

# Harnessing Large Language Models for Knowledge Graph Question Answering via Adaptive Multi-Aspect Retrieval-Augmentation

Derong Xu<sup>1,2</sup>, Xinhang Li<sup>1</sup>, Ziheng Zhang<sup>3</sup>, Zhenxi Lin<sup>3</sup>, Zhihong Zhu<sup>4</sup>, Zhi Zheng<sup>1</sup>, Xian Wu<sup>3\*</sup>,  
Xiangyu Zhao<sup>2\*</sup>, Tong Xu<sup>1\*</sup>, Enhong Chen<sup>1</sup>

<sup>1</sup>University of Science and Technology of China & State Key Laboratory of Cognitive Intelligence,

<sup>2</sup>City University of Hong Kong,

<sup>3</sup>Jarvis Research Center, Tencent YouTu Lab,

<sup>4</sup>Peking University

{derongxu,xinhangli,zhengzhi97}@mail.ustc.edu.cn, {zihengzhang, chalerislin, kevinxwu}@tencent.com,  
xianzhao@cityu.edu.hk, {tongxu, cheneh}@ustc.edu.cn, zhihongzhu@stu.pku.edu.cn

## Abstract

Large Language Models (LLMs) demonstrate remarkable capabilities, yet struggle with hallucination and outdated knowledge when tasked with complex knowledge reasoning, resulting in factually incorrect outputs. Previous studies have attempted to mitigate it by retrieving factual knowledge from large-scale knowledge graphs (KGs) to assist LLMs in logical reasoning and prediction of answers. However, this kind of approach often introduces noise and irrelevant data, especially in situations with extensive context from multiple knowledge aspects. In this way, LLM attention can be potentially misled from question and relevant information. In our study, we introduce an Adaptive Multi-Aspect Retrieval-augmented over KGs (Amar) framework. This method retrieves knowledge including entities, relations, and subgraphs, and converts each piece of retrieved text into prompt embeddings. The Amar framework comprises two key sub-components: 1) a self-alignment module that aligns commonalities among entities, relations, and subgraphs to enhance retrieved text, thereby reducing noise interference; 2) a relevance gating module that employs a soft gate to learn the relevance score between question and multi-aspect retrieved data, to determine which information should be used to enhance LLMs' output, or even filtered altogether. Our method has achieved state-of-the-art performance on two common datasets, WebQSP and CWQ, showing a 1.9% improvement in accuracy over its best competitor and a 6.6% improvement in logical form generation over a method that directly uses retrieved text as context prompts. These results demonstrate the effectiveness of Amar in improving the reasoning of LLMs.

**Code** — <https://github.com/Applied-Machine-Learning-Lab/AMAR>

## Introduction

Recently, large language models (LLMs) like GPT-4 (OpenAI 2024) and Llama (Touvron et al. 2023) have shown impressive performance improvements across a variety of natural language processing (NLP) tasks (WANG et al. 2024;

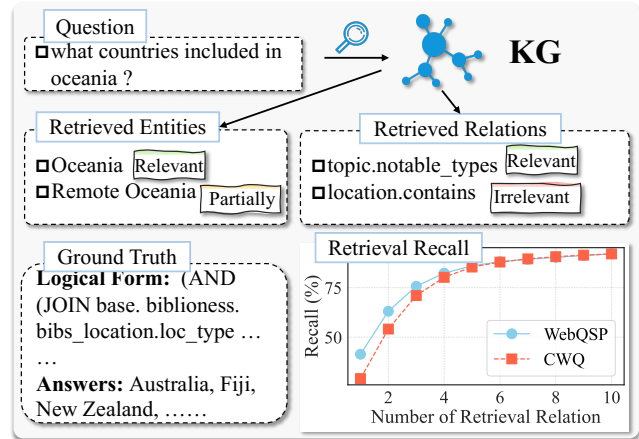


Figure 1: An example illustrates the retrieval results categorized as *relevant*, *irrelevant*, and *partially relevant*.

WEN et al. 2023; XU et al. 2024; ZENG et al. 2023). However, when dealing with specialized knowledge not in training corpus and complex knowledge reasoning, LLMs still struggle with outdated knowledge and hallucination problems (Ji et al. 2023). They thus may produce factually incorrect outputs, limiting their usefulness in the areas requiring high reliability, such as healthcare (He et al. 2023) and safety (Dong et al. 2024). To solve complex reasoning with specialized knowledge, a line of research (Luo et al. 2023; Sun et al. 2024) have explored Knowledge Graph Question Answering (KGQA), a task that improves logical reasoning and prediction of answers by retrieving reliable information from large-scale knowledge graphs (KGs) like Freebase (Bollacker et al. 2008) and Wikidata (Vrandečić and Krötzsch 2014). KGs store extensive factual knowledge in a structured format known as triplets, consisting of (*head entity*, *relation*, *tail entity*), which is seen as a potential solution for enhancing the interpretability of LLMs reasoning (Sun et al. 2024).

The recent advancements in KGQA can be broadly classified into two main categories. The first category involves using LLMs to transform input questions into structured log-

\*Corresponding authors.

ical forms (Luo et al. 2023), such as S-expressions, which can then be queried on KG using a graph database query language like SPARQL to obtain the final answers. These approaches leverage the structured nature of KG to extract precise information in response to complex queries. The second category of KGQA involves directly predicting answers using LLMs without extra logical forms. Some works in this category perform multi-hop reasoning on KG in a step-by-step manner by querying LLMs (Sun et al. 2024). Other works retrieve additional contextual information related to the question and use it as contextual input to enhance LLMs (He et al. 2024; LUO et al. 2024). These approaches aim to improve the accuracy of answers by incorporating more context into the reasoning process.

However, the retrieval process often yields a large amount of information, as shown in Figure 1, which may be irrelevant or only partially relevant to the question and ground truth. The retrieval recall curve demonstrates a positive correlation between the number of retrieved data and the recall score, indicating that the retrieved data contain valuable information. Therefore, it raises a question: *How to identify which retrieved knowledge is valuable and which is not?* This question becomes particularly crucial when a significant portion of retrieved knowledge does not provide substantial assistance. In terms of this, we observed that previous studies (Yu et al. 2022; Hu et al. 2022) failed to considering commonalities among different aspects of retrieved information, which can be beneficial in identifying crucial knowledge. And they also neglected adaptive learning to establish relevance between question and retrieved text.

To address these overlooked challenges, we propose an adaptive multi-aspect retrieval-augmented over KG (AMAR) framework. Instead of directly appending the retrieved information as context input, AMAR utilizes the retrieved information more flexibly, therefore enhancing LLMs reasoning. AMAR primarily includes two modules: 1) Self-alignment module, where multi-aspect retrieval data (including entities, relations, and subgraphs, all of which are linearized as text) is separately mapped into prompt embeddings from text, facilitating fine-grained tuning. Next, cross-attention and self-attention are applied to the multi-aspect embeddings to obtain consistency tokens. The aim is to align the commonalities among different pieces of information. For instance, if an entity and a subgraph both mention “*Oceania*”, this implies that they are probably consistent in conveying the same piece of information. In this way, we can enhance crucial knowledge, thereby reducing noise interference. 2) Relevance gating module, which measures the relevance between the embedding of question and consistency tokens. This module introduces siamese networks with shared weights to learn relevance score, which serves as a soft gate to adaptively decide which retrieved information is more useful for LLMs reasoning. Through the self-alignment and relevance gating modules, AMAR adaptively filters and selects multi-aspect retrieval knowledge, enabling a more rational utilization of context and avoiding interference from noise. We then fine-tune LLM in a parameter-efficient manner to leverage the retrieved knowledge, assisting in generating reasonable logical forms. Finally, we refine

the logical forms using the similarity of entities and relations and query KG to obtain final answers. Overall, this work makes three key contributions:

- To the best of our knowledge, this is the first study to enhance LLMs reasoning for KGQA tasks by utilizing multi-aspect KG information as prompt embeddings.
- We propose novel self-alignment and relevance gating modules, which enable LLMs to adaptively filter and select multi-aspect retrieval knowledge, allowing for more rational utilization of context while effectively avoiding interference from noise.
- The effectiveness of AMAR was extensively validated on two datasets across five metrics. AMAR demonstrated superior performance, achieving state-of-the-art (SOTA) results compared to 22 strong baselines.

## Preliminaries

We first introduce key concepts and notations for our task.

**Knowledge Graph (KG).** A KG is a collection of factual knowledge organized in form of triples:  $\mathcal{G} = \{(h, r, t)\} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$ , where  $\mathcal{E}$  represents the entity set and  $\mathcal{R}$  represents relation set. Each triple consists of three elements: a head entity  $h$ , a relation  $r$ , and a tail entity or a literal  $t$ .

**Logical Form.** The logical form is a structured language representation of a question. In this work, we adopt S-expression  $\mathcal{F}$  as our chosen logical expression, following (Luo et al. 2023; Yu et al. 2022). As shown in the examples provided in Figure 1, the S-expression utilizes functions (such as *JOIN*, *AND*) that operate on set-based semantics, which keep a balance between readability and compactness and thus is well-suited for KGQA (Gu et al. 2021).

**KGQA with LLM.** KGQA is a classical NLP task that has been further enhanced through the utilization of LLMs. Given a question  $q$ , this task aims to retrieve knowledge related to  $q$  from a KG  $\mathcal{G}$  and generate an S-expression  $\mathcal{F}$ . Since many KG storage engines support SPARQL, the generated S-expression  $\mathcal{F}$  is converted into a SPARQL query, which is further executed against  $\mathcal{G}$  to obtain the final answers. In our work, We retrieve multi-aspect knowledge (including entities  $k_e$ , relations  $k_r$ , and subgraphs  $k_s$ ) and design a model  $f$  that utilizes LLMs to generate  $\mathcal{F}$  based on question and retrieval knowledge as input, i.e.,  $\mathcal{F} = f(q, k_e, k_r, k_s)$ . By converting and executing the SPARQL query on the KG, we can obtain final answers denoted as  $a = query(convert(\mathcal{F})) \in \mathcal{A}_q$ , where  $\mathcal{A}_q \subseteq \mathcal{E}$ .

## Methodology

In this section, as shown in Figure 2, we begin with multi-aspect knowledge retrieval. We then delve into two specific modules, providing detailed explanations of their functionalities, and demonstrate the process of querying KG using logical form.

### Retrieval of Multi-Aspect Knowledge

Previous KGQA methods either focus on retrieving entities and relations (Hu et al. 2022) or retrieving multiple triplets to form subgraphs (Yu et al. 2022; He et al. 2024). In our

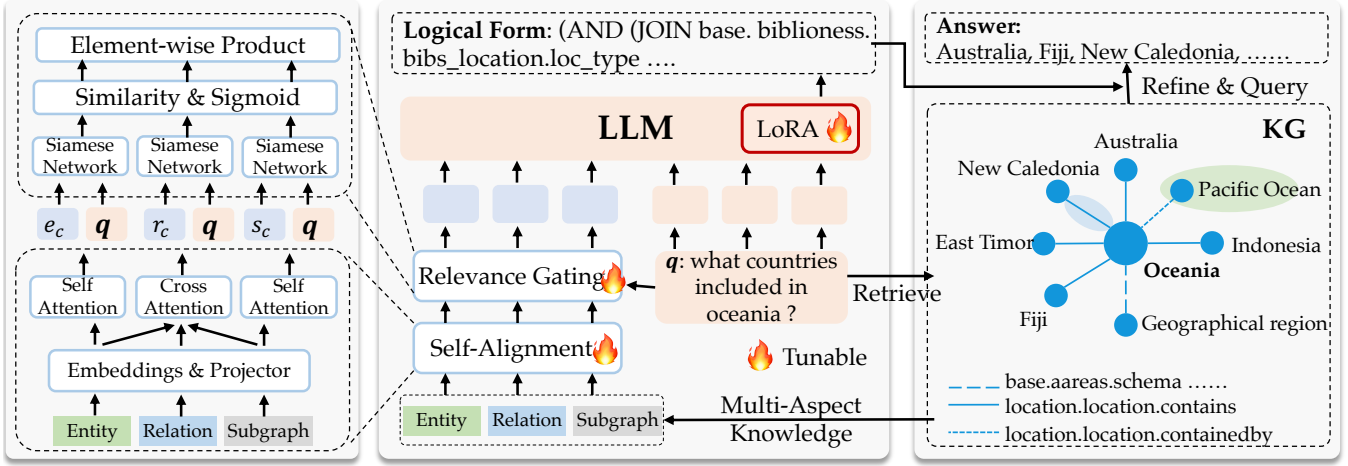


Figure 2: Overall framework of AMAR . We retrieve multi-aspect knowledge from KG and obtain weighted consistency prompt embeddings using Self-Alignment and Relevance Gating modules.

research, we propose to leverage three types of knowledge simultaneously. Each possesses varying aspects of information, which complement each other and also share commonalities. This enables us to effectively align crucial knowledge across three types. Due to the differences in their representations (e.g., length and format), we employ different retrieval methods for each type.

**Entity Retrieval.** One effective approach for retrieving candidate entities  $k_e$  is to conduct entity linking with question  $q$ . Following Hu et al. (2022), we employ the ELQ (Li et al. 2020) for question entity linking, which utilizes a bi-encoder to simultaneously detect mentions and link them to entities end-to-end. Then FACCC1 (Gabrilovich, Ringgaard, and Subramanya 2013) (a comprehensive Freebase annotation of corpora) is employed to identify entities that were not linked by ELQ, to enhance the range of candidate entities.

**Relation Retrieval.** In large-scale KG (e.g. Freebase), relations are typically organized hierarchically, such as the example *base.biblioness.bibs\_location.loc\_type*. Therefore, directly using question-based dense retrieval for similarity may not be effective. To address this, following Hu et al. (2022), we propose masking entity mentions detected during the candidate entity retrieval stage with a [BLANK] token for each question  $q$ . Building on the work of Hu et al. (2022); Das et al. (2021), we train two separate BERT models that encode questions and relations into a shared dense space. The objective of optimization is to maximize the score of the relevant relation compared to randomly sampled relations. To retrieve the nearest relations, we employ FAISS (Douze et al. 2024), a highly efficient vector database, which allows us to speed up the search process and obtain the most relevant results.

**Subgraph Retrieval.** One crucial consideration is the wealth of structural and semantic information contained within KG. Since KG data is typically stored as triplets, we linearize triplets by combining head entity, relation, and tail entity for retrieval. Following (Yu et al. 2022), we pro-

pose grouping linearized sentences with the same head entity into a document. To save computing resources, we only focus on 1-hop subgraphs to capture structural information. Furthermore, concerning the potential information loss when converting long documents into vectors, we employ sparse retrieval approaches that rely on keyword dependencies. Specifically, we employ techniques like BM25, which calculates TF-IDF scores based on sparse word matches between input questions and KB-linearized passages.

### Self-Alignment

Although multi-aspect retrieval offers comprehensive auxiliary information, it can also introduce irrelevant knowledge and noise, resulting in negative impacts. To address this concern, we propose to align the commonalities among multi-aspect information to improve informativeness.

We first utilize LLM’s embedding layer to convert the multi-aspect retrieval knowledge into text embeddings  $\mathbf{X}_e \in \mathbb{R}^{t_k \times l \times e}$ , where the  $t_k$  represent top  $k$  retrieval texts,  $l$  denotes the maximum length of text, and  $e$  indicate the embedding dim of the token. We then apply an averaging operation, followed by a projector network  $\mathcal{M}$ , which transforms these embeddings into prompt embeddings  $\mathbf{X}_t \in \mathbb{R}^{t_k \times e}$  (i.e., one piece of retrieval text is projected to one token embedding). For the sake of efficiency,  $\mathcal{M}$  is designed to consist of down-projection and up-projection layers, with a nonlinear layer situated between them, as follows:

$$\begin{aligned} \mathbf{X}_e &= \text{Embeddings}(T), \\ \mathbf{X}_t &= \mathcal{M}(\hat{\mathbf{X}}_e), \quad \hat{\mathbf{X}}_e = \frac{1}{l} \sum_{i=1}^l \mathbf{X}_e[:, i, :], \end{aligned} \quad (1)$$

where  $T$  represents the retrieval text of entities, relations, or subgraphs. Next, we apply self-attention separately to the prompt embeddings of entities and relations to obtain entity-consistency tokens  $e_c$  and relation-consistency tokens  $r_c$ :

$$\begin{aligned} e_c &= \text{Self-Attn}(\mathbf{E}_t) \in \mathbb{R}^{t_k \times e}, \\ r_c &= \text{Self-Attn}(\mathbf{R}_t) \in \mathbb{R}^{t_k \times e}, \end{aligned} \quad (2)$$

Here, we use  $S_t$ ,  $E_t$ , and  $R_t$  represent prompt embeddings  $X_t$  of subgraphs, entities, and relations, respectively. This attention is applied to individual retrieval text (e.g., one sentence) rather than individual tokens to learn the correlation and consistency between different retrieval information and their importance within the entire top  $k$  retrieval data.

In addition, we further focus on leveraging subgraphs information. To determine the significance of retrieval text, we employ entities and relations as alignment factors. As shown in Figure 2, by aligning triples in a subgraph that is highly consistent to relation ‘*location.location.contains*’, we can weight important data through commonalities of multi-aspect knowledge. To obtain subgraph-consistency tokens  $s_c$ , we perform cross-attention to entity-subgraph and relation-subgraph pairs, respectively:

$$\begin{aligned} s_c^e &= \text{Cross-Attn}(S_t, E_t, E_t) \in \mathbb{R}^{t_k \times e}, \\ s_c^r &= \text{Cross-Attn}(S_t, R_t, R_t) \in \mathbb{R}^{t_k \times e}, \end{aligned} \quad (3)$$

and sum the results to get  $s_c = s_c^e + s_c^r \in \mathbb{R}^{t_k \times e}$ . The consistency tokens of entity  $e_c$ , relation  $r_c$ , and subgraph  $s_c$  contain refined knowledge aligned between multi-aspect information, enhancing the utilization of the retrieved data.

## Relevance Gating

After obtaining the consistency tokens, we expect the model to learn the relevance between retrieval data and the question. The relevance is used to construct a soft gating mechanism, which will adaptively select the relevant retrieval information to be utilized. To achieve this, we design a siamese network for each type of consistency token to measure its relevance to the question embedding  $Q_e \in \mathbb{R}^{l \times e}$ . Each siamese network consists of a shared MLP  $\mathcal{M}_{share}$  networks that processes both question embedding and consistency tokens, the generated  $q_m \in \mathbb{R}^{l \times e}$  and  $x_c \in \mathbb{R}^{t_k \times e}$  are then used to calculate similarity score  $G_{sim} \in \mathbb{R}^{t_k \times l}$  through a batch matrix multiplication. This score is subsequently averaged, and a sigmoid activation function is applied to produce the final relevance score  $g$ , as described below:

$$\begin{aligned} q_m &= \mathcal{M}_{share}(Q_e), \quad x_c = \mathcal{M}_{share}(X_c), \\ g &= \text{Sigmoid}\left(\frac{1}{l} \sum_{i=1}^l G_{sim}[:, i]\right), \quad G_{sim} = x_c \cdot q_m^T. \end{aligned} \quad (4)$$

Here  $X_c$  can be denoted as entities, relations, or subgraphs consistency tokens. The relevance score of entity  $g_e$ , relations  $g_r$  and subgraph  $g_s$  serve as soft gates, are used to model the influence of each consistency tokens by element-wise product, respectively:

$$e_c^w = g_e \circ e_c, \quad r_c^w = g_r \circ r_c, \quad s_c^w = g_s \circ s_c. \quad (5)$$

To enhance generalization ability, we introduce randomly initialized soft tokens  $p \in \mathbb{R}^{l \times e}$ , which are concatenated with the weighted consistency tokens and the question embedding. These combined embeddings are then fed into LLMs to generate logical form  $\mathcal{F}$  as follows:

$$\mathcal{F} = f_{\theta, \phi_1, \phi_2}([p; e_c^w; r_c^w; s_c^w; Q_e]), \quad (6)$$

where  $[\cdot]$  represents concatenation operation, the parameters  $\theta$  of LLMs itself are frozen. The parameters that require back-propagation optimization include our model parameters  $\phi_1$  and LoRA parameters  $\phi_2$ .

## Query Execution

Due to the long-tail distribution of fine-tuned data and the lack of specific knowledge, LLM may not strictly adhere to the contextual information provided (Luo et al. 2023). As a result, the generated logical forms may contain non-existent entities or relations. For instance, when asked ‘*where was rihanna born and raised?*’, an LLM might generate the logical form ‘*(JOIN (R people.person.place\_of\_birth) Rihanna)*’ instead of the correct spelling ‘*Rihanna*’. This discrepancy renders the logical form non-executable on KG, and the same issue can also arise with relations.

To further refine the quality of entity and relation, we employ a similarity-based approach with KG. Specifically, we utilize an unsupervised SimCSE model to measure the similarity between each entity in the generated logical form  $\mathcal{F}$  and the labels of entities in the entity set  $\mathcal{E}$ . By setting a threshold, we retain the most relevant entities  $\mathcal{E}_{sub}$ . Additionally, we query the KG to identify relations  $\mathcal{R}_{2-hop}$  that are within 2 hops of the obtained subset of entities  $\mathcal{E}_{sub}$ . Similarly, we calculate the similarity between all relations in  $\mathcal{F}$  and the identified set of relations. This refinement process allows us to generate a new list of candidate logical forms  $\mathcal{F}_{new}$  that better align with the KG. After converting to SPARQL language, it can be used to query answers from KG:  $a = \text{query}(\text{convert}(\mathcal{F}_{new}))$ .

## Experiments

### Experiment Settings

**Datasets.** Our experiments were conducted using two well-known datasets: WebQuestionsSP (WebQSP) (Yih et al. 2016) and ComplexWebQuestions (CWQ) (Talmor and Berant 2018). Both datasets contain SPARQL queries that correspond to the questions and can be executed on Freebase to obtain answers.

**Baselines.** In this study, we evaluate performance with 22 baselines, which are categorized into four groups: embedding-based (EM-based), information retrieval-based (IR-based), semantic parsing-based (SP-based), and LLM-based methods. It is important to note that some methods, such as DecAF (Yu et al. 2022), can be classified as multiple groups, specifically IR-based and SP-based. To ensure fairness, we do not include the results of using the oracle entity linking annotations setting, such as RoG (LUO et al. 2024).

**Evaluation Metrics.** We use Hits@1, F1, and Acc as primary evaluation metrics following (Luo et al. 2023). Hits@1 assesses the accuracy of top-1 predicted answer, F1 considers the coverage of all possible answers, and Acc measures the strict exact-match accuracy. We further assess the quality of generated S-expressions by employing two metrics: the extract match ratio (EM) and the match after beam search ratio (BM) with ground-truth S-expressions, for analytical experiments.

Type	Model	WebQSP			CWQ		
		F1	Hits@1	Acc	F1	Hits@1	Acc
EM-based	KV-Mem (Miller et al. 2016)	34.5	46.7	-	15.7	18.4	-
	NSM <sub>+h</sub> (He et al. 2021)	67.4	74.3	-	44.0	48.8	-
	TransferNet (Shi et al. 2021)	-	71.4	-	-	48.6	-
	KGT5 (Saxena et al. 2022)	-	56.1	-	-	36.5	-
IR-based	GraftNet (Sun et al. 2018)	60.4	66.4	-	32.7	36.8	-
	PullNet (Sun et al. 2019)	-	68.1	-	-	45.9	-
	SR+NSM (Zhang et al. 2022)	64.1	68.9	-	47.1	50.2	-
	SR+NSM+E2E (Zhang et al. 2022)	64.1	69.5	-	46.3	49.3	-
	UniKGQA (Jiang et al. 2023b)	71.0	77.0	-	49.4	50.9	-
SP-based	CBR-KBQA (Das et al. 2021)	72.8	-	69.9	70.0	-	67.1
	GMT-KBQA (Hu et al. 2022)	76.6	-	73.1	77.0	-	72.2
	UnifiedSKG (Xie, Wu, and et. al 2022)	73.9	-	-	68.8	-	-
	RnG-KBQA (Ye et al. 2022)	75.6	-	-	-	-	-
	DecAF (Yu et al. 2022)	78.8	82.1	-	-	70.4	-
	FC-KBQA (Zhang et al. 2023)	76.9	-	-	56.4	-	-
LLM-based	KD-CoT (Wang et al. 2023)	52.5	68.6	-	-	55.7	-
	Pangu (Gu, Deng, and Su 2023)	79.6	-	-	-	-	-
	StructGPT (Jiang et al. 2023a)	-	72.6	-	-	-	-
	ChatKBQA (Luo et al. 2023)	79.8	83.2	73.8	77.8	82.7	73.3
	ToG-R (GPT-4) (Sun et al. 2024)	-	82.6	-	-	69.5	-
	G-Retriever (He et al. 2024)	-	70.1	-	-	-	-
	GNN-RAG (Mavromatis et al. 2024)	73.5	82.8	-	60.4	62.8	-
	AMAR (Ours)	<b>81.2</b> ±0.15	<b>84.3</b> ±0.16	<b>75.2</b> ±0.10	<b>78.5</b> ±0.11	<b>83.1</b> ±0.09	<b>74.5</b> ±0.07

Table 1: Performance comparison of different types of KGQA methods on WebQSP and CWQ datasets. We present the Mean scores and standard deviations (mean ± std) of five experiments with different random seeds. The best result is highlighted in **bold**, and the baseline results are taken from corresponding papers.

**Implementation Details.** Following Luo et al. (2023), we fine-tune LLaMA2-7B on WebQSP and LLaMA2-13B on CWQ using LoRA. We evaluate the impact of backbones and fine-tuning methods in our subsequent experiments. During inference, we utilize beam search to generate multiple logical forms. We select the executable logical form with the highest score to obtain answers. All experiments were done on NVIDIA A6000 GPUs. We only searched the number of retrievals  $k$  with values of  $\{4, 8, 16, 32, 64, 100\}$ .

## Main Results

As observed from Table 1, AMAR outperforms all baselines across all metrics on both datasets. Notably, on WebQSP dataset, accuracy has improved by 1.6% compared to the second-best baseline, ChatKBQA, marking new state-of-the-art performance. Specifically, AMAR surpasses sub-graph retrieval techniques such as SR+NSM and DecAF with 26% and 2.4% F1 improvements on WebQSP, respectively, as well as entity and relation retrieval methods like GMT-KBQA with 5.3% F1 improvement on WebQSP. This can be attributed to our proposed self-alignment mechanism, which effectively aligns multi-aspect knowledge. On the other hand, AMAR also outperforms other LLM-based approaches, such as G-Retriever and ChatKBQA, suggesting that our approach of learning prompt embeddings for retrieval can more flexibly leverage the capabilities of LLMs to utilize retrieval knowledge.

Model	WebQSP				
	F1	Hits@1	Acc	EM	BM
AMAR	<b>81.2</b>	<b>84.3</b>	<b>75.2</b>	<b>63.9</b>	76.4
w/o <i>SN</i>	80.1	83.3	74.3	63.6	<b>77.0</b>
w/o <i>SA</i>	79.5	82.6	73.8	63.0	75.7
w/o <i>RG</i>	79.4	82.2	74.3	63.4	76.7
w/o <i>SA&amp;RG</i>	78.7	81.0	72.5	62.1	73.6
w/o <i>SA&amp;SN</i>	79.1	82.0	73.8	63.1	76.2
w/o <i>ALL</i>	76.2	79.5	70.1	59.7	72.4

Table 2: Ablation study of sub-modules on WebQSP dataset.

## Ablation Study

In this section, we conduct a series of ablation studies to address the following question:

### How do the proposed modules improve performance?

Specifically, we conduct experiments against five variants: 1) w/o *SN*: without siamese network, and relevance score is calculated by vector inner product; 2) w/o *SA*: without self-alignment module; 3) w/o *RG*: without relevance gating module; 4) w/o *SA&RG*: without both self-alignment and relevance gating modules, where AMAR obtains prompt embeddings with only MLP. 5) w/o *ALL*: directly appending the retrieval knowledge as context instead of converting retrieval knowledge to prompt embeddings. As shown in Table 2, we observe that the performance on most metrics decreases when either *SN*, *SA*, or *RG* is removed. This validates the effectiveness of the proposed sub-modules. Fur-

Model	WebQSP				
	F1	Hits@1	Acc	EM	BM
AMAR	<b>81.2</b>	<b>84.3</b>	<b>75.2</b>	<b>63.9</b>	<b>76.4</b>
w/o Relation	78.6	81.7	73.1	63.0	74.9
w/o Entity	79.2	82.2	73.5	63.8	74.4
w/o Subgraph	79.7	82.9	73.8	63.5	75.1

Table 3: Quantitative comparison of the impacts of retrieval information on AMAR’s performance.

thermore, we find that performance significantly drops when the entire framework is removed, indicating that appending retrieved knowledge directly as context text introduces a large amount of noise, preventing LLMs from focusing on learning the mapping from question to logical form. Additionally, we notice that the BM is higher in AMAR w/o *SN* than in AMAR. This suggests that the ground-truth logical forms are mostly ranked within the top 2 or lower positions during beam search generation, leading to a lower EM.

**What impacts do different aspects of retrieval information have on performance?** To preserve the integrity of AMAR, we remove retrieval knowledge by replacing the text embedding with randomly initialized ones. As shown in Table 3, the results indicate obvious performance drops after removing retrieval knowledge, including ‘*subgraph*’, ‘*entity*’, and ‘*relation*’. This decline highlights the significance of different aspects of the retrieval knowledge on the overall performance. Further analysis reveals that removing the ‘*relation*’ component results in the largest drop in performance, suggesting that ‘*relation*’ plays a crucial role in generating logical expressions. Instead, while the ‘*subgraph*’ still contributes to performance, it appears to be less critical for logical forms than ‘*relation*’ or ‘*entity*’. These findings provide valuable insights for further optimization of the model.

### Number of Retrieval Analysis

In this section, we explore the impact of the number of retrieval knowledge. We compare the approach of directly inputting retrieved data as a *context prompt*. If the input exceeds the maximum context limit (i.e., 4096 for LLaMA2), we truncate the retrieved information from subgraphs. As shown in figure 3, it can be observed that when the amount of retrieved data is relatively small, our method does not significantly differ from *Context Prompt*, which suggests that useful information recalled is still limited. However, as the quantity of retrieved data increases (e.g., reaching 64 or 100), our method achieves a substantial performance improvement, while *context prompt* drastically declines. This demonstrates that introducing a long context results in substantial noise, making it difficult for LLMs to learn important data. In contrast, by treating retrieved information as an individual prompt embedding, we avoid the issue of excessively long inputs and better utilize the rich information.

### Efficiency of Fine-Tuning

In this section, we analyze the efficiency of our method combined with different fine-tuning methods, including Prompt

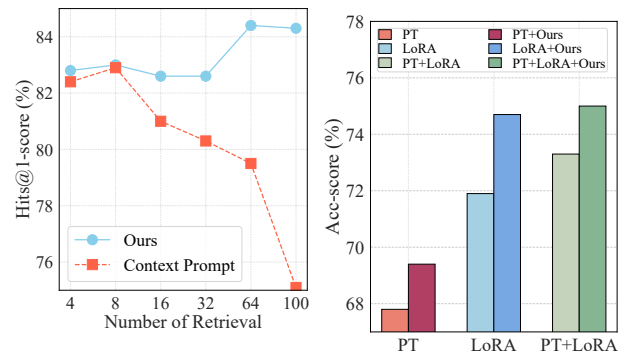


Figure 3: Performance vary with the number of Retrieval.

Figure 4: Combine with different fine-tuning methods.

Tuning (PT) (Lester et. al 2021) and Low-Rank Adaptation (LoRA). To ensure fairness, we conduct fine-tuning experiments without our module by concatenating the retrieved knowledge text with the input context. As shown in Figure 4, we observe a significant improvement in performance after incorporating AMAR, Regardless of whether PT or LoRA is used, our method consistently outperforms baselines. Notably, the combination of our method with PT+LoRA fine-tuning yields the best results. This highlights the capability of AMAR to effectively learn from retrieved information, while the direct concatenation of context introduces considerable noise. Furthermore, we find that LoRA fine-tuning outperforms PT fine-tuning. This can be attributed to the inherent complexity of logical form generation in the KGQA task. LoRA has more tunable parameters and can act on all project layers, thus enabling LLMs to better adapt to tasks.

### Case Study

In this section, we present a case study to illustrate how our model adaptively learns the importance of retrieval knowledge. We have not presented a case for subgraph retrieval due to its extensive length. As shown in Table 4, our method is proven effective in assigning high scores to correct entities and relations for entity retrieval and relation retrieval, while irrelevant or misleading information receives low scores. For instance, the entity ‘*Harper Lee*’ received a score of [0.9272], and relation ‘*people.person.education*’ received a score of [0.9844]. Nevertheless, when retrieved text is directly used as the context prompt, it is susceptible to interference from erroneous information in retrieved data. This can result in generating incorrect relations, such as ‘*education.educational.institution.school type*’. This example highlights that our method not only enhances the performance of LLM in KGQA but also improves the quality of retrieved information by setting weight.

### Analysis of LLM Backbones

In this section, we investigate the question: **Does the performance improvement of our method solely come from LLMs?** To answer this, we conduct experiments and compare results using different LLMs as backbones: LLaMA2-7B with fine-tuning, LLaMA2-13B with fine-tuning, and

Question	<i>what highschool did harper lee go to?</i>		
Entity Retrieval	① <i>Harper Lee</i> [0.9272], ② <i>Senior secondary education</i> [0.5441], ③ <i>Lee Remick</i> [0.6482], ④ <i>Secondary education</i> [0.5412], ⑤ <i>Barbara Kingsolver</i> [0.4320], ⑥ <i>High school movement</i> [0.6758]		
Relation Retrieval	① <i>people.person.education</i> [0.9844], ② <i>education.education.institution</i> [0.9844], ③ <i>education.educational_institution.school_type</i> [0.7281], ④ <i>education.school.lowest_grade_taught</i> [0.5469], ⑤ <i>education.school_mascot.school</i> [0.5431], ⑥ <i>common.topic.notable_types</i> [0.8252]		
Logical Form AMAR	by	(AND (JOIN <i>common.topic.notable_types</i> <i>High school</i> ) (JOIN (R <i>education.education.institution</i> ) (JOIN (R <i>people.person.education</i> ) <i>Harper Lee</i> ))) ✓	
Logical Form Context Prompt	by	(AND (JOIN <i>education.educational.institution.school_type</i> <i>School</i> ) (JOIN (R <i>education.education.institution</i> ) (JOIN (R <i>people.person.education</i> ) <i>Harper Lee</i> ))) ✗	

Table 4: A case study on WebQSP, where the ‘[float]’ represents the scores assigned to each retrieval information, indicating the level of influence it has on the model and the text with underline means erroneous generation.

Model	WebQSP		CWQ	
	Acc	Hits@1	Acc	Hits@1
<i>GPT-4</i>				
ToG-R	-	82.6	-	69.5
<i>LLaMA2-7B</i>				
G-Retriever	-	70.1	-	-
GNN-RAG	-	82.8	-	62.8
ChatKBQA	73.8	83.2	73.0	82.3
AMAR	<b>75.2</b>	<b>84.3</b>	73.4	82.9
<i>LLaMA2-13B</i>				
ChatKBQA	73.1	82.7	73.3	82.7
AMAR	74.7	83.3	<b>74.5</b>	<b>83.1</b>

Table 5: Analysis on different LLM backbones.

GPT-4 (frozen). In Table 5, we observe that AMAR with LLaMA2-7B and 13B as backbones outperform the baselines using the same backbone, such as ChatKBQA (Luo et al. 2023), and both even surpass ToG-R (Sun et al. 2024) using GPT-4 as the backbone on two datasets. These results suggest that the performance gain of AMAR is not solely attributed to the use of more capable LLMs but rather to the proposed utilization of the commonality between multi-aspect knowledge and the relevance of the question. We found that LLaMA2-13B perform worse than LLaMA2-7B on WebQSP. We believe the reason lies in dataset characteristics: the scale of WebQSP (1) is much smaller than that of CWQ, and (2) WebQSP has a maximum complexity, only consisting of 2-hop questions. This may cause the LLaMA2-13B to overfit, leading to reduced performance.

## Related Work

**Knowledge Graph Question Answering (KGQA).** KGQA aims to answer questions over KG, and previous methods are usually categorized as EM-based, IR-based, SP-based and LLM-based methods. EM-based methods encode the entities and relations in the embedding space and reason final answer using these structural embeddings (Shi et al. 2021). IR-based KBQA methods propose to retrieve and re-rank answers from KGs given information conveyed in the ques-

tion (Chen et al. 2019; Zhang et al. 2022). SP-based methods, focus on transforming question into a structural query, such as SPARQL and S-expression, and reason final answers using these executable queries (Liang et al. 2017; Lan and Jiang 2020). Besides, recent attempts have been made to utilize LLM-based methods for KGQA (Jiang et al. 2023c). For instance, LUO et al. (2024) presents a planning-retrieval-reasoning pipeline. ToG (Sun et al. 2024) proposes to interactively explore paths and reasoning on KGs using LLM as an agent. Despite their significant improvements, a substantial challenge persists when handling multi-aspect retrieved data as input, which may introduce irrelevant knowledge. However, our method is capable of aligning knowledge and implementing adaptive relevance gating with questions, thus addressing this issue.

**Large Language Model Reasoning.** Considering the impressive abilities of LLMs, some previous works focus on facilitating LLMs’ reasoning via prompting (He et al. 2021; Das et al. 2021; Jiang et al. 2023a; Wang et al. 2023). To overcome the unfaithful reasoning of LLMs, Jiang et al. (2023a) proposes StructGPT, an iterative reading-then-reasoning framework, to improve reasoning of LLMs when handling structured data. KD-CoT is designed to formulate chain-of-thought into a multi-round QA format and LLMs can retrieve external knowledge during interaction (Wang et al. 2023). However, LLMs reasoning is challenging when leveraging information present in lengthy texts (Wang et al. 2024). In contrast, our model facilitates efficient selection of information by utilizing prompt embeddings, thereby mitigating the issue of excessive context.

## Conclusion

In this work, we propose a novel approach to enhance LLMs reasoning and factual output by retrieving multi-aspect knowledge from KGs. By employing self-alignment and relevance gating modules, AMAR adaptively enhances and selects relevant information. It has proven more effective than simply appending retrieved text to input context, as it minimizes noise interference. Through extensive experiments, AMAR outperforms 22 baseline models and achieves a new state-of-the-art performance.

## Acknowledgments

This work was supported in part by the grants from National Science and Technology Major Project (No. 2023ZD0121104), and National Natural Science Foundation of China (No.62222213, 62072423). This research was partially supported by Research Impact Fund (No.R1015-23), APRC - CityU New Research Initiatives (No.9610565, Start-up Grant for New Faculty of CityU), CityU - HKIDS Early Career Research Grant (No.9360163), Hong Kong ITC Innovation and Technology Fund Midstream Research Programme for Universities Project (No.ITS/034/22MS), Hong Kong Environmental and Conservation Fund (No. 88/2022), and SIRG - CityU Strategic Interdisciplinary Research Grant (No.7020046), Tencent (CCF-Tencent Open Fund, Tencent Rhino-Bird Focused Research Program).

## References

- Bollacker, K.; Evans, C.; Paritosh, P.; Sturge, T.; and Taylor, J. 2008. Freebase: a collaboratively created graph database for structuring human knowledge. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, 1247–1250.
- Chen, Z.-Y.; Chang, C.-H.; Chen, Y.-P.; Nayak, J.; and Ku, L.-W. 2019. UHop: An Unrestricted-Hop Relation Extraction Framework for Knowledge-Based Question Answering. 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)*, 345–356. Minneapolis, Minnesota: Association for Computational Linguistics.
- Das, R.; Zaheer, M.; Thai, D.; Godbole, A.; Perez, E.; Lee, J.-Y.; Tan, L.; Polymenakos, L.; and Mccallum, A. 2021. Case-based Reasoning for Natural Language Queries over Knowledge Bases. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 9594–9611.
- Dong, Z.; Zhou, Z.; Yang, C.; Shao, J.; and Qiao, Y. 2024. Attacks, defenses and evaluations for llm conversation safety: A survey. *arXiv preprint arXiv:2402.09283*.
- Douze, M.; Guzhva, A.; Deng, C.; Johnson, J.; Szilvasy, G.; Mazaré, P.-E.; Lomeli, M.; Hosseini, L.; and Jégou, H. 2024. The Faiss library.
- Gabrilovich, E.; Ringgaard, M.; and Subramanya, A. 2013. FACC1: Freebase annotation of ClueWeb corpora, Version 1.
- Gu, Y.; Deng, X.; and Su, Y. 2023. Don't Generate, Discriminate: A Proposal for Grounding Language Models to Real-World Environments. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 4928–4949.
- Gu, Y.; Kase, S.; Vanni, M.; Sadler, B.; Liang, P.; Yan, X.; and Su, Y. 2021. Beyond IID: three levels of generalization for question answering on knowledge bases. In *Proceedings of the Web Conference 2021*, 3477–3488.
- He, G.; Lan, Y.; Jiang, J.; Zhao, W. X.; and Wen, J.-R. 2021. Improving multi-hop knowledge base question answering by learning intermediate supervision signals. In *Proceedings of the 14th ACM international conference on web search and data mining*, 553–561.
- He, K.; Mao, R.; Lin, Q.; Ruan, Y.; Lan, X.; Feng, M.; and Cambria, E. 2023. A survey of large language models for healthcare: from data, technology, and applications to accountability and ethics. *arXiv preprint arXiv:2310.05694*.
- He, X.; Tian, Y.; Sun, Y.; Chawla, N. V.; Laurent, T.; LeCun, Y.; Bresson, X.; and Hooi, B. 2024. G-retriever: Retrieval-augmented generation for textual graph understanding and question answering. *arXiv preprint arXiv:2402.07630*.
- Hu, X.; Wu, X.; Shu, Y.; and Qu, Y. 2022. Logical form generation via multi-task learning for complex question answering over knowledge bases. In *Proceedings of the 29th International Conference on Computational Linguistics*, 1687–1696.
- Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y. J.; Madotto, A.; and Fung, P. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12): 1–38.
- Jiang, J.; Zhou, K.; Dong, Z.; Ye, K.; Zhao, W. X.; and Wen, J.-R. 2023a. StructGPT: A General Framework for Large Language Model to Reason over Structured Data. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 9237–9251.
- Jiang, J.; Zhou, K.; Zhao, X.; and Wen, J.-R. 2023b. UniKGQA: Unified Retrieval and Reasoning for Solving Multi-hop Question Answering Over Knowledge Graph. In *The Eleventh International Conference on Learning Representations*.
- Jiang, J.; Zhou, K.; Zhao, X.; and Wen, J.-R. 2023c. UniKGQA: Unified Retrieval and Reasoning for Solving Multi-hop Question Answering Over Knowledge Graph. In *The Eleventh International Conference on Learning Representations*.
- Lan, Y.; and Jiang, J. 2020. Query Graph Generation for Answering Multi-hop Complex Questions from Knowledge Bases. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 969–974. Online: Association for Computational Linguistics.
- Lester et. al, B. 2021. The Power of Scale for Parameter-Efficient Prompt Tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 3045–3059. Online and Punta Cana, Dominican Republic: Association for Computational Linguistics.
- Li, B. Z.; Min, S.; Iyer, S.; Mehdad, Y.; and Yih, W.-t. 2020. Efficient One-Pass End-to-End Entity Linking for Questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 6433–6441.
- Liang, C.; Berant, J.; Le, Q.; Forbus, K. D.; and Lao, N. 2017. Neural Symbolic Machines: Learning Semantic Parsers on Freebase with Weak Supervision. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 23–33. Vancouver, Canada: Association for Computational Linguistics.
- Luo, H.; Tang, Z.; Peng, S.; Guo, Y.; Zhang, W.; Ma, C.; Dong, G.; Song, M.; Lin, W.; et al. 2023. Chatkbqa: A

- generate-then-retrieve framework for knowledge base question answering with fine-tuned large language models. *arXiv preprint arXiv:2310.08975*.
- LUO, L.; Li, Y.-F.; Haf, R.; and Pan, S. 2024. Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning. In *The Twelfth International Conference on Learning Representations*.
- Mavromatis et al., C. 2024. GNN-RAG: Graph Neural Retrieval for Large Language Model Reasoning. *arXiv preprint arXiv:2405.20139*.
- Miller, A.; Fisch, A.; Dodge, J.; Karimi, A.-H.; Bordes, A.; and Weston, J. 2016. Key-Value Memory Networks for Directly Reading Documents. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1400–1409.
- OpenAI. 2024. GPT-4 Technical Report. *arXiv:2303.08774*.
- Saxena et al., A. 2022. Sequence-to-Sequence Knowledge Graph Completion and Question Answering. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2814–2828.
- Shi, J.; Cao, S.; Hou, L.; Li, J.; and Zhang, H. 2021. TransferNet: An Effective and Transparent Framework for Multi-hop Question Answering over Relation Graph. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 4149–4158.
- Sun, J.; Xu, C.; Tang, L.; Wang, S.; Lin, C.; Gong, Y.; Ni, L.; Shum, H.-Y.; and Guo, J. 2024. Think-on-Graph: Deep and Responsible Reasoning of Large Language Model on Knowledge Graph. In *The Twelfth International Conference on Learning Representations*.
- Sun et al., H. 2018. Open Domain Question Answering Using Early Fusion of Knowledge Bases and Text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 4231–4242.
- Sun et al., H. 2019. PullNet: Open Domain Question Answering with Iterative Retrieval on Knowledge Bases and Text. 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)*, 2380–2390. Hong Kong, China: Association for Computational Linguistics.
- Talmor, A.; and Berant, J. 2018. The Web as a Knowledge-Base for Answering Complex Questions. In Walker, M.; Ji, H.; and Stent, A., eds., *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 641–651. New Orleans, Louisiana: Association for Computational Linguistics.
- Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Vrandečić, D.; and Krötzsch, M. 2014. Wikidata: a free collaborative knowledgebase. *Commun. ACM*, 57(10): 78–85.
- Wang, K.; Duan, F.; Wang, S.; Li, P.; Xian, Y.; Yin, C.; Rong, W.; and Xiong, Z. 2023. Knowledge-driven cot: Exploring faithful reasoning in llms for knowledge-intensive question answering. *arXiv preprint arXiv:2308.13259*.
- WANG, L.; MA, C.; FENG, X.; ZHANG, Z.; YANG, H.; ZHANG, J.; CHEN, Z.; TANG, J.; CHEN, X.; LIN, Y.; ZHAO, W. X.; WEI, Z.; and WEN, J. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6): 186345.
- Wang, X.; Salmani, M.; Omid, P.; Ren, X.; Rezagholizadeh, M.; and Eshaghi, A. 2024. Beyond the limits: A survey of techniques to extend the context length in large language models. *arXiv preprint arXiv:2402.02244*.
- WEN, M.; LIN, R.; WANG, H.; YANG, Y.; WEN, Y.; MAI, L.; WANG, J.; ZHANG, H.; and ZHANG, W. 2023. Large sequence models for sequential decision-making: a survey. *Frontiers of Computer Science*, 17(6): 176349.
- Xie, T.; Wu, C. H.; and et. al. 2022. UnifiedSKG: Unifying and Multi-Tasking Structured Knowledge Grounding with Text-to-Text Language Models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 602–631. Abu Dhabi, United Arab Emirates: Association for Computational Linguistics.
- XU, D.; CHEN, W.; PENG, W.; ZHANG, C.; XU, T.; ZHAO, X.; WU, X.; ZHENG, Y.; WANG, Y.; and CHEN, E. 2024. Large language models for generative information extraction: a survey. *Frontiers of Computer Science*, 18(6): 186357.
- Ye, X.; Yavuz, S.; Hashimoto, K.; Zhou, Y.; and Xiong, C. 2022. RING-KBQA: Generation Augmented Iterative Ranking for Knowledge Base Question Answering. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, 6032–6043. Dublin, Ireland: Association for Computational Linguistics.
- Yih, W.-t.; Richardson, M.; Meek, C.; Chang, M.-W.; and Suh, J. 2016. The value of semantic parse labeling for knowledge base question answering. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, 201–206.
- Yu, D.; Zhang, S.; Ng, P.; Zhu, H.; Li, A. H.; Wang, J.; Hu, Y.; Wang, W. Y.; Wang, Z.; and Xiang, B. 2022. DecAF: Joint Decoding of Answers and Logical Forms for Question Answering over Knowledge Bases. In *The Eleventh International Conference on Learning Representations*.
- ZENG, Y.; LI, Z.; CHEN, Z.; and MA, H. 2023. Aspect-level sentiment analysis based on semantic heterogeneous graph convolutional network. *Frontiers of Computer Science*, 17(6): 176340.
- Zhang, J.; Zhang, X.; Yu, J.; Tang, J.; Tang, J.; Li, C.; and Chen, H. 2022. Subgraph Retrieval Enhanced Model for Multi-hop Knowledge Base Question Answering. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5773–5784.
- Zhang, L.; Zhang, J.; Wang, Y.; Cao, S.; Huang, X.; Li, C.; Chen, H.; and Li, J. 2023. FC-KBQA: A Fine-to-Coarse Composition Framework for Knowledge Base Question Answering. In *The 61st Annual Meeting Of The Association For Computational Linguistics*.