyclic graph can be converted into a tree where the root node is v_7 . The tail entities and intermediate entities in complex queries are in one-to-one correspondence with the root nodes and leaf nodes in the generated calculation tree respectively. In view of the reverse order that the tail entity brings to the root node, the edges in the tree corresponding to the relationship in each query are child-to-parent directions. In the Appendix, We demonstrate a simple but systematic method for converting complex queries into computational trees.

Binary Tree Decomposition. We split each 1-to-n intersection/union node into n corresponding child nodes. Consider that merging union branches may result in a 1-to-n structure consisting of intersection and union edges, as shown in Figure 1. This can be properly addressed by separating v_7 into an intersection node structure (v'_3 and v'_5 in Figure 1) firstly, then decomposing the node that was decomposed in the previous layer into a deeper intersection node structure (v'_1 and v'_2 for v'_3 , v'_4 , v_3 for v'_5), where v' denotes an intermediate entity retrieved by Neighborhood Retrieval Algorithm (Seen in Appendix).

Reverse Level Traversal. Finally, we decompose the binary computation tree into independent branches. Since the root node of the calculation tree is the answer entity, we perform a hierarchical traversal of all non-leaf nodes of the binary tree in reverse. As shown in Figure 1, the complex EFO_1 query is decomposed into a sequence:

$$\begin{aligned} &[(v_1, r, v'_1), (v_2, r, v'_2), (v_3, r, v'_3), (v_4, r, v'_4), \\ &(v'_1, r, v'_3), \wedge, (v'_2, r, v'_3), \neg, (v'_4, r, v'_5), \wedge, (v_3, r, v'_5), \\ &(v'_3, r, v_7), \wedge, (v'_5, r, v_7)]. \end{aligned}$$

Data Augmentation. Now we can turn any loopless EFO_1 query into a series of separate subqueries. We use a defined template to integrate the decomposition process into the answers to the training corpus. So, the training corpus \mathcal{C} can be transferred into the following form:

$$\mathcal{C} = (\mathcal{S}, \mathcal{A}) = (\mathcal{D} \oplus \mathcal{X} \oplus f_1(q_\tau), f_2(V_{\tau, Decomposed})), \quad (6)$$

where $V_{\tau, Decomposed}$ indicates the answer corresponding to the logical query with the decomposition reasoning path.

Fine-tuning Enhanced by Curriculum Learning

As mentioned in previous sections, though decomposing into chain responses, complex queries still vary greatly in difficulty and complexity due to differences between query structures.

Naturally, we believe that these different types of samples should not be simply lumped together. Intuitively, we incorporate curriculum learning into our training. To be specific, given the particularity of complex reasoning data, when we decompose it into logical chains, naturally, we can use the

number of decomposed sub-logical queries as a natural difficulty discriminator to select different types of queries, e.g., a 1p query would be defined as difficulty-1, while a 2p query, which can be decomposed into two projection queries and an intersection query, would be defined as difficulty-3. The detailed difficulty discriminating process will be shown in the Appendix.

Finally, we divided samples into three parts: easy samples, medium samples and difficult samples according to the difficulty level. Correspondingly, our training process is also divided into three stages. After we did some exploratory experiments, we did not simply train three data sets in the order of easy-medium-difficult. On the contrary, we decided to first use 80% easy samples, 10% medium samples, and 10% difficult samples for the first stage of training and the subsequent two-stage training process is a Leto, and experimental results in the next few sections also proved that this is effective.

Reasoning Module

We retrieve relevant information and textualize the EFO_1 query, and finally, we populate it into the template in the Appendix to generate queries.

As shown in Appendix, we use the LLM to do a simple text generation task to get the answer. After fine-tuning, LACT can follow the output mode in the training stage in Figure 1, so we can extract final answers through simple regular expressions with the template in the Appendix. The specific prompt can be found in the Appendix.

Experiments

Training Settings

We opt for the most popular datasets: **FB15K**, **FB15K-237**, **NELL995**. We used the training set of the above dataset as the original training data. Detailed information about the dataset and training set is listed in the Appendix.

Experimental Settings

Baseline Methods For comparing with KGE, we chose the following representative methods as baselines: **GQE** (Hamilton et al. 2018), **Query2Box(Q2B)** (Choudhary and Reddy 2023), **BetaE** (Ren and Leskovec 2020), **CQD** (Arakelyan et al. 2020), **ConE** (Zhang et al. 2021), **GNN-QE** (Zhu et al. 2022), **QTO** (Bai et al. 2023). Moreover, we also compared our LACT with the LLM-based method **LARK** (Choudhary and Reddy 2023).

Evaluation Protocol Following previous works (Ren, Hu, and Leskovec 2019), we choose the Mean Reciprocal Rank (MRR) with standard evaluation protocol for evaluating complex reasoning on KG. The detailed setting can be found in the Appendix.

Main Results

The main experiment results of three datasets are shown in Table 1. The baseline training datasets contain five different queries including 1p/2p/3p/2i/3i, so another four queries are considered as OOD types, so the mean value of OOD types of queries is recorded as avg_{ood} . In our experiments,

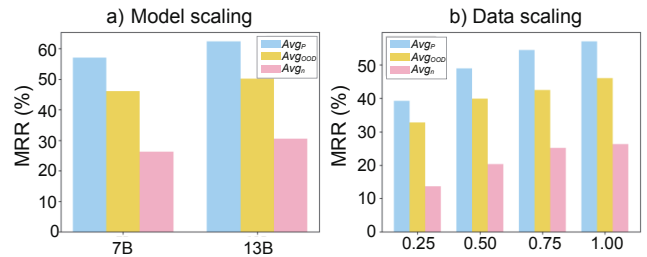


Figure 2: Performance of **scaling** LACT on FB15K-237 with different a) **model** and b) **data** scales.

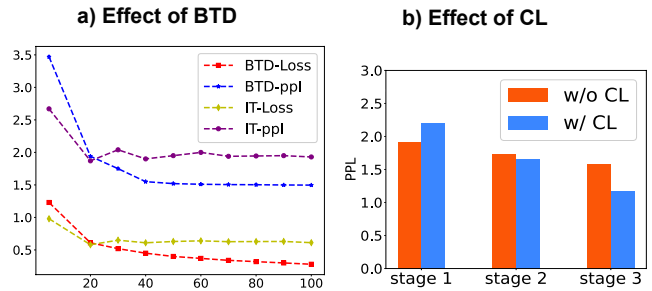


Figure 3: Results of **Ablation Studies**. (a) Comparison **PPL and train loss** results of whether to **use BTD** based on FB15k. (b) Comparison **PPL results** of whether use **CL** based on FB15K high-difficulty queries.

LACT consistently outperforms baseline methods across all datasets. Notably, LACT yields an average gain of 7.3%, 2.9%, and 6.3% on avg_p , avg_{ood} , and avg_n , compared to the previous SOTA method, especially more challenging datasets like FB15K-237 and NELL995. Our method exhibits superior reasoning capability and effectively captures a wide range of internal relations, leading to enhanced performance on complex queries.

Ablation Studies

To validate the efficacy of the LACT module, we conduct a two-part ablation study. Firstly, we will investigate the impact of logical chain decomposition, while the second part assesses the effectiveness of curriculum learning.”

Effect of Binary Tree Decomposition (BTD). As shown in Figure 3, logical chain decomposition can stimulate LLM’s ability of logical decomposition, thereby greatly improving the performance of difficult queries.

From a training perspective, as shown in Figure 3, although perplexity (PPL) and training loss of decomposed queries before training were slightly higher than that of ordinary queries, we found that as training progresses, the loss and PPL of decomposed queries will quickly decrease to levels much lower than ordinary queries, proving that chain decomposition is effective to reduce the difficulty of learning complex queries.

Effect of Curriculum Learning. Curriculum learning, as illustrated in Table 3, greatly alleviates the gap between difficult training samples and the understanding ability of LLMs.

Method	avg_p	avg_{ood}	avg_n	1p	2p	3p	2i	3i	pi	ip	2u	up	2in	3in	inp	pin	pni
Dataset				FB15K													
GQE	28.0	20.1	-	54.6	15.3	10.8	39.7	51.4	27.6	19.1	22.1	11.6	-	-	-	-	-
Query2Box	38.0	29.3	-	68.0	21.0	14.2	55.1	66.5	39.4	26.1	35.1	16.7	-	-	-	-	-
BetaE	41.6	34.3	11.8	65.1	25.7	24.7	55.8	66.5	43.9	28.1	40.1	25.2	14.3	14.7	11.5	6.5	12.4
CQD-CO	46.9	35.3	-	89.2	25.3	13.4	74.4	78.3	44.1	33.2	41.8	21.9	-	-	-	-	-
CQD-Beam	58.2	49.8	-	89.2	54.3	28.6	74.4	78.3	58.2	67.7	42.4	30.9	-	-	-	-	-
ConE	49.8	43.4	14.8	73.3	33.8	29.2	64.4	73.7	50.9	35.7	55.7	31.4	17.9	18.7	12.5	9.8	15.1
GNN-QE	72.8	68.9	38.6	88.5	69.3	58.7	79.7	83.5	69.9	70.4	74.1	61.0	44.7	41.7	42.0	30.1	34.3
QTO	74.0	71.8	49.2	89.5	67.4	58.8	80.3	83.6	75.2	74.0	76.7	61.3	61.1	61.2	47.6	48.9	27.5
LARK	56.1	43.1	18.4	72.8	50.7	36.2	66.9	60.4	56.1	23.5	52.4	40.6	16.2	5.7	33.7	26.1	10.0
LACT	82.6	71.9	56.9	93.5	73.5	59.6	92.3	82.3	76.8	75.9	74.6	60.4	81.2	61.6	52.0	43.5	41.7
Dataset				FB15K-237													
GQE	16.3	10.3	-	35.0	7.2	5.3	23.3	34.6	16.5	10.7	8.2	5.7	-	-	-	-	-
Query2Box	20.1	15.7	-	40.6	9.4	6.8	29.5	42.3	21.2	12.6	11.3	7.6	-	-	-	-	-
BetaE	20.9	14.3	5.5	39.0	10.9	10.0	28.8	42.5	22.4	12.6	12.4	9.7	5.1	7.9	7.4	3.5	3.4
CQD-CO	21.8	15.6	-	46.7	9.5	6.3	31.2	40.6	23.6	16.0	14.5	8.2	-	-	-	-	-
CQD-Beam	22.3	15.7	-	46.7	11.6	8.0	31.2	40.6	21.2	18.7	14.6	8.4	-	-	-	-	-
FuzzQE	24.0	17.4	7.8	42.8	12.9	10.3	33.3	46.9	26.9	17.8	14.6	10.3	8.5	11.6	7.8	5.2	5.8
ConE	23.4	16.2	5.9	41.8	12.8	11.0	32.6	47.3	25.5	14.0	14.5	10.8	5.4	8.6	7.8	4.0	3.6
GNN-QE	26.8	19.9	10.2	42.8	14.7	11.8	38.3	54.1	31.1	18.9	16.2	13.4	10.0	16.8	9.3	7.2	7.8
QTO	33.5	27.6	15.5	49.0	21.4	21.2	43.1	56.8	38.1	28.0	22.7	21.4	16.8	26.7	15.1	13.6	5.4
LARK	50.7	41.0	10.6	73.6	40.5	26.8	46.1	43.1	49.9	22.9	62.8	28.3	6.5	3.4	23.2	16.5	3.2
LACT	57.0	44.4	21.9	76.5	54.3	30.3	56.0	54.5	54.6	36.9	56.5	29.7	17.6	33.1	27.1	19.8	11.2
Dataset				NELL995													
GQE	18.6	12.5	-	32.8	11.9	9.6	27.5	35.2	18.4	14.4	8.5	8.8	-	-	-	-	-
Query2Box	22.9	15.2	-	42.2	14.0	11.2	33.3	44.5	22.4	16.8	11.3	10.3	-	-	-	-	-
BetaE	24.6	14.8	5.9	53.0	13.0	11.4	37.6	47.5	24.1	14.3	12.2	8.5	5.1	7.8	10.0	3.1	3.5
CQD-CO	28.8	20.7	-	60.4	17.8	12.7	39.3	46.6	30.1	22.0	17.3	13.2	-	-	-	-	-
CQD-Beam	28.6	19.8	-	60.4	20.6	11.6	39.3	46.6	25.4	23.9	17.5	12.2	-	-	-	-	-
FuzzQE	27.0	18.4	7.8	47.4	17.2	14.6	39.5	49.2	26.2	20.6	15.3	12.6	7.8	9.8	11.1	4.9	5.5
ConE	27.2	17.6	6.4	53.1	16.1	13.9	40.0	50.8	26.3	17.5	15.3	11.3	5.7	8.1	10.8	3.5	3.9
GNN-QE	28.9	19.6	9.7	53.3	18.9	14.9	42.4	52.5	30.8	18.9	15.9	12.6	9.9	14.6	11.4	6.3	6.3
QTO	32.9	24.0	12.9	60.7	24.1	21.6	42.5	50.6	31.3	26.5	20.4	17.9	13.8	17.9	16.9	9.9	5.9
LARK	52.9	26.9	12.4	87.8	45.7	33.5	51.3	48.7	23.1	22.2	20.6	41.1	9.9	5.9	24.5	13.3	7.3
LACT	60.1	32.0	17.2	91.4	53.6	40.6	62.2	54.9	31.4	34.8	27.0	34.0	16.0	21.2	21.0	16.3	11.6

Table 1: MRR scores (%) on complex reasoning for **14 types of queries**. avg_p represents the mean score of nine ordinary queries including 1p/2p/3p/2i/3i and pi/ip/2u/ip; avg_{ood} is the mean score of out of distribution (OOD) queries, which consist pi/ip/2u/ip queries; avg_n is the mean score of queries including negation operation.

We can observe from Figure 3 that compared with random shuffle sequence training, difficult training samples under curriculum learning gradually become easier to understand. It is worth mentioning that we found that the gain of curriculum learning on training corpus that has not been decomposed by logical chains is very small, which supports our theory from the side. It is difficult for LLMs to understand the difficulty difference between undecomposed samples, so curriculum learning is also difficult to take effect.

Transferability Study

Considering the diversity of complex reasoning tasks, we can divide transferability into two levels, task-level and dataset-level transferability.

Task-level transferability. The results in Table 1 show that our LACT achieves a relative gain of 9.9% on the OOD task, which demonstrates the strong generalization of our fine-

tuning framework. Even in the OOD queries, as shown in the Appendix, more than 95% of test samples can still follow logical chain reasoning. These phenomena indicate the strong generalization ability of LACT.

Dataset-level transferability. In fact, almost all KGE methods, even if some of the optimization methods claim not to require training, require a KGE base model adapted to a specific dataset, which leads to the inherent defect of extremely poor Transferability of the KGE method. However, as previous research has shown, fine-tuning LLMs is mainly to stimulate the knowledge and capabilities of potentially relevant downstream tasks contained in LLM pre-training. This has also become the theoretical basis for the transferability of fine-tuning methods for LLMs. The results in the Appendix show that the reasoning ability stimulated by one dataset can still be demonstrated in another dataset, which reflects well in the query performance which only dropped less than 5%.

Method	1p	2p	3p	2i	3i	pi	ip	2u	up	2in	3in	inp	pin	pni
Llama2	67.2	42.3	38.3	61.6	44.8	34.1	36.9	44.2	28.4	44.7	38.5	36.9	32.1	30.0
+ IT	94.6	68.8	60.2	84.5	66.7	56.0	60.2	69.5	42.3	66.5	54.4	41.0	38.8	36.9
+BTD	91.5	72.3	65.6	89.7	75.2	60.1	65.4	72.5	49.9	74.3	66.4	48.9	46.5	42.5

Table 2: **Accuracy** results (%) evaluated on **FB15k** of whether to **use BTD** for *hard* complex query on all query types.

	BTD	CL	1p	2p	3p	2i	3i	pi	ip	2u	up	2in	3in	inp	pin	pni
Llama2 w/IT			94.6	68.8	60.2	84.5	66.7	56.0	60.2	69.5	42.3	66.5	54.4	41.0	38.8	36.9
	✓		94.2	72.3	65.6	89.7	75.2	60.1	65.4	72.5	49.9	74.3	66.4	48.9	46.5	42.5
		✓	94.6	70.9	61.3	84.4	72.7	58.2	63.2	69.5	47.3	68.5	64.0	42.3	40.9	37.1
	✓	✓	94.8	78.7	69.2	94.8	88.1	79.3	80.5	80.7	67.1	90.6	70.4	59.3	53.6	46.7

Table 3: **Accuracy** results (%) evaluated on **FB15k** of whether to **use CL** for *hard* complex query on all query types.

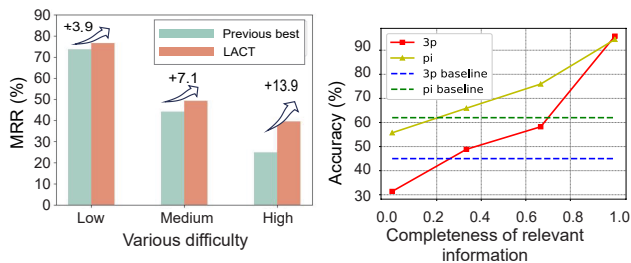


Figure 4: (left) **MRR** performance of LACT and previous SOTA methods at **different difficulties** based on FB15K. (right) The **correlation between related information completeness and accuracy** evaluated on FB15K-237, we selected a 3p query and pi query with the same inference path length as the task types. We assume that the completeness of all simple queries is 1.

Model-level transferability. We tried analytical experiments with different base models to determine whether our LACT was universal. Obviously, The results in the Appendix show that all types of queries have been improved to a certain extent due to the progress of the base model. Experimental results show that our LACT is suitable for different pedestal models and has strong generalization.

Scalability Study

To verify the scalability of LACT, we scale LACT to different model sizes and data volumes.

Performance on different model size. We tried scaling model size to see if LACT would have an impact when operating on a larger scale. As Figure 2 shows, the performance of our method improves as the model size increases.

Performance on different data size. We conducted experiments on different ratios of training data to verify the robustness of LACT.

Discussion

When and Where Does LACT Work?

The performance of LACT would be related to the following two aspects: *I. The completeness of relevant information extracted from KG. II. Sophistication of complex reasoning.*

LACT performs consistently better with more complete information. We take the form of a posteriori that set the completeness of relevant triplets to the proportion of triplets in the inference path of complex reasoning in the provided context, (For example, if the inference path of a 2i query is (Turing Award, winner, Yoshua Bengio), (Canada, citizen, Yoshua Bengio), and we can only retrieve (Turing Award, winner, Yoshua Bengio) in incomplete KG, then our completeness of relevant triplets to 1/2.) and set the completeness of simple queries that can be directly inferred to 1, to obtain the relation between Accuracy and correlation information completeness. As seen in Figure 4, LACT obtains a significant gain when the completeness of relevant information increases, though, with zero relevant information, it remains a certain amount of complex reasoning ability.

LACT performs consistently better on higher difficulties. As mentioned before, we simply divide the difficulty of the query by the number of hops in the query. The results in Figure 4 show that our model yields more gain in tasks of higher-level difficulty and complexity, which benefits from our unique and sophisticated fine-tuning framework.

Conclusion

In this paper, we present a simple and effective fine-tuning framework LACT to boost the complex logical reasoning ability over KGs. LACT is a two-part method that consists of both Binary Tree Decomposition and Curriculum Learning and can be applied to various size LLMs with different data sizes. We empirically demonstrate the effectiveness and universality of the LACT on a series of widely used knowledge graph datasets. Further analyses delve into the underlying mechanism of our LACT, and investigate When and Why Does LACT Work. We hope that our work can inspire more research on combining knowledge graphs and LLMs.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Arakelyan, E.; Daza, D.; Minervini, P.; and Cochez, M. 2020. Complex Query Answering with Neural Link Predictors. In *International Conference on Learning Representations*.
- Arakelyan, E.; Minervini, P.; Daza, D.; Cochez, M.; and Augenstein, I. 2024. Adapting neural link predictors for data-efficient complex query answering. In *Advances in Neural Information Processing Systems*.
- Bai, Y.; Lv, X.; Li, J.; and Hou, L. 2023. Answering complex logical queries on knowledge graphs via query computation tree optimization. In *International Conference on Machine Learning*.
- Bai, Y.; Ying, Z.; Ren, H.; and Leskovec, J. 2021. Modeling heterogeneous hierarchies with relation-specific hyperbolic cones. In *Advances in Neural Information Processing Systems*.
- Bengio, Y.; Louradour, J.; Collobert, R.; and Weston, J. 2009. Curriculum learning. In *International Conference on Machine Learning*.
- Bordes, A.; Usunier, N.; Garcia-Duran, A.; Weston, J.; and Yakhnenko, O. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems*.
- Chen, X.; Hu, Z.; and Sun, Y. 2022. Fuzzy logic based logical query answering on knowledge graphs. In *the AAAI Conference on Artificial Intelligence*.
- Choudhary, N.; Rao, N.; Katariya, S.; Subbian, K.; and Reddy, C. K. 2021. Self-Supervised Hyperboloid Representations from Logical Queries over Knowledge Graphs. In *the Web Conference 2021*.
- Choudhary, N.; and Reddy, C. K. 2023. Complex Logical Reasoning over Knowledge Graphs using Large Language Models. *arXiv preprint arXiv:2305.01157*.
- Dalvi, N.; and Suciu, D. 2007. Efficient query evaluation on probabilistic databases. *The VLDB Journal*.
- Ding, L.; Wang, L.; and Liu, S. 2023. Recurrent graph encoder for syntax-aware neural machine translation. *International Journal of Machine Learning and Cybernetics*.
- Ding, L.; Wang, L.; Liu, X.; Wong, D. F.; Tao, D.; and Tu, Z. 2021. Progressive Multi-Granularity Training for Non-Autoregressive Translation. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*.
- Elman, J. L. 1993. Learning and development in neural networks: the importance of starting small. *Cognition*.
- Guo, C.; Tian, Z.; Tang, J.; Li, S.; and Wang, T. 2025. Multi-pattern retrieval-augmented framework for Text-to-SQL with Poincaré-Skeleton retrieval and meta-instruction reasoning. *Information Processing & Management*, 62(3): 103978.
- Guu, K.; Miller, J.; and Liang, P. 2015. Traversing Knowledge Graphs in Vector Space. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*.
- Hamilton, W.; Bajaj, P.; Zitnik, M.; Jurafsky, D.; and Leskovec, J. 2018. Embedding logical queries on knowledge graphs. In *Advances in neural information processing systems*.
- Krueger, K. A.; and Dayan, P. 2009. Flexible shaping: How learning in small steps helps. *Cognition*.
- Liu, J.; Zhong, Q.; Ding, L.; Jin, H.; Du, B.; and Tao, D. 2023. Unified instance and knowledge alignment pretraining for aspect-based sentiment analysis. *IEEE/ACM transactions on audio, speech, and language processing*.
- Liu, L.; Wang, Z.; Qiu, R.; Ban, Y.; and Tong, H. 2024a. Logic Query of Thoughts: Guiding Large Language Models to Answer Complex Logic Queries with Knowledge Graphs. *arXiv preprint arXiv:2404.04264*.
- Liu, Y.; Liu, J.; Shi, X.; Cheng, Q.; and Lu, W. 2024b. Let's Learn Step by Step: Enhancing In-Context Learning Ability with Curriculum Learning. *arXiv preprint arXiv:2402.10738*.
- Lu, Q.; Qiu, B.; Ding, L.; Xie, L.; and Tao, D. 2023. Error analysis prompting enables human-like translation evaluation in large language models: A case study on chatgpt. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2024*.
- Luo, L.; Ju, J.; Xiong, B.; Li, Y.-F.; Haffari, G.; and Pan, S. 2023. Chatrule: Mining logical rules with large language models for knowledge graph reasoning. *arXiv preprint arXiv:2309.01538*.
- Luo, L.; Li, Y.-F.; Haffari, G.; and Pan, S. 2024. Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning. In *International Conference on Learning Representations*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*.
- Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; and Wu, X. 2024a. Unifying large language models and knowledge graphs: A roadmap. *IEEE Transactions on Knowledge and Data Engineering*.
- Pan, S.; Tian, Z.; Ding, L.; Huang, Z.; Wen, Z.; and Li, D. 2024b. POMP: Probability-driven Meta-graph Prompter for LLMs in Low-resource Unsupervised Neural Machine Translation. *arXiv preprint arXiv:2401.05596*.
- Peng, K.; Ding, L.; Zhong, Q.; Shen, L.; Liu, X.; Zhang, M.; Ouyang, Y.; and Tao, D. 2023. Towards Making the Most of ChatGPT for Machine Translation. In Bouamor, H.; Pino, J.; and Bali, K., eds., *Findings of the Association for Computational Linguistics: EMNLP 2023*.
- Ren, H.; Hu, W.; and Leskovec, J. 2019. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. In *International Conference on Learning Representations*.

- Ren, H.; and Leskovec, J. 2020. Beta embeddings for multi-hop logical reasoning in knowledge graphs. In *Advances in Neural Information Processing Systems*.
- Sachan, M.; and Xing, E. 2016. Easy Questions First? A Case Study on Curriculum Learning for Question Answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Wang, H.; Zhao, S.; Qiang, Z.; Li, Z.; Xi, N.; Du, Y.; Cai, M.; Guo, H.; Chen, Y.; Xu, H.; et al. 2023a. Knowledge-tuning Large Language Models with Structured Medical Knowledge Bases for Reliable Response Generation in Chinese. *arXiv preprint arXiv:2309.04175*.
- Wang, S.; Wei, Z.; Han, M.; Fan, Z.; Shan, H.; Zhang, Q.; and Huang, X. 2023b. Query Structure Modeling for Inductive Logical Reasoning Over Knowledge Graphs. *arXiv preprint arXiv:2305.13585*.
- Wang, S.; Wei, Z.; Xu, J.; Li, T.; and Fan, Z. 2023c. Unifying Structure Reasoning and Language Pre-training for Complex Reasoning Tasks. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Wang, Y.; Zhang, Y.; Li, P.; and Liu, Y. 2023d. Gradual Syntactic Label Replacement for Language Model Pre-Training. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Wang, Y.; Zhong, W.; Li, L.; Mi, F.; Zeng, X.; Huang, W.; Shang, L.; Jiang, X.; and Liu, Q. 2023e. Aligning large language models with human: A survey. *arXiv preprint arXiv:2307.12966*.
- Wang, Z.; Song, Y.; Wong, G. Y.; and See, S. 2023f. Logical message passing networks with one-hop inference on atomic formulas. *arXiv preprint arXiv:2301.08859*.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*.
- Xu, B.; Zhang, L.; Mao, Z.; Wang, Q.; Xie, H.; and Zhang, Y. 2020. Curriculum Learning for Natural Language Understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Xu, Y.; He, S.; Cai, L.; Liu, K.; and Zhao, J. 2023. Prediction and Calibration: Complex Reasoning over Knowledge Graph with Bi-directional Directed Acyclic Graph Neural Network. In *Findings of the Association for Computational Linguistics: ACL 2023*.
- Yin, H.; Wang, Z.; and Song, Y. 2023. Rethinking Complex Queries on Knowledge Graphs with Neural Link Predictors. *arXiv preprint arXiv:2304.07063*.
- Zeng, A.; Liu, X.; Du, Z.; Wang, Z.; Lai, H.; Ding, M.; Yang, Z.; Xu, Y.; Zheng, W.; Xia, X.; et al. 2022. GLM-130B: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414*.
- Zhang, R.; Su, Y.; Trisedya, B. D.; Zhao, X.; Yang, M.; Cheng, H.; and Qi, J. 2023a. AutoAlign: Fully Automatic and Effective Knowledge Graph Alignment enabled by Large Language Models. *IEEE Transactions on Knowledge and Data Engineering*.
- Zhang, Y.; Chen, Z.; Zhang, W.; and Chen, H. 2023b. Making Large Language Models Perform Better in Knowledge Graph Completion. *arXiv preprint arXiv:2310.06671*.
- Zhang, Y.; Li, Y.; Cui, L.; Cai, D.; Liu, L.; Fu, T.; Huang, X.; Zhao, E.; Zhang, Y.; Chen, Y.; et al. 2023c. Siren’s song in the AI ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.
- Zhang, Y.; and Yao, Q. 2022. Knowledge graph reasoning with relational digraph. In *Proceedings of the ACM web conference 2022*.
- Zhang, Z.; Wang, J.; Chen, J.; Ji, S.; and Wu, F. 2021. Cone: Cone embeddings for multi-hop reasoning over knowledge graphs. In *Advances in Neural Information Processing Systems*.
- Zhao, D.; Wan, G.; Zhan, Y.; Wang, Z.; Ding, L.; Zheng, Z.; and Du, B. 2023. KE-X: Towards subgraph explanations of knowledge graph embedding based on knowledge information gain. *Knowledge-Based Systems*.
- Zhong, Q.; Ding, L.; Liu, J.; Du, B.; Jin, H.; and Tao, D. 2023a. Knowledge graph augmented network towards multiview representation learning for aspect-based sentiment analysis. *IEEE Transactions on knowledge and data engineering*.
- Zhong, Q.; Ding, L.; Liu, J.; Du, B.; and Tao, D. 2023b. Can chatgpt understand too? a comparative study on chatgpt and fine-tuned bert. *arXiv preprint arXiv:2302.10198*.
- Zhou, L.; Ding, L.; Duh, K.; Watanabe, S.; Sasano, R.; and Takeda, K. 2021. Self-Guided Curriculum Learning for Neural Machine Translation. In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*.
- Zhu, Y.; Wang, X.; Chen, J.; Qiao, S.; Ou, Y.; Yao, Y.; Deng, S.; Chen, H.; and Zhang, N. 2023. LLMs for Knowledge Graph Construction and Reasoning: Recent Capabilities and Future Opportunities. *arXiv preprint arXiv:2305.13168*.
- Zhu, Z.; Galkin, M.; Zhang, Z.; and Tang, J. 2022. Neural-symbolic models for logical queries on knowledge graphs. In *International Conference on Machine Learning*.

Acknowledgments

This research is supported by NSFC (62225113, 62206202), National Key RD Program of China (2023YFC2705700), The Innovative Research Group Project of Hubei Province under Grants (2024AFA017), the National Research Foundation, Singapore, and the CyberSG R&D Programme Office (“CRPO”), under the National Cybersecurity R&D Programme (“NCRP”), RIE2025 NCRP Funding Initiative (Award CRPO-GC1-NTU-002).