

Retrieval, Reasoning, Re-ranking: A Context-Enriched Framework for Knowledge Graph Completion

Muzhi Li^{2,1*}, Cehao Yang^{3,1*}, Chengjin Xu^{1*†}, Xuhui Jiang¹,
Yiyan Qi¹, Jian Guo^{1†}, Ho-fung Leung[‡], Irwin King^{2†}

¹IDEA Research, International Digital Economy Academy

²Department of Computer Science & Engineering, The Chinese University of Hong Kong

³Artificial Intelligence Thrust, Hong Kong University of Science and Technology (Guangzhou)

{limuzhi, yangcehao, xuchengjin, jiangxuhui, qiyiyan, guojian}@idea.edu.cn

{mzli, king}@cse.cuhk.edu.hk, ho-fung.leung@outlook.com

cyang289@connect.hkust-gz.edu.cn

Abstract

The Knowledge Graph Completion (KGC) task aims to infer the missing entity from an incomplete triple. Existing embedding-based methods rely solely on triples in the KG, which is vulnerable to specious relation patterns and long-tail entities. On the other hand, text-based methods struggle with the semantic gap between KG triples and natural language. Apart from triples, entity contexts (e.g., labels, descriptions, aliases) also play a significant role in augmenting KGs. To address these limitations, we propose KGR³, a context-enriched framework for KGC. KGR³ is composed of three modules. Firstly, the *Retrieval* module gathers supporting triples from the KG, collects plausible candidate answers from a base embedding model, and retrieves context for each related entity. Then, the *Reasoning* module employs a large language model to generate potential answers for each query triple. Finally, the *Re-ranking* module combines candidate answers from the two modules mentioned above, and fine-tunes an LLM to provide the best answer. Extensive experiments on widely used datasets demonstrate that KGR³ consistently improves various KGC methods. Specifically, the best variant of KGR³ achieves absolute Hits@1 improvements of 12.3% and 5.6% on the FB15k237 and WN18RR datasets.

1 Introduction

Knowledge Graphs (KGs) are graph-structured knowledge bases (KBs) that organize factual knowledge as triples in the form of (*head entity, relation, tail entity*). Recently, KGs have become a crucial foundation for various downstream applications, such as recommendation systems (Chen et al., 2022a,b), question answering (Sun et al., 2024), and sentiment analysis (Wang and Shu, 2023). Nev-

*Equal contribution.

†Corresponding authors.

‡Independent researcher.

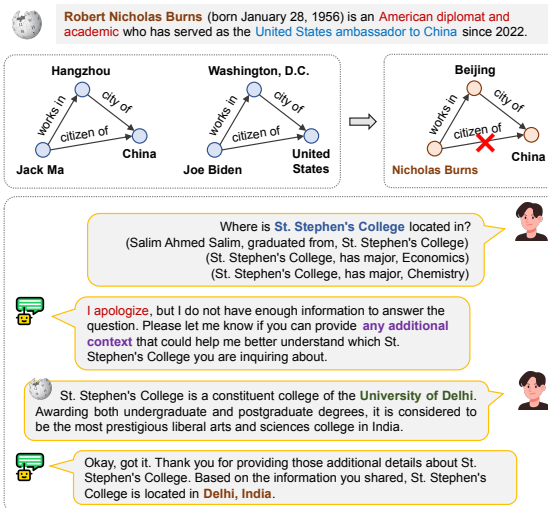


Figure 1: Limitations of existing embedding-based (top) and LLM-based (bottom) KGC methods.

ertheless, mainstream KGs such as Freebase (Bollacker et al., 2008) and Wordnet (Miller, 1995) suffer from serious incomplete issues. This problem highlights the importance of the *Knowledge Graph Completion* (KGC) task, which aims to predict the missing entity from an incomplete triple.

Existing KGC methods can be roughly categorized into embedding-based methods (Bordes et al., 2013; Yang et al., 2015; Sun et al., 2019; Cao et al., 2022) and text-based methods (Yao et al., 2019; Wang et al., 2021, 2022). Embedding-based methods implicitly learn rules based on relation patterns observed in triples and make predictions based on the likelihood of these patterns occurring (Wu et al., 2023). For example, from triple (X , works in, Y) and (Y , city of, Z), it is very likely to deduce that (X , citizen of, Z). However, these methods ignore the contextual semantics supporting these triples, leading to conclusions that do not align with the facts. Text-based methods employ pre-trained language models (PLMs) to embed entities and relations with their labels and descriptions. How-

ever, these methods still cannot surpass the latest embedding-based counterparts (Ge et al., 2023) due to *the substantial semantic gap between structural KG triples and natural language sentences*.

Large language models (LLMs), trained by extensive corpora, demonstrate emergent semantic understanding and in-context learning (ICL) capabilities. Recent studies (Wei et al., 2023; Liu et al., 2024) have proposed utilizing LLMs for the KGC task, as these models harbor general knowledge that can be leveraged to mitigate information scarcity for long-tail entities. However, the application of LLMs in KGC tasks encounters several limitations. Firstly, if the pre-training corpora of the LLMs lack adequate contextual information on specific entities, the LLMs may produce hallucinated or biased responses. Secondly, the structured nature of KG triples limits the ability of LLMs to effectively capture and leverage contextual information from the graph structure. These shortcomings necessitate a strong reliance on a considerable amount of in-context demonstrations (Wei et al., 2023) or external structured embeddings (Liu et al., 2024), which inevitably limit the performance and generality of existing approaches.

Considering the aforementioned challenges, we propose a context-enriched KGC framework named **KGR³**, which consists of three modules: *Retrieval*, *Reasoning*, and *Re-ranking*. Given a query triple, the *retrieval* module gathers semantically relevant supporting triples with the same relation and similar entities, and extracts plausible candidate answers from a base KGC model. To ensure that the LLM attains a fundamental understanding of the retrieved information, this module also collects and augments relevant contextual information to entities in supporting triples and the candidate answer list. Then, the *reasoning* module exploits the semantic understanding capability of pre-trained LLM to suggest several potential answers based on in-context demonstrations and the description of the known entity. Finally, the re-ranking module fine-tunes the LLM to select out the corrupted entity of the training triple from a set of candidate entities, enabling it to process structured knowledge. **KGR³** possesses strong plug-and-play capability, making it compatible with all base KGC methods without costly re-training. During inference, the re-ranking module integrates the candidate answers derived from the base KGC model and the reasoning module, and then instructs the LLM to output the entity label that best completes the query triple.

We validate the proposed framework on two conventional KGC datasets: FB15k237 and WN18RR. The extensive experiments show that KGR³ significantly and consistently outperforms all baseline methods with different types of base KGC models and backbone LLMs, showing its superiority. Notably, the best variant of KGR³ achieves state-of-the-art performance with absolute Hits@1 improvements of 12.3% and 5.6% on the two datasets. Our contributions are summarized as follows:

- We propose a novel KGR³ framework for the KGC task, which systematically retrieves relevant supporting contexts, conducts semantic reasoning, and re-ranks candidate answers.
- We notice the semantic gap between KG triples and natural language sentences, and seamlessly bridge this gap with entity contexts.
- We conduct extensive experiments and ablation studies to evaluate the effectiveness of the KGR³ framework, and discuss the importance of incorporating entity contexts and LLMs.

2 Related Works

Embedding-based methods Embedding-based methods are fundamental in KGC, which focuses on learning a set of low-dimensional embeddings for entities and relations with certain geometric or mathematical constraints. Most typically, TransE (Bordes et al., 2013) assumes the translated head embedding of a triple is close to the embedding of the tail. DistMult (Yang et al., 2015) aims to maximize the Hadamard product of the head, relation, and tail embeddings of each triple. To model symmetric and anti-symmetric relations, ComplEx (Trouillon et al., 2016) generalizes DistMult by introducing complex embeddings with Hermitian dot product. RotatE (Sun et al., 2019) interprets relation as a rotation operation in complex space, which can effectively infer inversion and composition patterns. ATTH (Chami et al., 2020) and GIE (Cao et al., 2022) further leverage hyperbolic embeddings and operations to capture the intrinsic hierarchical structure in KGs. Recently, Graph Neural Networks (GNNs) have emerged as powerful methods for graph embedding (Zhou et al., 2020; Song et al., 2023b,a), with various applications (Ma et al., 2023; Song et al., 2022, 2024). Considering the heterogeneity of KGs, RGCN (Schlichtkrull et al., 2018), WGCN (Zhao et al., 2021), and CompGCN (Vashishth et al.,

2020) adapted to consider relation types in their message-passing functions. KBGAT (Nathani et al., 2019) proposes a two-layer attentional network to encode each triple and to measure its importance to the tail entity. In addition, NBF-Net (Zhu et al., 2021), RED-GNN (Zhang and Yao, 2022) and A*Net (Zhu et al., 2023b) integrate structural information from paths between the two entities, which also supports inductive KGC. Despite the simplicity and high scalability, embedding methods suffer from the long-tail entity distribution in KGs, and hence, cannot generate semantic expressive embeddings for boundary entities, which limits the performance of KGC.

Text-based methods. Apart from the graph structure, textual information in KGs also entails rich semantic knowledge. DKRL proposes to initialize entity embeddings with a convolutional neural network. KG-BERT (Yao et al., 2019) tokenizes triples with textual descriptions of entities and relations, and utilizes BERT (Devlin et al., 2019) to assess their plausibility. KEPLER (Wang et al., 2021) jointly finetunes the pre-trained BERT with the KG embedding and MLM objectives, showcasing improved KGC accuracy. SimKGC (Wang et al., 2022) introduces a contrastive learning strategy, which reduces the computational complexity by re-using in-batch and pre-batch entities as negative samples. CoLE (Liu et al., 2022) extends KG-BERT with co-distillation learning. Inspired by the semantic understanding and reasoning capability of LLMs, KICGPT (Wei et al., 2023) proposes a GPT-based in-context learning (ICL) paradigm for the KGC task. However, due to the large semantic gap between KG triples and natural language sentences, all approaches mentioned above cannot outperform the latest embedding-based methods (Ge et al., 2023; Shang et al., 2024). Recently, DIFT (Liu et al., 2024) devises a supervised finetuning (SFT) solution to guide the LLM in completing triples. Despite achieving state-of-the-art performance, DIFT necessitates costly re-training to adapt to different base KGC models, thereby limiting its compatibility and generality.

3 Problem Specifications

A Knowledge Graph (denoted as $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{T}\}$) can be represented as a set of triples in the form of $(h, r, t) \in \mathcal{T}$, where $h, t \in \mathcal{E}$, $r \in \mathcal{R}$. The notations h and t denote the head and the tail entity of a triple. $\mathcal{E}, \mathcal{R}, \mathcal{T}$ are the set of entities, relations,

and triples, respectively. Besides, KGs are usually associated with knowledge bases (KBs), such as YAGO (Suchanek et al., 2007), Wikidata (Vrandečić and Krötzsch, 2014), and DBpedia (Lehmann et al., 2014). KBs are renowned for representing general knowledge about real-world objects, including people, organizations, places, products, and among others. Apart from triples, KBs also store rich contextual information (or “contexts”) for entities in KGs, which includes *entity labels*, *entity descriptions*, *aliases*, etc. We argue that these entity contexts contain valuable semantic knowledge for the KGC task. In this paper, we exploit LLMs to reveal the missing entity of an incomplete triple based on relevant triples and the contexts of entities involved.

4 Methods

In this section, we introduce our proposed context-enriched KGC framework KGR³, which consists of three components: (1) *Retrieval*, (2) *Reasoning*, and (3) *Re-ranking*.

4.1 Stage 1: Retrieval

The retrieval module focuses on gathering structural and semantic knowledge that may contribute to the completion of certain incomplete triples.

4.1.1 Supporting Triple Retrieval

In KGs, the attributes of an entity are represented in structural triples. Different entities connected by the same relation often share common salient properties. The internal knowledge inherent in the graph structure provides the most direct support to the validity of a triple. Given an incomplete query triple in the form of $(h, r, ?)$ or $(?, r, t)$, we aim to retrieve k supporting triples that are the most semantically similar to the incomplete query triple. Intuitively, we prioritize triples with the same entity and relation from the training set. If the number of available triples is less than k , we broaden our choices to triples with the same relation, and with entities that are semantically similar to the known one in the query triple.

4.1.2 Textual Context Retrieval

We note that there is a significant semantic gap between structural triples and natural language sentences. For example, in Figure 2, entity “Kasper Schmeichel” is originally represented by a unique entity id “/m/07h1h5” while relation “plays for sports teams” is originally represented

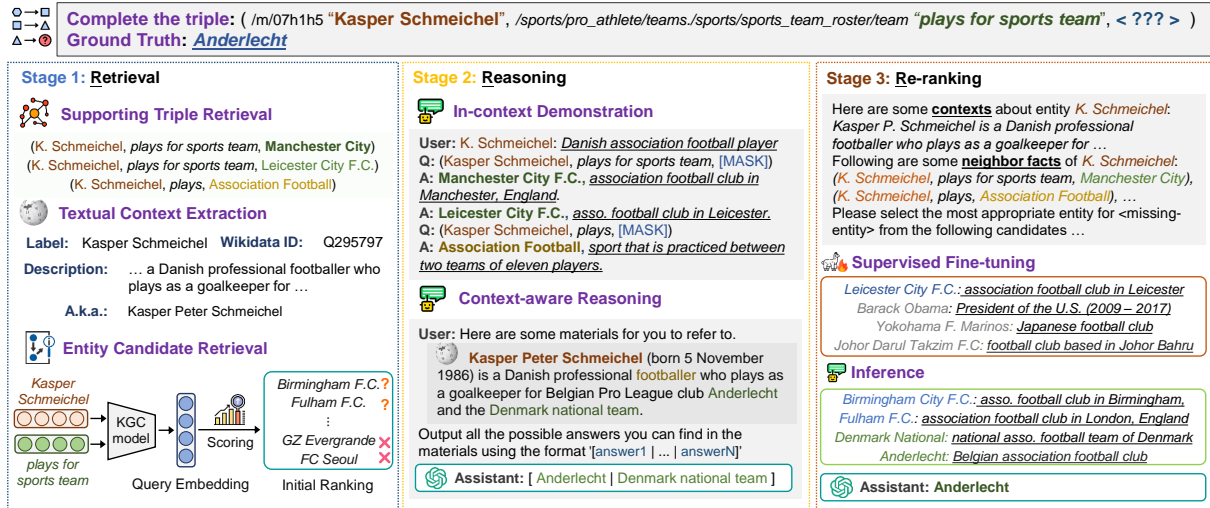


Figure 2: The end-to-end architecture of the proposed KGR³ framework, which consists of three modules: 1) Reasoning (left), 2) Reasoning (middle), and 3) Re-ranking (right). (Please see Appendix A for detailed prompts.)

as `"/sports/pro_athlete/teams./sports/sports_team_roster/team"`. Such a structured format is difficult for LLMs to process. To fully leverage the semantic understanding capabilities of LLMs, we extract relevant contexts related to entities in the query triple and supporting triples from Wikidata knowledge base (Vrandečić and Krötzsch, 2014).

In mainstream KGs, entities are represented in numerical or textual IDs. Each entity ID acts as an index to the data frame in its corresponding knowledge base. Apart from triples, the data-frame of an entity contains significant contextual information such as its entity label. Since Google Freebase (Bollacker et al., 2008) is deprecated and migrated to Wikidata (Vrandečić and Krötzsch, 2014), we map the entity IDs in the FB15k237 dataset to corresponding Wikidata QIDs with official data dumps. We then collect the textual entity label, the short description, and aliases from Wikidata URIs. As for entities in the WN18RR dataset, we adopt the same set of entity labels and descriptions used in KGBERT (Yao et al., 2019) and SimKGC (Wang et al., 2022).

4.1.3 Candidate Answer Retrieval

The widely adopted ranking-based evaluation for the KGC task requires the model to score the plausibility of each entity in the KG as a potential replacement for the missing entity in the query triple. However, given the vast number of entities in the KG, employing LLMs to score and rank each entity is computationally expensive and impractical. Inspired by (Lovell et al., 2021; Wei et al., 2023; Li et al., 2024), we employ a base KGC model to

initialize the scoring and ranking of entities within the KG. Formally, we denote the ranked entity list \mathcal{A}_{KGC} as following:

$$\mathcal{A}_{KGC} = [e_1^{(k)}, e_2^{(k)}, \dots, e_{|\mathcal{E}|}^{(k)}], \quad (1)$$

$$\text{where } e_i^{(k)} = \arg \max_{e \in \{\mathcal{E} \setminus \{e_{<i}^{(k)}\}\}} f_r(h, e) \text{ or } f_r(e, t),$$

$f_r(h, t)$ is the scoring function the KGC model evaluating the plausibility of a triple (h, r, t) . Then, we extract the top- n entities with the highest scorings from \mathcal{A}_{KGC} as candidate answers, and extract their labels and descriptions based on their entity IDs. The plausibility of these candidate answers will be re-evaluated in the re-ranking module.

4.2 Stage 2: Reasoning

In the second stage, we first exploit the supporting triples to direct the LLM in performing the KGC task. In addition, we employ the LLM to generate several possible answers drawing upon the description of the known entity in the query triple.

4.2.1 Supporting Triple Demonstrations

In this sub-section, we present the design of our prompt tailored for the demonstration, which is crucial in facilitating in-context learning (Wei et al., 2023). For each supporting triple, we first provide LLM with the description of the head entity. The entity description serves three objectives: (1) *disambiguate entities with the same label*, (2) *rejuvenate the LLM’s memory about known entities*, and (3) *provide essential information for entities that are not included in the LLM pre-training corpus*.

Following the convention of (Zhu et al., 2023a; Wei et al., 2023), we corrupt the neighboring entity with a “[MASK]” token. In order to narrow the semantic gap between structural triples and natural language sentences, we translate the masked triple into a natural language question. Subsequently, we ask the LLM to generate answers based on its semantic understanding of the short description of the known entity and the question. Different from KICGPT (Wei et al., 2023), all supporting triples included in our demonstration prompts are sampled from the training set, which avoids potential information leakage. Finally, we provide the label of the corrupted entity and its entity description as our explanation. Considering multiple supporting triples, the LLM can also discern salient properties of adjacent neighbors connected by the same relation and similar entities. These salient properties play a vital role in helping the LLM to find out proper answers for the query triple.

4.2.2 Context-aware Reasoning

We anticipate that LLMs can harness their information extraction and semantic understanding capabilities by utilizing comprehensive contextual information about the known entity, thereby generating potential answers. Similarly, we pass the description of the known entity and the question translated from the query triple to the LLM. The LLM is then instructed to output a list of answers in its response. It should be noted that generative LLMs do not guarantee that output answers will conform to entities in the KG. Therefore, we post-process the LLM output by replacing entity aliases with entity labels and filtering out invalid and unreliable answers that do not appear within the top- δ positions of \mathcal{A}_{KGC} . Finally, we obtain a list of m answers, which are formally denoted as:

$$\begin{aligned} \mathcal{A}_{\text{LLM}} &= [e_1^{(l)}, e_2^{(l)}, \dots, e_m^{(l)}] \\ &= f_{\text{LLM}}(q, c_e^{(q)}, \mathcal{D}(q)) \cap \mathcal{A}_{\text{KGC}}[0 : \delta], \end{aligned} \quad (2)$$

where $q, c_e^{(q)}$ and $\mathcal{D}(q)$ denote the question, description of the known entity, and supporting triple demonstrations. Entities in \mathcal{A}_{LLM} are ensured to be simultaneously supported by the LLM and the base KGC model.

4.3 Stage 3: Re-ranking

Motivated by the complementary nature of semantic and structural knowledge, we aim to exploit the

candidate answer list generated by the LLM and the base KGC model to compose our final rankings.

To better enable the LLM in utilizing entity descriptions and structured neighbor facts for ranking candidate answers to query triples, we introduce supervised fine-tuning (SFT) with LoRA adaptation (Chao et al., 2024). Inspired by DIFT (Liu et al., 2024), the training objective of SFT is to find out the missing entity of an incomplete triple from a set of candidate answers. Specifically, we construct training samples by corrupting the tail (or head) entity of each triple in the training set. For each corrupted triple, we randomly sample $n - 1$ negative samples from the entity set, where half of them are connected by the same relation as the corrupted ground truth entity. Incorporating these hard negative samples helps the LLM to distinguish between different entities with the same property, which is crucially important since candidate entities suggested by base KGC models usually yield similar characteristics. After that, we translate the masked triple to a natural language question, and gather the label and description for each candidate entity. Finally, we provide the question q , retrieved neighbor facts $\mathcal{N}(q)$, and the description of the known entity $c_e^{(q)}$, along with candidate answers \mathcal{A} and their descriptions $c(\mathcal{A})$ to the LLM, and fine-tune the LLM to output the label y of the ground truth entity. Formally, we have the SFT loss:

$$\mathcal{L}_{\text{SFT}} = - \sum_{i=1}^{|\mathcal{T}|} \log(y|q, \mathcal{N}(q), c_e^{(q)}, \mathcal{A}, c(\mathcal{A})). \quad (3)$$

It is important to emphasize that the SFT process does not rely on the prior inference results from existing KGC approaches. This ensures that the KGR³ framework can be implemented as a plug-and-play solution.

During the inference stage, we construct a candidate answer set \mathcal{A}_c , which is composed of top- p entities from \mathcal{A}_{KGC} and top- $(n - p)$ entities in \mathcal{A}_{LLM} that are not previously encountered. Formally, we have:

$$\mathcal{A}_c = \mathcal{A}_{\text{KGC}}[0 : p] \cup \{\mathcal{A}_{\text{LLM}} \setminus \mathcal{A}_{\text{KGC}}[0 : p]\}[0 : n - p]. \quad (4)$$

If \mathcal{A}_{LLM} contains fewer than $(n - p)$ entities (e.g. $m < n - p$), we supplement the candidate answer set with additional entities from \mathcal{A}_{KGC} to reach a total of n entities. Similarly, we instruct the fine-tuned LLM to select the most appropriate candidate entity to complete the query triple, and output its

entity label. The selected entity e_{ans} is then prioritized and moved to the front of the candidate answer list \mathcal{A}_c . Finally, we construct the re-ordered entity list \mathcal{A}_{RR} for performance evaluation. Here, we have

$$\mathcal{A}_{\text{RR}} = [e_{\text{ans}} \parallel \mathcal{A}_c \setminus \{e_{\text{ans}}\} \parallel \mathcal{A}_{\text{KGC}}[n : |\mathcal{E}|]], \quad (5)$$

where $[\cdot \parallel \cdot]$ denotes the concatenation operation.

5 Experiments

In this section, we assess the effectiveness of the KGR³ framework in the KGC task.

5.1 Datasets

We utilize two widely-used benchmark datasets, namely FB15k237 and WN18RR to evaluate the proposed method. FB15k237 is derived from Freebase (Bollacker et al., 2008), an encyclopedic KG containing general knowledge about topics such as celebrities, organizations, movies and sports. WN18RR is a subset of WordNet (Miller, 1995), a lexical KG with knowledge about English morphology. To prevent potential data leakage, FB15k237 and WN18RR excludes reversible relations from their backend KGs. Detailed statistics of the two datasets are shown in Table 1.

Dataset	FB15k237	WN18RR
#Entities	14,541	40,943
#Relations	237	11
#Train	272,115	86,835
#Valid	17,535	3,034
#Test	20,466	3,134

Table 1: Statistics of Datasets

5.2 Baselines and Evaluation Metrics

We compare the KGR³ framework with four types of baseline methods: (1) traditional KG Embedding methods TransE (Bordes et al., 2013), ComplEx (Trouillon et al., 2016), RotatE (Sun et al., 2019), TuckER (Balazevic et al., 2019); (2) GNN-based embedding methods CompGCN (Vashishth et al., 2020), NBF-Net (Zhu et al., 2021); (3) text-based methods KG-BERT (Yao et al., 2019), MEM-KGC (Choi et al., 2021), SimKGC (Wang et al., 2022), CoLE (Liu et al., 2022); and (4) LLM-based methods KICGPT (Wei et al., 2023), DIFT (Liu et al., 2024). Among these baselines, we select TransE, RotatE, GIE, SimKGC, CoLE, and NBF-Net as our base KGC models because these methods are highly representative and exhibits strong performance.

We utilize the widely adopted evaluation metrics, namely Hits@ k ($k = 1, 3, 10$) and MRR to evaluate our proposed method. Hits@ k measures the proportion of query triples which the ground truth entities are ranked within the top- k position. MRR measures the mean reciprocal rank for each ground truth entities. Higher results indicates a better performance. Implementation details of baseline models and our models are described in Appendix B.

5.3 Main Results

Table 2 summarizes the performance of the KGR³ framework on six different base KGC methods. The experiment results show that the best variant of KGR³ significantly outperforms all baseline methods among all evaluation metrics. Compared to the previous state-of-the-art baseline, KGR³, with Llama3-8B as the backbone LLM, achieves absolute Hits@1 improvements of 11.1% and 4.0% on the FB15k237 and the WN18RR datasets, respectively. * It demonstrates that KGR³ is highly effective for the KGC task.

Notably, the improvement in Hits@1 is more substantial than that in Hits@3 and Hits@10. This indicates that the KGR³ framework is particularly effective at identifying the most accurate answers. Since our framework primarily focuses on re-ordering top- n entities from the initial ranked entity list, the upper bound of Hits@1, Hits@3, and Hits@10 are implicitly constrained by the Hits@ n performance of the base KGC model. Given that Hits@1 is typically further from this upper bound, the potential for improvement will be greater. Additionally, by leveraging semantic knowledge from entity contexts, the LLM gains a more comprehensive understanding of the entities, thereby enabling more precise inferences, particularly for top-ranked candidate answers.

Compared to the selected base KGC models, the corresponding variants of KGR³ consistently boost the performance on both datasets to a large margin. It shows that KGR³ is compatible with various types of KGC models, confirming its strong generality and plug-and-play capability. In contrast, KICGPT fails to outperform the state-of-the-art GNN-based or text-based methods, which underutilizes the power of the LLM. In addition, variants of KGR³ consistently outperform counterparts of the LLM-based baseline DIFT (Liu et al., 2024) with the same base KGC models. The performance

*KGR³ obtains better performance with Qwen2-7B LLM.

Methods	FB15k237				WN18RR			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
<i>KG embedding methods</i>								
TransE	0.279	0.198	0.376	0.441	0.243	0.043	0.441	0.532
CompLEx	0.247	0.158	0.275	0.428	0.440	0.410	0.460	0.510
RotatE	0.338	0.241	0.375	0.533	0.476	0.428	0.492	0.571
TuckER	0.358	0.266	0.394	0.544	0.470	0.443	0.482	0.526
GIE	0.362	0.271	0.401	0.552	0.491	0.452	0.505	0.575
HittER	0.373	0.279	0.409	0.558	0.503	0.462	0.516	0.584
<i>Graph neural network-based methods</i>								
CompGCN	0.355	0.264	0.390	0.535	0.479	0.443	0.494	0.546
NBF-Net	0.415	0.321	0.450	0.599	0.551	0.497	0.573	0.666
<i>Text-based methods</i>								
KG-BERT	-	-	-	0.420	0.216	0.041	0.302	0.524
MEM-KGC	0.346	0.253	0.381	0.531	0.557	0.475	0.604	0.704
SimKGC	0.338	0.252	0.364	0.511	0.671	0.595	0.719	0.802
CoLE	0.389	0.294	0.429	0.572	0.593	0.538	0.616	0.701
<i>Large language model-based methods</i>								
ChatGPT (0-shot)	-	0.237	-	-	-	0.190	-	-
ChatGPT (1-shot)	-	0.267	-	-	-	0.212	-	-
KICGPT	0.412	0.327	0.448	0.581	0.564	0.478	0.612	0.677
DIFT + TransE	0.389	0.322	0.408	0.525	0.491	0.462	0.496	0.560
DIFT + SimKGC	0.402	0.338	0.418	0.528	<u>0.686</u>	<u>0.616</u>	<u>0.730</u>	<u>0.806</u>
DIFT + CoLE	0.439	0.364	0.468	0.586	0.617	0.569	0.638	0.708
KGR³ + TransE	0.456	0.414	0.474	0.550	0.506	0.487	0.515	0.556
KGR³ + RotatE	0.456	0.400	0.476	0.569	0.520	0.495	0.520	0.550
KGR³ + GIE	0.463	0.400	0.485	0.581	0.558	0.520	0.580	0.615
KGR³ + SimKGC	0.471	0.429	0.490	0.557	0.717	0.656	0.759	0.809
KGR³ + CoLE	<u>0.507</u>	<u>0.455</u>	<u>0.537</u>	<u>0.612</u>	0.635	0.579	0.676	0.723
KGR³ + NBF-Net	0.535	0.475	0.564	0.635	0.641	0.605	0.662	0.695

Table 2: Experiment results of the KGC task on FB15k-237 and WN18RR datasets. The best results are in **bold** and the second-best ones are underlined. All results of baseline methods are referred from corresponding original papers. For KGR³, we adopt LLama3-8B as the backbone LLM of the Reasoning and Re-ranking module.

improvements can be primarily attributed to the incorporation of the reasoning module and the inclusion of entity descriptions for candidate answers. It should also be noted that DIFT necessitates an expensive re-training process and pre-trained KG embeddings to adapt to various KGC models (Liu et al., 2024). Without such a process, DIFT cannot guarantee the experimental results as claimed.

5.4 Ablation Studies

We verify the effectiveness of each component in the KGR³ framework by answering the following research questions (RQs). Table 3 and 4 shows the experimental results for ablation studies.

RQ1: Does the reasoning module improve the final inference performance? To address this question, we simply apply the re-ranking module to candidate answers retrieved from the base KGC models. The performance drop in the “KGR³ w/o Reasoning” variant demonstrates that the pre-trained LLM can provide plausible answers that are

not initially ranked at the top positions by the base models. This capability effectively breaks through the limitations of base KGC models and increases the performance ceiling of our method.

RQ2: Whether entity descriptions contribute to enhancing KG completion? In the “KGR³ w/o Entity Descriptions” variant, we remove the descriptions for the known entity of a query triple and each of its candidate answers, resulting in sub-optimal experiment results. The performance decline re-confirms our hypothesis that LLMs may lack a fundamental understanding of certain entities within the KG, showing the importance of retrieving and leveraging textual contexts.

RQ3: Can LLMs generate desirable KGC results without the help of KG triples? We observe a significant performance decline with the “KGR³ w/o Neighbor Facts” variant, particularly for Hits@1. In this case, the LLM can only rely on limited semantic knowledge derived from entity descriptions and its inherent knowledge base, which

Settings	FB15k237				WN18RR			
	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
KGR³ (Llama3-8B)	0.535	0.475	0.564	0.635	0.717	0.656	0.759	0.809
w/o. Reasoning	0.531	0.472	0.559	0.629	0.674	0.596	0.725	0.804
w/o. Entity Descriptions	0.523	0.460	0.554	0.627	0.675	0.601	0.723	0.801
w/o. Neighbor Facts	0.405	0.295	0.467	0.590	0.646	0.519	0.753	0.807

Table 3: Results for ablation studies with the removal of reasoning module, entity descriptions, or neighbor facts. We adopt NBF-Net and SimKGC as the base KGC models for the FB15k237 and WN18RR datasets, respectively.

LLMs		FB15k237				WN18RR			
Re-ranking	Reasoning	MRR	Hits@1	Hits@3	Hits@10	MRR	Hits@1	Hits@3	Hits@10
Llama2-7B	Llama2-7B	0.524	0.462	0.555	0.627	0.709	0.645	0.754	0.803
Llama2-7B	GPT3.5	0.530	0.466	0.562	0.635	0.710	0.644	0.757	0.808
Llama3-8B	Llama3-8B	0.535	0.475	0.564	0.634	<u>0.717</u>	<u>0.656</u>	0.759	0.809
Llama3-8B	GPT3.5	0.536	0.477	0.565	0.636	<u>0.717</u>	0.655	0.759	0.807
Qwen2-1.5B	Qwen2-1.5B	0.526	0.465	0.555	0.627	0.706	0.641	0.751	0.803
Qwen2-1.5B	GPT3.5	0.531	0.469	0.560	0.648	0.706	0.637	0.753	0.806
Qwen2-7B	Qwen2-7B	<u>0.539</u>	<u>0.482</u>	<u>0.566</u>	0.634	0.724	0.672	<u>0.754</u>	0.805
Qwen2-7B	GPT3.5	0.543	0.487	0.570	<u>0.637</u>	0.724	0.672	<u>0.754</u>	0.807

Table 4: Results for ablation studies with different combinations of LLMs in the Reasoning and Re-ranking stages. We adopt the best base KGC model, namely NBF-Net and SimKGC for the FB15k237 and the WN18RR datasets.

proves insufficient for generating precise predictions. This underscores that KG triples provide accurate and irreplaceable structural knowledge that is not inherently present in the LLM. In general, experimental results in Table 3 reconfirm the complementary relationship between textual and structural contexts.

Case study. In addition, we conduct a case study on the hallucination case that may happen during the reasoning stage. For example, when being instructed to complete the incomplete triple (`<missing-entity>, /olympics/olympic_games/participating_countries, /m/04vjh` [Mauritania]) which queries the Olympic Games in which Mauritania has participated, the pre-trained LLM (GPT) incorrectly responds “*Mauritania has never participated in the Olympic Games*”. This error highlights how hallucinations can arise from the model’s lack of critical factual knowledge.

Nevertheless, triple (`/m/04vjh` [Mauritania], `/olympics/olympic_participating_country/athletes - /olympics/olympic_athlete_affiliation/Olympics, /m/06sks6` [2012 Summer Olympics]) shows that Mauritania does attend 2012 Summer Olympic games, which can help the LLM to figure out Mauritania also participated in [2008 Summer Olympics] during the re-ranking stage. This case study also shows that access to accurate factual knowledge is essential for LLMs to perform reasoning tasks like KGC.

Settings	MRR	Hits@1	Hits@3	Hits@10
Reordering (RotatE)	0.382	0.293	0.417	0.559
KGR ³ (RotatE)	0.456	0.400	0.476	0.569
Reordering (GIE)	0.391	0.301	0.426	0.573
KGR ³ (GIE)	0.463	0.400	0.485	0.581

Table 5: Ablation Experiments on FB15k-237 dataset with different re-ranking strategies.

RQ4: Can KGR³ ensure desirable performance with different LLMs? Experimental results in Table 4 show that KGR³ consistently outperforms all baseline methods across all 8 LLM combinations. Notably, KGR³ achieves the state-of-the-art performance when using Qwen2-7B and GPT3.5 as the LLMs of the re-ranking and the reasoning module. Compared to variants using pre-trained open-source LLMs, those employing GPT3.5 in reasoning stage produce better predictions, showcasing its stronger semantic understanding and instruction following capabilities.

It should be noted that KGR³, when using the same LLM, Llama2-7B, still outperforms DIFT by a large margin. Hence, we cannot simply attribute the performance improvements to the power of LLMs. In addition, KGR³ can also produce plausible predictions with an 1.5B model, which reduces the average SFT time from 28.3h to 9.05h, demonstrating its strong robustness and effectiveness in low-resource settings.

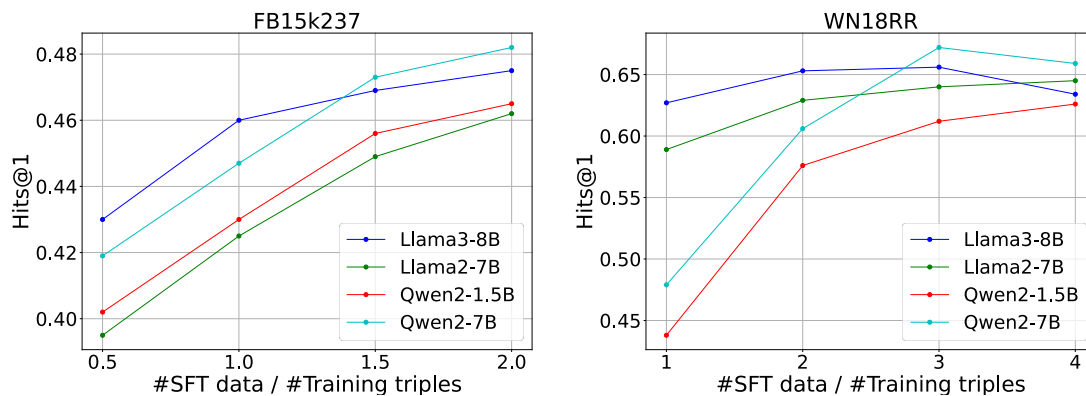


Figure 3: Hits@1 performance over the amount of SFT data on the two datasets with different LLMs.

RQ5: How does the ranking strategy affect the performance of KGR³? Rather than instructing the LLM to “sort” or “reorder” the entire candidate entity list, KGR³ guides the LLM to identify and place the most likely candidate at the top, while preserving the relative order of the remaining candidates. As shown in Table 5, we conduct additional experiments to compare these two ranking strategies. Under the reorder setting, the LLM is fine-tuned to follow the original ranking produced by the corresponding base KGC models. However, the results remain suboptimal, suggesting that LLMs are not well-suited for sorting entire entity lists. This limitation can be attributed to the nature of KGs, which do not inherently provide ground truth rankings of entities for each query triple. As a result, it is not intuitive to claim that ground truth tail entity A is more relevant than ground truth tail entity B for a given query $(h, r, ?)$. Additionally, the experimental results demonstrate that constraining LLMs to mimic the ranking patterns of prior KGC methods does not alleviate their inherent limitations.

Furthermore, in some cases, LLMs may select an incorrect, or a false negative candidate answer, which reduces the rankings of the ground truth entity. For example, during the re-ranking stage, the Llama3-8B model suggests “Solihull” as the answer of query triple $(England, location\ contains, <missing-entity>)$, where the ground truth is “Pontefract”. In fact, both “Solihull” and “Pontefract” are towns in England. Due to the inherent defect of the KG, the former is judged as a wrong answer. The re-ranking strategy adopted in KGR³ ensures that when an LLM makes an incorrect judgment, the ranking of the ground truth answer drops by no more than 1 position, and hence, results in desirable performance improvements.

RQ6: How does the volume of SFT data affect the performance of KGR³? From Figure 3, we can conclude that increasing SFT data generally improves the model performance. We attribute the boundary effects observed in the WN18RR dataset to the characteristics of its entities, which consist of common English words that can be readily interpreted by LLMs. It is also noteworthy that, even with a limited number of training samples, KGR³ still achieves desirable results. Given the substantial performance improvements, taking additional computational costs is deemed justifiable.

6 Conclusion

In this paper, we propose KGR³, an LLM-based context-enriched KGC framework with three modules: Retrieval, Reasoning, and Re-ranking. By leveraging contextual information, KGR³ effectively bridges the semantic gap between structural KG triples and natural language. Experimental results show that incorporating supporting triples and entity descriptions with LLM in-context learning and SFT significantly improves the KGC performance. Future work will focus on adapting LLMs to other KG reasoning tasks such as inductive KGC and knowledge-based question answering.

Acknowledgement

The work described in this paper was partially supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK 14222922, RGC GRF 2151185). We would also like to express our sincere gratitude to the reviewers and area chairs for their constructive comments and suggestions.

Limitations

Although the proposed framework achieves a significant breakthrough in the KGC task, it still has some remaining issues to be resolved in the future. Firstly, the proposed KGR³ framework is not capable of handling the KGC task under an “inductive setting”. The KGC task discussed in this paper, along with most related works, operates under a “transductive setting”, where entities in test triples also exist in the training set. We plan to tackle unseen entities that are not present in the KG in the future. Secondly, the commonly adopted evaluation metrics Hits@*k* necessitate ranking the plausibility of all entities within the KG for each query triple. Due to the limitations imposed by the maximum sequence length of LLMs, it is impractical to rank tens of thousands of entities within a KG. To *meet the evaluation criteria*, one may have to select a proper subset of entity candidates using existing KGC models. Hence, the base KGC model may inevitably constrain the final performance of the proposed framework.

Ethics Statement

To the best of our knowledge, this work does not involve any discrimination, social bias, or private data. All the datasets are constructed from open-source KGs such as Wikidata. Therefore, we believe that our study complies with the ACL Ethics Policy.

References

- Ivana Balazevic, Carl Allen, and Timothy Hospedales. 2019. [Tucker: Tensor factorization for knowledge graph completion](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5185–5194, Hong Kong, China. Association for Computational Linguistics.
- Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. 2008. [Freebase: A collaboratively created graph database for structuring human knowledge](#). In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data, SIGMOD '08*, page 1247–1250, New York, NY, USA. Association for Computing Machinery.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems 26*, volume 26. Curran Associates, Inc.
- Zongsheng Cao, Qianqian Xu, Zhiyong Yang, Xiaochun Cao, and Qingming Huang. 2022. [Geometry interaction knowledge graph embeddings](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(5):5521–5529.
- Ines Chami, Adva Wolf, Da-Cheng Juan, Frederic Sala, Sujith Ravi, and Christopher Ré. 2020. [Low-dimensional hyperbolic knowledge graph embeddings](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6901–6914, Online. Association for Computational Linguistics.
- Wenshuo Chao, Zhi Zheng, Hengshu Zhu, and Hao Liu. 2024. [Make large language model a better ranker](#). *Preprint*, arXiv:2403.19181.
- Yankai Chen, Menglin Yang, Yingxue Zhang, Mengchen Zhao, Ziqiao Meng, Jianye Hao, and Irwin King. 2022a. Modeling scale-free graphs with hyperbolic geometry for knowledge-aware recommendation. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM)*, pages 94–102.
- Yankai Chen, Yaming Yang, Yujing Wang, Jing Bai, Xiangchen Song, and Irwin King. 2022b. Attentive knowledge-aware graph convolutional networks with collaborative guidance for personalized recommendation. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 299–311. IEEE.
- Bonggeun Choi, Daesik Jang, and Youngjoong Ko. 2021. [Mem-kgc: Masked entity model for knowledge graph completion with pre-trained language model](#). *IEEE Access*, 9:132025–132032.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Xiou Ge, Yun Cheng Wang, Bin Wang, and C.-C. Jay Kuo. 2023. [Compounding geometric operations for knowledge graph completion](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6947–6965, Toronto, Canada. Association for Computational Linguistics.
- Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, and Christian Bizer. 2014. [Dbpedia - a large-scale, multilingual knowledge base extracted from wikipedia](#). *Semantic Web Journal*, 6.
- Muzhi Li, Minda Hu, Irwin King, and Ho fung Leung. 2024. [The integration of semantic and structural knowledge in knowledge graph entity typing](#). *Preprint*, arXiv:2404.08313.

- Yang Liu, Zequn Sun, Guangyao Li, and Wei Hu. 2022. [I know what you do not know: Knowledge graph embedding via co-distillation learning](#). In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, CIKM '22*, page 1329–1338, New York, NY, USA. Association for Computing Machinery.
- Yang Liu, Xiaobin Tian, Zequn Sun, and Wei Hu. 2024. [Finetuning generative large language models with discrimination instructions for knowledge graph completion](#). *Preprint*, arXiv:2407.16127.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Justin Lovelace, Denis Newman-Griffis, Shikhar Vashishth, Jill Fain Lehman, and Carolyn Rosé. 2021. [Robust knowledge graph completion with stacked convolutions and a student re-ranking network](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1016–1029, Online. Association for Computational Linguistics.
- Yueen Ma, Zixing Song, Xuming Hu, Jingjing Li, Yifei Zhang, and Irwin King. 2023. [Graph component contrastive learning for concept relatedness estimation](#). In *AAAI*, pages 13362–13370. AAAI Press.
- George A. Miller. 1995. [Wordnet: a lexical database for english](#). *Commun. ACM*, 38(11):39–41.
- Deepak Nathani, Jatin Chauhan, Charu Sharma, and Manohar Kaul. 2019. [Learning attention-based embeddings for relation prediction in knowledge graphs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4710–4723, Florence, Italy. Association for Computational Linguistics.
- Michael Schlichtkrull, Thomas Kipf, Peter Bloem, Rianne Berg, Ivan Titov, and Max Welling. 2018. [Modeling Relational Data with Graph Convolutional Networks](#), pages 593–607.
- Bin Shang, Yinliang Zhao, Jun Liu, and Di Wang. 2024. [Mixed geometry message and trainable convolutional attention network for knowledge graph completion](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(8):8966–8974.
- Zixing Song, Yueen Ma, and Irwin King. 2022. [Individual fairness in dynamic financial networks](#). In *NeurIPS 2022 Workshop: New Frontiers in Graph Learning*.
- Zixing Song, Ziqiao Meng, and Irwin King. 2024. [A diffusion-based pre-training framework for crystal property prediction](#). In *AAAI*, pages 8993–9001. AAAI Press.
- Zixing Song, Yifei Zhang, and Irwin King. 2023a. [No change, no gain: Empowering graph neural networks with expected model change maximization for active learning](#). In *NeurIPS*.
- Zixing Song, Yifei Zhang, and Irwin King. 2023b. [Optimal block-wise asymmetric graph construction for graph-based semi-supervised learning](#). In *NeurIPS*.
- Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. 2007. [Yago: A core of semantic knowledge](#). In *Proceedings of the 16th International Conference on World Wide Web, WWW '07*, page 697–706, New York, NY, USA. Association for Computing Machinery.
- Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M. Ni, Heung-Yeung Shum, and Jian Guo. 2024. [Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph](#). *Preprint*, arXiv:2307.07697.
- Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2019. [Rotate: Knowledge graph embedding by relational rotation in complex space](#). In *International Conference on Learning Representations*.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Eric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2071–2080, New York, New York, USA. PMLR.
- Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha Talukdar. 2020. [Composition-based multi-relational graph convolutional networks](#). In *International Conference on Learning Representations*.
- Denny Vrandečić and Markus Krötzsch. 2014. [Wiki-data: A free collaborative knowledgebase](#). *Commun. ACM*, 57(10):78–85.
- Haoran Wang and Kai Shu. 2023. [Explainable claim verification via knowledge-grounded reasoning with large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6288–6304, Singapore. Association for Computational Linguistics.
- Liang Wang, Wei Zhao, Zhuoyu Wei, and Jingming Liu. 2022. [SimKGC: Simple contrastive knowledge graph completion with pre-trained language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4281–4294, Dublin, Ireland. Association for Computational Linguistics.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021. [KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation](#). *Transactions of the Association for Computational Linguistics*, 9:176–194.

- Yanbin Wei, Qiushi Huang, Yu Zhang, and James Kwok. 2023. [KICGPT: Large language model with knowledge in context for knowledge graph completion](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 8667–8683, Singapore. Association for Computational Linguistics.
- Hong Wu, Zhe Wang, Kewen Wang, Pouya Ghiasnejad Omran, and Jiangmeng Li. 2023. [Rule Learning over Knowledge Graphs: A Review](#). *Transactions on Graph Data and Knowledge*, 1(1):7:1–7:23.
- 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 *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. [Kgbert: Bert for knowledge graph completion](#).
- Yongqi Zhang and Quanming Yao. 2022. [Knowledge graph reasoning with relational digraph](#). In *Proceedings of the ACM Web Conference 2022, WWW '22*, page 912–924, New York, NY, USA. Association for Computing Machinery.
- Yunxiang Zhao, Jianzhong Qi, Qingwei Liu, and Rui Zhang. 2021. [Wgcn: Graph convolutional networks with weighted structural features](#). In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21*, page 624–633, New York, NY, USA. Association for Computing Machinery.
- Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyang Luo, and Yongqiang Ma. 2024. [Llamafactory: Unified efficient fine-tuning of 100+ language models](#). *arXiv preprint arXiv:2403.13372*.
- Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. 2020. [Graph neural networks: A review of methods and applications](#). *AI Open*, 1:57–81.
- Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhi Yao, Shumin Deng, Huajun Chen, and Ningyu Zhang. 2023a. [Llms for knowledge graph construction and reasoning: Recent capabilities and future opportunities](#). *arXiv preprint arXiv:2305.13168*.
- Zhaocheng Zhu, Xinyu Yuan, Michael Galkin, Louis-Pascal Xhonneux, Ming Zhang, Maxime Gazeau, and Jian Tang. 2023b. [Aast net: A scalable path-based reasoning approach for knowledge graphs](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 59323–59336. Curran Associates, Inc.
- Zhaocheng Zhu, Zuobai Zhang, Louis-Pascal Xhonneux, and Jian Tang. 2021. [Neural bellman-ford networks: A general graph neural network framework for link prediction](#). In *Advances in Neural*

A Appendix

A.1 Prompt templates of Reasoning stage

Table 6 shows the prompt templates of the **Reasoning** stage and provides an example.

In-context Demonstrations:

User: **Washington County: county in Pennsylvania, U.S.** The question is to predict the head entity [MASK] from the given ([MASK], *location adjoining_relationship adjoins*, Washington County) by completing the sentence 'Washington County is the adjoins of what location? The answer is '.

Assistant: The answer is Westmoreland County, so the [MASK] is Westmoreland County. **Westmoreland County: county in Pennsylvania, United States**

User: **Rockland County: suburban county in New York, United States** The question is to predict the head entity [MASK] from the given ([MASK], *location adjoining_relationship adjoins*, Rockland County) by completing the sentence 'Rockland County is the adjoins of what location? The answer is '.

Assistant: The answer is Bergen County, so the [MASK] is Bergen County. **Bergen County: county in New Jersey, United States**

Query triple: ([MASK], */location/adjoining_relationship/adjoins*, Champaign)

Reasoning prompts:

User: The question is to predict the head entity [MASK] from the given ([MASK], *location adjoining_relationship adjoins*, Champaign) by completing the sentence 'Champaign is the adjoins of what location? The answer is '. Output all some possible answers based on your own knowledge, using the format '[answer1, answer2, ..., answerN]' and please start your response with 'The possible answers:'. Do not output anything except the possible answers.

Context-aware Reasoning:

User: Here are some materials for you to refer to:

Champaign: Champaign is a city in Champaign County, Illinois, United States. The population was 88,302 at the 2020 census. It is the tenth-most populous municipality in Illinois and the fourth most populous city in the state outside the Chicago metropolitan area. It is a principal city of the Champaign–Urbana metropolitan area, which had 236,000 residents in 2020.

The question is to predict the head entity [MASK] from the given ([MASK], *location adjoining_relationship adjoins*, Champaign) by completing the sentence 'Champaign is the adjoins of what location? The answer is '. Output all the possible answers you can find in the materials using the format '[answer1, answer2, ..., answerN]' and please start your response with 'The possible answers:'. Do not output anything except the possible answers. If you cannot find any answer, please output some possible answers based on your own knowledge.

Reasoning results:

Assistant: The possible answers: [Urbana, Champaign County, Illinois Silicon Prairie, Parkland College]

Table 6: Prompt Template of context-aware reasoning.

A.2 Prompt templates of Re-Ranking stage

Table 7 shows the prompt templates of the **Re-ranking** stage and gives an example. It is noteworthy that this case empirically shows the effectiveness of the **Reasoning** and **Re-ranking** processes. The ground truth answer “Urbana” does not rank in a leading position by the KGC model. However, the LLM provides plausible candidates including the ground truth answer “Urbana”, by analyzing the context of the known entity “Champaign” in the incomplete triple during the **Reasoning** process. In addition, the fine-tuned LLM succeeds in selecting the correct answer from the candidate list based on relevant textual contexts and neighbor facts during the **Re-ranking** process.

Query triple: (< ??? >, /location/adjoining_relationship/adjoins, Champaign)

Re-Ranking prompts:

User: Here is an incomplete triple with missing head entity <missing-entity>: (<missing-entity>, location location adjoin_s. location adjoining_relationship adjoins, Champaign).

Following are some contexts about tail entity Champaign:

Champaign is a city in Champaign County, Illinois, United States. The population was 88,302 at the 2020 census. It is the tenth-most populous municipality in Illinois and the fourth most populous city in the state outside the Chicago metropolitan area. It is a principal city of the Champaign-Urbana metropolitan area, which had 236,000 residents in 2020.

Following are some triple facts of entity Champaign:

(Illinois, location location contains, Champaign)

(Ludacris, people person place of birth, Champaign)

(Champaign County, location location contains, Champaign)

Please select the most appropriate entity for <missing-entity> from the candidate answer list:

Champaign County: county in Illinois, United States,

McHenry County: county in Illinois, United States,

Lake County: county in Illinois, United States,

Cook County: county in Illinois, United States,

Madison County: county in Illinois, United States,

St. Clair County: county in Illinois, United States,

DuPage County: county in Illinois, United States,

McLean County: county in Illinois, United States,

Champaign: city in Champaign County, Illinois, United States,

Peoria County: county in Illinois, United States of America,

Kane County: county in Illinois, United States,

Tazewell County: county in Illinois, United States,

Oak Park: village in Cook County, Illinois, United States; suburb of Chicago, Illinois,

Will County: county in Illinois, United States,

Lake Forest: city in Lake County, Illinois, United States

Springfield: city in and county seat of Sangamon County and Illinois federated state capital city

Aurora: city in Kane County, Illinois, United States

Urbana: town in and county seat of Champaign County, Illinois, United States

Alton: city in Madison County, Illinois, United States

Kankakee: city in Kankakee County, Illinois, United States.

Re-Ranking Results:

Assistant: Urbana

Evaluation: The ground truth “Urbana” hits at 1.

Table 7: Prompt Template of context-aware ranking.

B Implementation Details

We conduct all our experiments on a Linux server equipped with 2 Intel Xeon Platinum 8358 processors and 8 A100-SXM4-40GB GPUs. [†] For TransE and CoLE, we adopt the candidate entities provided in (Liu et al., 2024). For other base KGC methods, candidate entities are retrieved using their publicly available implementations. We tested our proposed method using 4 different LLMs, specifically Llama2-7B, Llama3-8B, Qwen2-1.5B, and Qwen2-7B. The time costs of SFT and re-ranking are reported in Table 9. It is noteworthy that the duration of the experiments is subject to specific configurations, which can be further reduced through parallel inference or the application of vLLM service. Additionally, we also evaluated OpenAI’s gpt-3.5-turbo-0125 model in the reasoning stage. During the re-ranking stage, we employ the LLaMA-Factory framework (Zheng et al., 2024) to fine-tune the mentioned LLMs. For LoRA adjustments, we set the rank to 16 and the alpha value to 32. We utilize the AdamW optimizer (Loshchilov and Hutter, 2017) to train our model, with an initial learning rate of $1.0e-4$, a per-device batch size of 2 and a gradient accumulation step of 4 iterations. We adopted BF16 precision to reduce the GPU memory usage. For each query triple, we retrieve 3 supporting triples for in-context demonstration and provide $n = 20$ candidate answers (during training, we provide 1 ground truth entity and 19 negative samples) to the LLM. Detailed hyperparameter settings are listed in Table 8.

Hyperparameters	Settings
SFT learning rate	1e-4
Per-device batch size	2
# GPUs used	2
Gradient accumulation step	4
LoRA rank	16
LoRA α value	32
Re-ranking scope n	20
# Easy negative samples	9
# Hard negative samples	10
δ	50
p	10

Table 8: Hyperparameter settings.

Dataset	Qwen2-1.5B	Qwen2-7B	Llama2-7B	Llama3-8B	Test (Llama3-8B)
FB15k237	12h 12min	36h 50min	43h 42min	38h 27min	0.180s / query
WN18RR	5h 54min	22h 33min	19h 30min	18h 12min	0.152s / query

Table 9: Time costs of supervised fine-tuning and re-ranking.

[†]Only 2 GPUs are used in our experiments.