

Large Language Models Can Better Understand Knowledge Graphs Than We Thought

Anonymous ACL submission

Abstract

As the parameter scale of large language models (LLMs) grows, jointly training knowledge graph (KG) embeddings with model parameters to enhance LLM capabilities becomes increasingly costly. Consequently, the community has shown interest in developing prompt strategies that effectively integrate KG information into LLMs. However, the format for incorporating KGs into LLMs lacks standardization; for instance, KGs can be transformed into linearized triples or natural language (NL) text. Current prompting methods often rely on a trial-and-error approach, leaving researchers with an incomplete understanding of which KG input format best facilitates LLM comprehension of KG content. To elucidate this, we design a series of experiments to explore LLMs’ understanding of different KG input formats within the context of prompt engineering. Our analysis examines both literal and attention distribution levels. Through extensive experiments, we indicate a counter-intuitive phenomenon: when addressing fact-related questions, unordered linearized triples are more effective for LLMs’ understanding of KGs compared to fluent NL text. Furthermore, noisy, incomplete, or marginally relevant subgraphs can still enhance LLM performance. Finally, different LLMs have distinct preferences for different formats of organizing unordered triples.

1 Introduction

Currently, researchers commonly utilize databases containing extensive factual knowledge, such as knowledge graphs (KGs), to reduce hallucination in language models and enhance the quality of their generated content (Pan et al., 2023). In the era of pre-trained language models (PLMs), integrating KGs during the training process has garnered significant interest within the community (Zhang et al., 2019; Xiong et al., 2019; Su et al., 2021; Arora et al., 2022; Chen et al., 2022). As language models evolve, training large-scale language

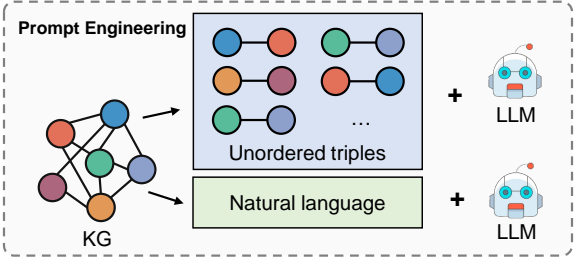


Figure 1: KG is processed into different input formats to provide knowledge to LLM.

models (LLMs) with billions of parameters using KG data may encounter limitations, such as severe resource constraints or lack of public access to model architectures, training data, or training methods (Ufuk, 2023). To address these problems, recent researches increasingly focus on injecting external knowledge into LLMs through prompt engineering techniques (Sorensen et al., 2022; White et al., 2023; Li et al., 2023; Wen et al., 2023). This lightweight approach leverages KGs to enhance LLMs and demonstrates efficacy in addressing various challenges related to factual knowledge. Some studies indicate that LLMs are highly sensitive to input patterns, and different input formats can impact model performance (Sclar et al., 2023; Voronov et al., 2024; Zhan et al., 2024).

As a kind of highly structural knowledge, KG can be input into LLM in various formats. As shown in Figure 1, recent works process KGs into unordered linearized triples and directly feed them into LLMs (Baek et al., 2023; Sen et al., 2023). Some other research employs KG-to-text generation approaches to convert structural knowledge prompts into natural language (NL) text, aiming to bridge the semantic gap between them (Ma et al., 2022; Xiong et al., 2022; Wu et al., 2023; Guo et al., 2023). However, generating text from KGs becomes a significant challenge when dealing with subgraphs containing numerous triples (tens or

even hundreds). Furthermore, in some tasks where factual accuracy is crucial, the necessity of narrowing the semantic gap between different formats of knowledge has not been further discussed, leaving researchers to rely solely on conjecture and experimentation. This situation underscores a significant opportunity in the field: comprehending which input format is the most conducive to LLMs’ understanding of KGs. Such advancements could pave the way for designing more reliable and efficient KG-related prompting strategies.

In this study, we explore what KG input formats can help LLMs better understand structural knowledge in the context of prompt engineering (PE) (Sahoo et al., 2024). We ask LLMs (e.g., ChatGPT¹ and Vicuna (Zheng et al., 2024)) to answer complex questions based on different KG input formats. Compared to other KG-related tasks (such as KG-to-text generation and link prediction), the question-answering (QA) task is not only more difficult but also offers a unique advantage in evaluation. Our questions encompass entity enumeration, counting, ranking, comparison, and truthfulness assessment, thereby offering a thorough evaluation of LLMs’ ability to comprehend externally injected KG. Moreover, the evaluation method for QA is objective, enabling a direct comparison of answers and thereby accurately assessing LLM responses.

Specifically, we convert KG triples from multiple knowledge graph question answering (KGQA) datasets into NL text using a generative model. This transformation from triples to NL text is termed *Triple-to-Text*. This raises a concern: will the performance of LLMs in understanding KG affected by potential errors in the model-generated text? Therefore, leveraging a document relation extraction dataset DocRED (Yao et al., 2019), we also transform the organization of knowledge from NL text into unordered linearized triples, and dub the transformation *Text-to-Triple*. The results of these two complementary experiments highlight a critical insight: **When addressing questions related to KG facts, compared to NL text, LLMs can better understand unordered structural knowledge and perform better.**

To elucidate this finding, we analyze LLMs at both literal and attention distribution levels. Intriguingly, our experiments reveal that providing noisy or incomplete subgraphs can still enhance LLM performance. Even when we remove all rea-

soning paths directly related to the answer and only provide adjacent subgraphs partially related to the answer, we observe performance improvements. Lastly, we observe that different LLMs show distinct preferences for prompt formats that are used for organizing unordered triples. We hope these findings provide the community with insights for better designing KG-related prompt strategies to enhance LLM performance. Our main findings are as follows:

- *When using external knowledge to answer fact-intensive questions, LLMs prefer unordered and abstract structural data over fluent natural language text.* Even in the absence of apparent logical expressions, the model can still identify key information relevant to the answer.
- *Noisy, incomplete and even marginally related knowledge can still yield beneficial results in model reasoning.* When the knowledge does not contain the key information necessary to solve the question, but only a subgraph of neighboring nodes to the key information, the model’s reasoning ability can still be partially enhanced.
- *A prompt strategy that performs well on one model may not be effective on another.* This suggests that researchers need to conduct meticulous experiments for more universally applicable knowledge injection prompts.

2 Related Works

Injecting KG knowledge into LLM during training. Injecting knowledge from KGs into LLMs during training has been extensively researched. This approach enables LLMs to grasp the semantics of KG embeddings through collaborative training (Zhang et al., 2019; Xiong et al., 2019; Liu et al., 2020; Su et al., 2021; Zha et al., 2022; Chen et al., 2022). Although these methods have shown progress in smaller LLMs, their applicability to larger-scale LLMs (such as ChatGPT) presents challenges that require careful consideration of model architecture, training methods, and other aspects. In addition, injecting knowledge within training may result in insufficient and incorrect internalized knowledge in LLMs (Ji et al., 2023).

Integrating KG knowledge for LLM using prompt engineering techniques. Integrating external structural knowledge into prompts to enhance

¹<https://openai.com/chatgpt>

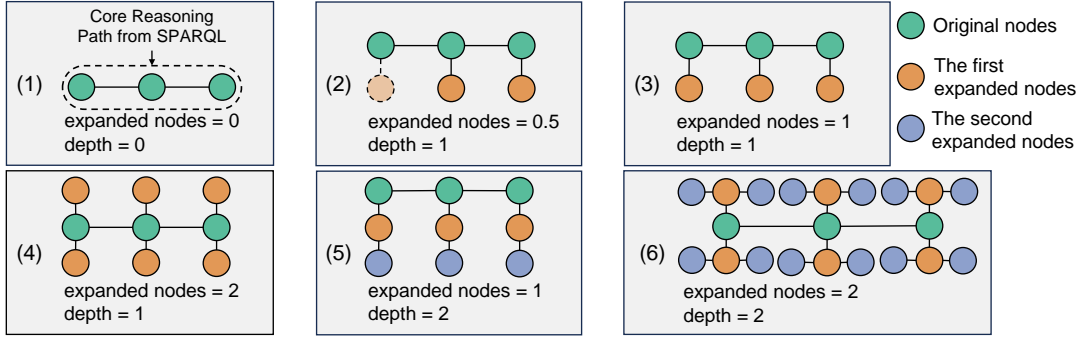


Figure 2: Six categories of our expansion method: (1) only providing core reasoning paths; (2) expanding each node on each core path by 1 neighboring node, with a 50% probability of deleting this expansion node (deleted nodes are represented by dashed lines); (3) expanding each node on each core path by 1 neighboring node and retaining all of them; (4) expanding each node on each core path by 2 neighboring nodes and retaining all of them; (5) starting from nodes on the core path, expanding to 2-hop neighboring nodes, expanding 1 node at a time; (6) starting from nodes on the core path, expanding to 2-hop neighboring nodes, expanding 2 nodes at a time. According to this rule, (4) and (5) contain the same number of nodes in the subgraph but exhibit different structures.

LLM capabilities has become a common strategy. Some studies (Baek et al., 2023; Sen et al., 2023) directly provide LLMs with linearized structural knowledge as part of the prompts. (Chen, 2023) linearizes structural data into unified table rows, inputting them into the LLM to generate answers based on contextual examples. StructGPT (Jiang et al., 2023) first selects entity and relation candidates, then uses multiple rounds of retrieving interface calls and LLMs (such as ChatGPT) as rankers to obtain answers.

However, other research suggests that inputs in NL format may more model-friendly. These studies first convert structural knowledge into NL and then use it as prompts. UDT-QA (Ma et al., 2022) treats structural data as a form of knowledge expansion, converting it into NL text and adding it to the document repository for retrieval. (Xiong et al., 2022) transform subgraphs extracted from SPARQL queries into NL paragraphs, incorporating them into prompts to drive LLMs in generating NL questions. (Wu et al., 2023) believe that converting structural knowledge into high-quality NL text can substantially reduce the semantic gap between them. They employ KG-to-text generation models to rewrite structural knowledge and use it to help LLMs answer questions. KnowledgeNavigator (Guo et al., 2023) performs efficient reasoning on KGs and uses templates to convert structural reasoning paths into NL, guiding LLM reasoning.

3 Methods

We emphasize in Section 1 that QA tasks have unique advantages in evaluating LLM’s understand-

ing of KGs. Thus, in our work, we define the performance of LLMs in understanding KGs as: which input format of KGs can better aid LLMs in accurately answering questions? Thus, we introduce two complementary analytical methods (*Triple-to-Text* and *Text-to-Triple*) to examine how LLMs understand KG knowledge in different formats. In Section 3.1, we gradually transform KGs into NL text to test the LLMs’ understanding of knowledge in various formats. In Section 3.2, we use human-written documents to create KGs and conduct the same evaluation to avoid potential errors from text generated by KG-to-text models.

In Section 3.3, we further analyze this phenomenon at both the literal level and attention distribution level to understand the reasons behind this result. In Section 3.4, we assess how noisy, incomplete, subgraphs lacking reasoning paths impact the performance of LLMs. Finally, in Section 3.5, we explore various prompt strategies based on unordered linearized triples to enhance LLMs. Our objective is to observe whether different LLMs exhibit consistent preferences for these methods.

3.1 LLM’s Understanding of *Triple-to-Text*

In this section, we aim to explore how LLM understands KG in different scales and input formats. Therefore, during the experiments, we control two variables: subgraph size and KG input format.

To quantitatively investigate the impact of the subgraph scale for LLMs, we propose a controllable expansion method to regulate the scale of the whole subgraph and the proportion of core reasoning paths within the subgraph. Specifically, we

use the Wikidata terminal to retrieve all constraint variable IDs from the SPARQL corresponding to each question in the KGQA dataset. SPARQL, as a query language that can retrieve the exact answer to a question in the KG, can help us obtain all answers and corresponding multiple reasoning paths. Moreover, to mitigate potential data leakage issues (Balloccu et al., 2024), we convert the retrieved variable IDs into English labels. As shown in Figure 2, the number of neighboring nodes expanded from the core reasoning path and the number of hops is controlled by the parameters *expanded nodes* and *depth*. We categorize our expansion method into six distinct types, which aid in evaluating the reasoning performance of LLMs on these subgraphs with different structures. We describe more details in Appendix A.1 for core paths generation.

To investigate the impact of different knowledge formats for LLMs, we refer to some intermediate steps in recent works and devise five levels of KG injection methods: (1) Omitting KG subgraphs (Without Triples). (2) Providing unordered linearized triples (Unordered Triples). We randomly shuffle all triples in the KG. (3) Combining triples that share the same head or tail entities to form meta paths (Gao et al., 2020) (Meta Paths). (4) Utilizing heuristic rules to convert meta paths into NL texts (Wang et al., 2022) (Natural Language (By rules)). (5) Employing a text generation model to convert meta paths into NL texts (Tang et al., 2023) (Natural Language (By model)). Different combinations of extension and injection modes can form 25 possible combinations (Omitting KG subgraphs is treated as one method, irrespective of the KG size). In the experimental section, we will analyze the outcomes of these combinations.

3.2 LLM’s Understanding of *Text-to-Triple*

The quality of NL text generated from KG may contain errors, potentially affecting the ability of LLMs to understand KGs in NL format. To eliminate such worry, it is essential to provide LLMs with human-written NL text and establish a mapping from text to KG. Thus, we utilize the document-level relation extraction dataset DocRED (Yao et al., 2019) as a basis to construct a dataset for evaluating LLMs’ understanding of KGs. All documents in DocRED are human-written, and all mapped triples are manually annotated and aligned with entities and relations in Wikidata. This dataset emphasizes cross-sentence reasoning, allowing the mapped triples within each document to form a complete, small-

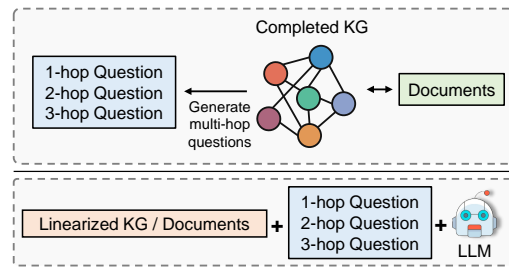


Figure 3: We employ a completed KG mapping with a document for generating fact-related questions, and evaluate the performance of LLM in answering these questions with different knowledge formats.

scale KG. We consider this small-scale complete KG as a complete structural representation of all entities and relations involved in the document. Based on these small-scale KGs, we can extract reasoning paths of different hops to generate questions.

Specifically, to generate QA pairs from each document, we refer to the dataset construction approach from LC-QuAD 2.0 (Dubey et al., 2019). We first fill the triples into various multi-hop templates to construct template questions, then use ChatGPT to paraphrase these questions to ensure diversity and complexity. We generate 1-hop, 2-hop, and 3-hop questions to evaluate the LLM’s ability based on a complete KG. As shown in Figure 3, the input format for the KG is categorized into unordered triples and human-written NL documents. We separately provide the unordered triples and human-written NL documents to the LLM to answer the questions. We present some example questions generated from the dataset in Appendix A.2.

3.3 Analysis of LLM’s Understanding

To conduct a deeper analysis of this phenomenon, we analyze LLM’s response at both the literal level and the attention distribution level.

At the literal level, we directly inquire about LLM’s preference for different types of knowledge. For fact-related questions, LLM indicates a preference for choosing linear triples. In contrast, for open-domain questions that require divergent thinking or more comprehensive descriptions, LLM may exhibit a preference for textual knowledge. The responses from LLMs suggest that they retrieve facts more easily and provide accurate answers from unordered triples compared to NL text. In Appendix A.3, we present a case of LLM’s response to a fact-related question.

At the attention distribution level, we observe the proportion of attention that the models allocate

to the correct answers within the input. The input knowledge is divided into triples and NL text. In the final layer of the model, we sum the attention between each token and the final token (i.e., the model-generated prediction label q) across all heads, and normalized all values to obtain the attention proportion of q to each token.

We focus on the attention proportion of q towards the answers within the distribution. It reflects the extent to which the LLM concerns the answers in the provided knowledge. A higher proportion suggests a greater likelihood that the model will include the answer in its generated response. Considering that the number of answer occurrences may vary across different formats of knowledge, we calculate the average attention proportion for each answer to ensure fairness. The model’s average attention proportion towards the answer is $\overline{Att} = \frac{1}{n} \sum_{i=1}^n Att_{ans_i,q}$, where $Att_{ans_i,q}$ represents the attention proportion between the tokens in the answer i and the prediction label q , and n is the total number of occurrences of the answer tokens in the knowledge text. We employ two fair comparison knowledge prompting modes: (1) providing the LLM with a single type of knowledge separately (triples or NL), and (2) providing the LLM with both types of knowledge simultaneously (triples and NL). These two prompting strategies are detailed in Appendix A.4.

3.4 LLM’s Robustness to Noisy or Incomplete Sub-graphs

In order to assess the capacity of LLM to comprehend noisy sub-graphs, we systematically sabotaged the external knowledge within the injecting sub-graphs. Specifically, we employ three approaches to alter the sub-graphs: (1) Nodes are proportionally deleted randomly. (2) Nodes are proportionally replaced with random irrelevant KG nodes. (3) All core reasoning paths in the sub-graph are removed. To ensure a smooth transition when replacing and deleting elements based on percentages, we employ the max expanded sub-graph ($expanded\ nodes=2, depth=2$, details can be found in Appendix A.1). The deletion and replacement ratio ranges from 10% to 90%.

The way to delete the core reasoning path is to delete all the green nodes in Figure 3. Deleting the core reasoning path can be viewed as a complete destruction of the path from the start-point node to the target node, preventing LLM use other nodes as intermediate steps to reach the target.

3.5 KG Prompt Strategies

Based on the findings of this study, we propose some strategies that leverage KG information to enhance the prompting of LLMs and evaluate their impact on LLM performance. In practical scenarios, triples serve as external data, which can be utilized to provide relevant knowledge to LLMs through retrieval-augmented generation (RAG) techniques (Yu et al., 2023; Asai et al., 2023; Tang and Yang, 2024). Consequently, we employ BERT as a scorer to retrieve triples from KG. The relevance score S of a triple t to a question q is calculated as $S = BERT(t, q)$.

In order to find a universal prompting strategy suitable for LLMs, we select the top 100 triples from the KG with the highest S as candidates and perform the following operations on these triples: 1) Grouping: Triples are divided into three groups based on two thresholds, namely most relevant, moderately relevant, and least relevant triples. 2) Ranking: We sort the candidate triples in descending order of S without displaying the scores. 3) Scoring: We append the S to the triples to indicate their related confidence to the question, providing a reference for the LLM. Appendix A.5 illustrates these three types of inputs and the training detail of BERT-based scorer.

4 Experiment

In this section, we first set up the experiments, then answer five questions about how LLM understands KGs using the methods introduced in Section 3.

4.1 Experimental Setup

4.1.1 Evaluation metrics

We ask the LLMs to respond in a unified format to facilitate evaluation, subsequently, we calculate the string similarity between the predicted answers and the gold standard answers. The accuracy of responses to each question is defined as: $Accuracy = N_{correct}/N_{total}$. Detailed information regarding the input format and metric settings are discussed in Appendix A.6.

4.1.2 Datasets and LLMs

Wikidata (Vrandečić and Krötzsch, 2014) is a large-scale, high-quality KG that is frequently updated. We select three KGQA datasets based on Wikidata, which include SPARQL queries: QALD-7 (Usbeck et al., 2017), LC-QuAD 2.0 (Dubey et al., 2019), and KQAPro (Cao et al., 2022). These datasets

	Expanded Node Ratio	Expanded Nodes	Hop	Knowledge Injection Method														
				Without Triples			Unordered Triples			Meta Paths			Natural Language (By rules)			Natural Language (By model)		
				ChatGPT	Vicuna 7b	Vicuna 13b	ChatGPT	Vicuna 7b	Vicuna 13b	ChatGPT	Vicuna 7b	Vicuna 13b	ChatGPT	Vicuna 7b	Vicuna 13b	ChatGPT	Vicuna 7b	Vicuna 13b
QALD-7 (64)	0%	0	0	71.45	14.35	60.82	84.01	74.59	78.02	84.01	70.00	73.29	75.07	60.71	71.44	73.62	60.33	65.65
	33.33%	0.5	1	71.45	14.35	60.82	84.00	66.44	79.42	84.00	66.38	76.18	76.04	53.61	74.59	69.28	46.81	68.79
	50%	1	1	71.45	14.35	60.82	84.01	72.41	79.90	84.01	72.19	73.14	71.21	39.76	61.54	57.68	38.60	52.85
	66.67%	2	1	71.45	14.35	60.82	84.01	67.34	81.11	84.01	60.05	77.00	68.55	47.67	64.69	59.86	47.00	55.27
	85.72%	2	2	71.45	14.35	60.82	82.56	58.65	78.21	84.01	58.55	73.14	68.31	48.26	57.68	54.54	38.84	52.76
LC-QuAD 2.0 (2000)	0%	0	0	16.42	2.60	13.47	50.80	36.76	46.62	50.29	33.47	39.53	23.28	11.32	17.34	21.12	10.04	16.88
	33.33%	0.5	1	16.42	2.60	13.47	51.61	36.37	45.75	50.16	31.88	39.12	26.98	14.13	21.10	23.32	10.54	18.12
	50%	1	1	16.42	2.60	13.47	51.71	35.95	42.43	48.55	30.73	38.22	24.47	12.12	20.47	20.49	10.79	17.70
	66.67%	2	1	16.42	2.60	13.47	51.24	35.18	43.73	47.57	29.38	36.29	23.13	12.59	18.91	19.34	9.74	17.04
	85.72%	2	2	16.42	2.60	13.47	51.35	33.97	44.96	48.82	30.11	37.66	23.85	11.52	19.69	21.00	9.44	17.27
KQAPro (2000)	0%	0	0	15.77	3.55	10.29	49.33	23.10	36.69	48.10	22.09	27.74	34.93	12.19	18.88	27.15	10.88	15.77
	33.33%	0.5	1	15.77	3.55	10.29	51.45	25.39	33.33	49.97	23.35	28.61	29.89	12.08	18.79	23.57	10.68	15.54
	50%	1	1	15.77	3.55	10.29	52.54	25.64	31.60	50.78	24.72	28.22	28.80	12.39	19.07	23.60	10.51	15.30
	66.67%	2	1	15.77	3.55	10.29	52.46	28.47	31.04	50.00	25.42	27.24	28.19	12.36	18.93	22.54	9.70	14.77
	85.72%	2	2	15.77	3.55	10.29	54.19	26.85	32.72	51.45	25.39	27.77	28.61	11.27	16.89	23.33	8.47	12.67
				15.77	3.55	10.29	54.03	27.63	32.75	51.51	25.70	28.58	22.60	10.23	15.30	17.20	7.49	11.49

Table 1: The understanding of KG subgraphs by LLMs.

serve as the foundation for generating our datasets. QALD-7 contains 215 training questions and 50 test questions. LC-QuAD 2.0 comprises 24k training questions and 6046 test questions. KQAPro includes 94k training questions and 10k test questions. From these datasets, we use a SPARQL endpoint to retrieve answers from Wikidata, filtering out questions with incorrect or unanswerable results. We also delete the questions where the core reasoning path cannot be extended to two hops, such as when the path contains numerical or other attribute information. After filtering, QALD-7 retains 64 questions. For LC-QuAD 2.0 and KQAPro, considering the cost of calling the API, we randomly select 2000 questions from each dataset.

To address concerns raised in Section 3.1 regarding the decline in LLM performance due to the quality of text generated by models, we utilize human-written documents from the DocRED (Yao et al., 2019) as the NL representation of the KG. This dataset consists of 5053 Wikipedia documents, each associated with a set of manually annotated KG triples. We select 800 documents that contain at least 3-hop paths and generate 1-hop, 2-hop, and 3-hop questions for each document.

We employ ChatGPT, GPT-4o², LLaMA3³, Vicuna 7B and 13B (Zheng et al., 2024) to evaluate the data, with all model parameters fixed. ChatGPT and Vicuna represent two mainstream series of LLMs, both demonstrating decent performance. Additionally, our hardware resources allow us to analyze attention distribution in Vicuna models with 7B and 13B parameters, enabling a performance comparison of different scale models. GPT-4o and LLaMA3 are LLMs released after 2023. We con-

sider GPT-4o as the latest iteration of ChatGPT and LLaMA3 as the newest iteration of Vicuna. Our aim is to use these new models to observe whether their ability to understand KGs of different input formats has changed over time.

4.2 Does unordered structural knowledge facilitate a better understanding of KG for LLMs than fluent NL text?

We conduct a horizontal comparison of the data in Table 1. The best results, highlighted in bold, indicate that the unordered triples consistently outperform other methods. The limited size (only 64 questions) of the QALD-7 dataset initially obscures the advantages of our method. However, as the complexity of the questions and the scale of the data increase (as seen in LC-QuAD 2.0 and KQAPro), the unordered triples demonstrate better performance in the knowledge prompt. Specifically, ChatGPT, when using unordered triple knowledge to answer questions, achieves over a 30% gain on LC-QuAD 2.0 and at least a 20% gain on KQAPro compared to using NL text. The Vicuna series models also exhibit significant gains. This suggests that unordered triple knowledge can aid LLMs in better answering multi-hop fact-intensive questions.

800 questions	ChatGPT	Vicuna 7b	Vicuna 13b	GPT4o	LLaMA3
1-hop Text	25.25	14.88	27.62	29.16	34.58
1-hop Triple	73.38	50.13	73.26	73.38	78.97
2-hop Text	14.25	9.50	15.37	14.77	19.72
2-hop Triple	19.88	11.00	16.38	19.46	21.76
3-hop Text	14.00	8.63	13.87	12.67	17.45
3-hop Triple	18.25	10.50	14.75	19.38	25.71

Table 2: The performance of LLM in answering multi-hop questions using complete KGs and NL text.

To dispel doubts about the quality of text generated by the model, we construct a QA dataset based on DocRED. The data in Table 2 indicate that, com-

²<https://openai.com/index/hello-gpt-4o>

³<https://llama.meta.com/llama3>

pared to NL text, LLMs achieve significant gains when utilizing triples. Notably, ChatGPT achieves a remarkable 48.13% enhancement in 1-hop questions by using linearized triple knowledge prompts. For 2-hop and 3-hop questions, the improvements are 5.63% and 4.25%, respectively.

Furthermore, to dispel doubts that this phenomenon might change with the development of the models' version, we further evaluate all data on GPT-4o and LLaMA3. Our findings reveal that the advantage of triple knowledge in understanding KGs by LLMs has not diminished with the advancement of these models. This further eliminates our concern: the decline in performance on textual knowledge is not due to poor quality knowledge representation. Moreover, the benefits of triple knowledge for understanding KGs remain evident even in the latest LLMs.

4.3 Does the scale of subgraphs affect the reasoning capabilities of LLMs in reasoning?

We conduct a vertical comparison of the data in Table 1. The best results, underlined, represent the maximum values for each subgraph expansion within each dataset. We find that subgraphs with a lower proportion of core reasoning paths (i.e. larger subgraphs) do not necessarily result in the weakest model reasoning performance. For instance, on KQAPro, the performance of LLMs gradually decreases with the increase in the subgraph scale for knowledge input in an NL format. However, this trend does not hold for LC-QuAD 2.0. This indicates that within the permissible input length, LLMs can retrieve useful knowledge from the background to answer questions. Additionally, although the configurations of 'expanded nodes=2' with 'depth=1' and 'expanded nodes=1' with 'depth=2' expand different subgraph structures, our results suggest LLM remains insensitive to the structure variation.

4.4 Why do LLMs show a understanding preference for specific input formats of KG?

We need to further analyze the underlying causes of this phenomenon. At the literal level, by querying the LLM itself, we observe that it has a clear preference for the type of knowledge input when solving fact-intensive questions. LLMs are capable of retrieving answers in unordered triples. The model's detailed responses are provided in Appendix A.3.

	Vicuna 7b				Vicuna 13b			
	KQAPro		LCQUAD 2.0		KQAPro		LCQUAD 2.0	
	Triple	Text	Triple	Text	Triple	Text	Triple	Text
Single	3.80	2.46	2.67	2.08	4.69	2.78	3.78	2.10
Double	1.31	1.12	2.38	0.89	2.72	1.62	2.57	1.40

Table 3: Average attention proportion to the answer for all questions.

At the attention distribution level, the experimental results of the method used in Section 3.3 are presented in Table 3. For each question, there are two input modes: providing only one type of knowledge in the prompt (Single) and providing both types of knowledge simultaneously (Double). The experimental results indicate that, regardless of whether both types of knowledge are provided simultaneously or separately, the model exhibits a higher attention proportion to the answer in linear triples. This indicates that the LLMs do not reject less fluent knowledge prompts and are capable of retrieving answers from unordered knowledge. Therefore, utilizing linearized triples as knowledge facilitates the appearance of the correct answer.

4.5 How do LLMs perform when handling subgraphs of KGs that are noisy, incomplete, or lack core reasoning paths?

We perform deletion and replacement operations on subgraphs based on the parameters of *expanded nodes=2* and *depth=2* from Section 3.1. The deletion simulates an incomplete KG scenario, and the replacement substitutes some nodes with unrelated nodes in KG to generate false factual information, thereby simulating noisy subgraphs.

Based on Table 4, we have two preliminary findings: 1) The random replacement of nodes in KG has a more significant impact on the inference performance of LLM compared to the random deletion. Incorrect facts are more likely to result in erroneous model outputs. 2) Despite larger models demonstrating superior answering performance, they exhibit a greater performance loss when subjected to random replacement and deletion of KG knowledge. There exists an inverse proportionality between a model's robustness and its size.

More importantly, we find that even when we delete and replace up to 90% of the triples, the ability of the LLM to answer questions still improves compared to having no information at all. Thus, in Table 6, we specifically removed all core reasoning paths within a subgraph and observed that as long

	ChatGPT						Vicuna 7b						Vicuna 13b					
	QALD-7		LC-QuAD 2.0		KQAPro		QALD-7		LC-QuAD 2.0		KQAPro		QALD-7		LC-QuAD 2.0		KQAPro	
Ratio	Delete	Replace	Delete	Replace	Delete	Replace	Delete	Replace	Delete	Replace	Delete	Replace	Delete	Replace	Delete	Replace	Delete	Replace
0%	82.56	82.56	50.60	50.60	54.03	54.03	65.70	65.70	32.84	32.84	27.63	27.63	77.00	77.00	42.98	42.98	32.75	32.75
10%	82.08	80.62	47.84	47.74	51.57	52.04	64.44	58.89	31.07	32.20	26.12	26.23	72.85	78.21	41.66	40.55	31.32	31.24
20%	82.08	80.62	46.23	46.02	49.66	49.91	70.00	58.65	31.83	29.59	26.09	26.17	75.51	81.16	39.43	39.43	30.96	31.07
30%	80.39	79.90	44.18	43.60	46.95	46.87	62.95	59.08	29.51	28.10	23.52	24.58	74.59	78.70	38.57	38.85	28.94	29.81
40%	80.39	79.90	42.49	42.41	43.82	44.10	55.27	57.68	27.83	27.47	23.60	23.94	77.00	73.84	36.12	36.83	27.93	28.19
50%	80.39	80.39	40.78	40.28	41.02	41.33	66.86	59.66	27.78	26.03	21.95	21.87	76.52	74.35	34.57	35.09	26.26	27.88
60%	80.39	80.39	36.25	35.99	37.28	37.08	63.24	61.30	24.50	22.56	21.37	19.16	75.31	78.21	33.50	32.94	25.53	25.36
70%	75.56	75.56	31.87	31.72	32.33	32.02	63.00	48.74	21.82	19.44	18.06	16.69	73.62	68.07	30.54	27.36	23.01	23.38
80%	65.85	63.91	26.28	25.82	27.21	27.46	50.00	47.25	18.54	16.50	14.57	13.49	65.17	66.14	25.56	23.79	20.41	20.25
90%	62.42	54.30	19.16	17.64	19.91	19.77	46.64	45.02	13.21	11.98	10.49	9.82	57.51	56.23	18.48	17.51	14.15	14.90
Degradation Ratio	20.14	28.26	31.44	32.96	34.12	34.26	19.06	20.68	19.63	20.86	17.14	17.81	<u>19.49</u>	<u>20.77</u>	<u>24.50</u>	<u>25.47</u>	<u>18.60</u>	<u>17.85</u>

Table 4: Randomly delete and replace nodes in the subgraph. The Degradation Ratio quantifies the discrepancy between the model’s peak performance and its poorest performance. Values in bold denote the maximum, while underlined values signify the second highest. Larger models exhibit the most performance degradation when faced with attacks.

Data Set	ChatGPT			Vicuna 7b			Vicuna 13b		
	Grouping	Ranking	Scoring	Grouping	Ranking	Scoring	Grouping	Ranking	Scoring
QALD-7	84.11	84.11	84.11	63.84	64.54	53.64	75.52	77.81	72.4
LC-QuAD 2.0	48.71	50.01	52.48	33.49	35.72	26.14	45.1	45.13	42.57
KQAPro	50.25	52.29	54.03	27.74	31.32	24.92	36.05	37.64	35.12

Table 5: KG prompt engineering. Distinct models exhibit unique preferences towards various prompting methods.

Dataset	ChatGPT	Vicuna 7b	Vicuna 13b
QALD-7	79.42 (71.45)	53.73 (14.35)	72.13 (60.82)
LC-QuAD 2.0	27.51 (16.42)	18.37 (2.60)	23.17 (13.47)
KQAPro	30.09 (15.77)	18.82 (3.55)	22.18 (10.29)

Table 6: LLMs use subgraph without reasoning path to answer questions. In brackets are the performance of the model without any knowledge (form Table 1).

as some triples were present (even if they weren’t particularly relevant), the QA performance of LLM remained superior to providing no information. We hypothesize that these triples still stimulate LLM thinking with question-related information, thus driving the model’s ability to generate accurate answers. Some cases are shown in the Appendix A.7.

4.6 What factors should be considered when designing prompt frameworks that leverage KGs as external knowledge?

The results of the knowledge prompt injection methods, designed in Section 3.5, are shown in Table 5. For grouping, we require the LLM to focus on higher relevant triples, thereby aiding in narrowing the search scope. For ranking, we expect the LLM to prioritize the foremost information in a sequence of triples. For scoring, we hope that the score will assist LLMs in retrieving relevant triples from the knowledge.

We observed that the performance of various prompt strategies is inconsistent across different models. ChatGPT favors the knowledge injection

method that incorporates confidence scores, while the Vicuna series exhibits a preference for ranking methods. This discrepancy may be attributed to variations in the training data and inherent tendencies of the respective models. This finding indicates that when designing prompting methods, the applicability of a given lightweight method across multiple models should be considered. Additionally, it is essential to conduct comprehensive experiments to demonstrate the effectiveness of the prompts.

5 Conclusion

In this study, we make an extensive study on various input formats of KG injecting into the LLM and obtain key insights: for fact-intensive questions, LLMs prefer unordered linear triples over NL text. Indeed, by filtering out irrelevant information and focusing on details, LLMs can enhance accuracy beyond our initial expectations. Interestingly, unordered linear triple inputs can improve model performance even when the background knowledge does not contain information that directly aids reasoning. Moreover, the varying preferences on KG-related prompt strategies among different models should be a focal point for the community, necessitating thorough experimental validation when proposing prompt frameworks. These findings offer valuable guidance for refining KG-related prompt strategies and underscore the importance of linear triple knowledge in LLMs handling fact-related tasks.

6 Limitations

Despite extensive research on the understanding of LLMs to KGs, this study has some limitations.

Due to the resource cost associated with API calls and constraints on the length of the paper, it is difficult for us to evaluate the dataset designed in Table 1 using the latest LLMs. Although Table 2 suffices to illustrate the main arguments presented in this paper, a more granular subgraph expansion and KG input format would facilitate a deeper investigation into the current versions of LLMs and their attention to KG knowledge. In future work, we plan to employ multiple versions of the models to evaluate all our experiments comprehensively.

We utilize the QALD-7 dataset (Usbeck et al., 2017) in our research. However, the limited number of questions in this dataset may introduce bias into the evaluation results. Moreover, our study is restricted to datasets based on Wikidata, inherently limiting the choice of KGQA datasets. Nonetheless, as the largest continuously updated and developed knowledge base, Wikidata remains the optimal choice. Other KGs, including DBpedia (Lehmann et al., 2015) and YAGO (Pellissier Tanon et al., 2020), have ceased updates. In future work, we will consider evaluating our model on different KG-based datasets.

References

Simran Arora, Sen Wu, Enci Liu, and Christopher Ré. 2022. Metadata shaping: A simple approach for knowledge-enhanced language models. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1733–1745.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. *arXiv preprint arXiv:2310.11511*.

Jinheon Baek, Alham Fikri Aji, and Amir Saffari. 2023. Knowledge-augmented language model prompting for zero-shot knowledge graph question answering. In *Proceedings of the First Workshop on Matching From Unstructured and Structured Data*, pages 70–98.

Simone Balloccu, Patrícia Schmidová, Mateusz Lango, and Ondrej Dusek. 2024. Pleak, cheat, repeat: Data contamination and evaluation malpractices in closed-source LLMs. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics*, volume 1, pages 67–93.

Shulin Cao, Jiaxin Shi, Liangming Pan, Lunyu Nie, Yutong Xiang, Lei Hou, Juanzi Li, Bin He, and Hanwang Zhang. 2022. KQA pro: A dataset with explicit

compositional programs for complex question answering over knowledge base. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*, volume 1, pages 6101–6119.

Qianglong Chen, Feng-Lin Li, Guohai Xu, Ming Yan, Ji Zhang, and Yin Zhang. 2022. Dictbert: dictionary description knowledge enhanced language model pre-training via contrastive learning. *arXiv preprint arXiv:2208.00635*.

Wenhu Chen. 2023. Large language models are few(1)-shot table reasoners. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1120–1130.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Mohnish Dubey, Debayan Banerjee, Abdelrahman Abdelkawi, and Jens Lehmann. 2019. Lc-quad 2.0: A large dataset for complex question answering over wikidata and dbpedia. In *Proceedings of the 18th International Semantic Web Conference*, pages 69–78.

Hanning Gao, Lingfei Wu, Po Hu, and Fangli Xu. 2020. Rdf-to-text generation with graph-augmented structural neural encoders. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, pages 3030–3036.

Tiezheng Guo, Qingwen Yang, Chen Wang, Yanyi Liu, Pan Li, Jiawei Tang, Dapeng Li, and Yingyou Wen. 2023. Knowledgenavigator: Leveraging large language models for enhanced reasoning over knowledge graph. *arXiv preprint arXiv:2312.15880*.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys (Volume 55)*, 55(12):1–38.

Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Xin Zhao, and Ji-Rong Wen. 2023. 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*, pages 9237–9251.

Jens Lehmann, Robert Isele, Max Jakob, Anja Jentzsch, Dimitris Kontokostas, Pablo N Mendes, Sebastian Hellmann, Mohamed Morsey, Patrick Van Kleef, Sören Auer, et al. 2015. Dbpedia—a large-scale, multilingual knowledge base extracted from wikipedia. *Semantic web*, 6(2):167–195.

Shiyang Li, Yifan Gao, Haoming Jiang, Qingyu Yin, Zheng Li, Xifeng Yan, Chao Zhang, and Bing Yin. 2023. Graph reasoning for question answering with triplet retrieval. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3366–3375.

739	Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. K-bert: Enabling language representation with knowledge graph. In <i>Proceedings of the 34th AAAI Conference on Artificial Intelligence</i> , volume 34, pages 2901–2908.	795
740		796
741		797
742		
743		
744	Kaixin Ma, Hao Cheng, Xiaodong Liu, Eric Nyberg, and Jianfeng Gao. 2022. Open domain question answering with a unified knowledge interface. In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics</i> , volume 1, pages 1605–1620.	798
745		799
746		800
747		801
748		802
749		803
750	Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jipu Wang, and Xindong Wu. 2023. Unifying large language models and knowledge graphs: A roadmap. <i>IEEE Transactions on Knowledge and Data Engineering</i> , 36:3580–3599.	804
751		805
752		806
753		807
754		
755	Thomas Pellissier Tanon, Gerhard Weikum, and Fabian Suchanek. 2020. Yago 4: A reason-able knowledge base. In <i>Proceedings of the 17th International Conference</i> , pages 583–596.	808
756		809
757		810
758		
759	Pranab Sahoo, Ayush Kumar Singh, Sriparna Saha, Vinija Jain, Samrat Mondal, and Aman Chadha. 2024. A systematic survey of prompt engineering in large language models: Techniques and applications. <i>arXiv preprint arXiv:2402.07927</i> .	811
760		812
761		813
762		814
763		815
764		816
765	Melanie Sclar, Yejin Choi, Yulia Tsvetkov, and Alane Suhr. 2023. Quantifying language models’ sensitivity to spurious features in prompt design or: How i learned to start worrying about prompt formatting. <i>arXiv preprint arXiv:2310.11324</i> .	817
766		818
767		819
768		820
769	Priyanka Sen, Sandeep Mavadia, and Amir Saffari. 2023. Knowledge graph-augmented language models for complex question answering. <i>Proceedings of the 1st Workshop on Natural Language Reasoning and Structured Explanations</i> , pages 1–8.	821
770		822
771		823
772		824
773		825
774	Taylor Sorensen, Joshua Robinson, Christopher Rytting, Alexander Shaw, Kyle Rogers, Alexia Delorey, Mahmoud Khalil, Nancy Fulda, and David Wingate. 2022. An information-theoretic approach to prompt engineering without ground truth labels. In <i>Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics</i> , volume 1, pages 819–862.	826
775		827
776		828
777		829
778		830
779		
780		
781		
782	Yusheng Su, Xu Han, Zhengyan Zhang, Yankai Lin, Peng Li, Zhiyuan Liu, Jie Zhou, and Maosong Sun. 2021. Cokebert: Contextual knowledge selection and embedding towards enhanced pre-trained language models. <i>AI Open</i> , 2:127–134.	831
783		832
784		833
785		834
786		835
787		836
788	Tianyi Tang, Junyi Li, Wayne Xin Zhao, and Ji-Rong Wen. 2023. Mvp: Multi-task supervised pre-training for natural language generation. In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 8758–8794.	837
789		838
790		839
791		840
792		841
793	Yixuan Tang and Yi Yang. 2024. Multihop-rag: Benchmarking retrieval-augmented generation for multi-hop queries. <i>arXiv preprint arXiv:2401.15391</i> .	842
794		843
		844
		845
		846
		847
	Furkan Ufuk. 2023. The role and limitations of large language models such as chatgpt in clinical settings and medical journalism. <i>Radiology</i> , 307(3):230–276.	
	Ricardo Usbeck, Axel-Cyrille Ngonga Ngomo, Bastian Haarmann, Anastasia Krithara, Michael Röder, and Giulio Napolitano. 2017. 7th open challenge on question answering over linked data (qald-7). In <i>Proceedings of the 4th SemWebEval Challenge at ESWC 2017</i> , pages 59–69.	
	Anton Voronov, Lena Wolf, and Max Ryabinin. 2024. Mind your format: Towards consistent evaluation of in-context learning improvements. <i>arXiv preprint arXiv:2401.06766</i> .	
	Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: a free collaborative knowledgebase. <i>Communications of the ACM</i> , 57(10):78–85.	
	Jianing Wang, Wenkang Huang, Minghui Qiu, Qiuhui Shi, Hongbin Wang, Xiang Li, and Ming Gao. 2022. Knowledge prompting in pre-trained language model for natural language understanding. In <i>Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing</i> , pages 3164–3177.	
	Yilin Wen, Zifeng Wang, and Jimeng Sun. 2023. Mindmap: Knowledge graph prompting sparks graph of thoughts in large language models. <i>arXiv preprint arXiv:2308.09729</i> .	
	Jules White, Quchen Fu, Sam Hays, Michael Sandborn, Carlos Olea, Henry Gilbert, Ashraf Elnashar, Jesse Spencer-Smith, and Douglas C Schmidt. 2023. A prompt pattern catalog to enhance prompt engineering with chatgpt. <i>arXiv preprint arXiv:2302.11382</i> .	
	Yike Wu, Nan Hu, Guilin Qi, Sheng Bi, Jie Ren, Anhuan Xie, and Wei Song. 2023. Retrieve-rewrite-answer: A kg-to-text enhanced llms framework for knowledge graph question answering. <i>arXiv preprint arXiv:2309.11206</i> .	
	Guanming Xiong, Junwei Bao, Wen Zhao, Youzheng Wu, and Xiaodong He. 2022. Autoqgs: Auto-prompt for low-resource knowledge-based question generation from sparql. In <i>Proceedings of the 31st ACM International Conference on Information & Knowledge Management</i> , page 2250–2259.	
	Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. 2019. Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model. In <i>The 10th International Conference on Learning Representations</i> .	
	Yuan Yao, Deming Ye, Peng Li, Xu Han, Yankai Lin, Zhenghao Liu, Zhiyuan Liu, Lixin Huang, Jie Zhou, and Maosong Sun. 2019. DocRED: A large-scale document-level relation extraction dataset. In <i>Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics</i> , pages 764–777.	

848 Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu,
849 Mingxuan Ju, Soumya Sanyal, Chenguang Zhu,
850 Michael Zeng, and Meng Jiang. 2023. Generate
851 rather than retrieve: Large language models are
852 strong context generators. In *The 11th International
853 Conference on Learning Representations*.

854 Hanwen Zha, Zhiyu Chen, and Xifeng Yan. 2022. In-
855 ductive relation prediction by bert. In *Proceedings of
856 the 36th AAAI conference on artificial intelligence*,
857 pages 5923–5931.

858 Pengwei Zhan, Zhen Xu, Qian Tan, Jie Song, and
859 Ru Xie. 2024. Unveiling the lexical sensitivity of
860 llms: Combinatorial optimization for prompt en-
861 hancement. *arXiv preprint arXiv:2405.20701*.

862 Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang,
863 Maosong Sun, and Qun Liu. 2019. ERNIE: En-
864 hanced language representation with informative en-
865 tities. In *Proceedings of the 57th Annual Meeting of
866 the Association for Computational Linguistics*, vol-
867 ume 1, pages 1441–1451.

868 Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan
869 Zhuang, Zhonghao Wu, Yonghao Zhuang, Zi Lin,
870 Zhuohan Li, Dacheng Li, Eric Xing, et al. 2024.
871 Judging llm-as-a-judge with mt-bench and chatbot
872 arena. *Advances in Neural Information Processing
873 Systems*, 36:46595–46623.

874 A Appendix

875 A.1 Core Reasoning Path Generation

876 We first find all the reasoning paths in the "where"
877 based on the SPARQL⁴ query corresponding to a
878 question. Because there are variables in this rea-
879 soning path, we assign values to all variables that
880 meet the conditions through the SPARQL query. If
881 the answer to a question is an entity, then accord-
882 ing to our query assignment method, the answer
883 must exist in the reasoning path. For example, for
884 the question "Who are the actors in The Big Bang
885 Theory?" we will get multiple reasoning paths (be-
886 cause it is a simple question, there is only one triple
887 in the reasoning path): (The Big Bang Theory, cast
888 member, Jim Parsons), (The Big Bang Theory, cast
889 member, Johnny Galecki), (The Big Bang Theory,
890 cast member, Kaley Cuoco), etc. These answers
891 exist in the reasoning path.

892 We employ a multi-hop question to demonstrate
893 our approach. For the question: "Which is the
894 island nation for the country of pound sterling?"
895 The corresponding SPARQL query is articulated as
896 follows: `select distinct ?obj where { wd:Q25224
897 wdt:P17 ?obj . ?obj wdt:P31 wd:Q112099 }`

⁴https://www.wikidata.org/wiki/Wikidata:SPARQL_tutorial

898 In this query, "wd:Q25224" represents 898
899 "pound sterling," "wdt:P17" represents "country," 899
900 "wdt:P31" represents "instance of," and 900
901 "wd:Q112099" represents "island country." 901
902 Thus, the query seeks to find entities that are 902
903 both countries using the pound sterling and 903
904 island nations. Upon execution, we extract all 904
905 pertinent variables from the query and assign them 905
906 appropriate values, such as attributing "United 906
907 Kingdom" to "?obj." Consequently, we derive the 907
908 logical sequence: [`<pound sterling>` `<country>`
909 `<United Kingdom>` , `<United Kingdom>` `<instance`
910 `of>` `<island country>`]. Then we can query the 910
911 neighbors and expand the entity nodes [`<pound`
912 `sterling>`, `<United Kingdom>`, `<island country>`]
913 in this reasoning path according to the method
914 described in Section 3.1.

915 A.2 Question Generation From DocRED

916 We employ ChatGPT to re-express (paraphrase) 916
917 all questions, which were initially devised using 917
918 manual templates. ChatGPT was not involve in the 918
919 direct generation of multi-hop questions or in 919
920 answering them during the dataset construction phase. 920
921 We construct some question templates for the multi- 921
922 hop scenario and designate the terminal node of the 922
923 multi-hop path as the answer, thereby establishing 923
924 the question-answer pair before ChatGPT's involve- 924
925 ment. To guarantee the diversity and fluency of the 925
926 questions, ChatGPT is utilized to rephrase the pre- 926
927 generated template questions. Subsequently, we 927
928 conduct manual reviews of the paraphrased outputs 928
929 to ensure the high quality of the question-answer 929
930 pairs. An example is shown in Table 7.

931 A.3 Literal Level Analysis

932 We first asked ChatGPT about its preference for 932
933 knowledge type.

- 934 • **Q:** Which external knowledge would you 934
935 choose to help you accurately answer the ques- 935
936 tion: unordered structured triples or fluent 936
937 text? 937
- 938 • **LLM Response:** As an AI language model, 938
939 I don't have personal preferences. However, 939
940 when it comes to accurately answering ques- 940
941 tions, both unordered structured triples and 941
942 fluent text can be useful sources of informa- 942
943 tion depending on the nature of the question 943
944 and the available data. When receiving struc- 944
945 tured triples, I can directly search for and re- 945
946 trieve the relationships between relevant en- 946

	Reasoning Path	Template Question	Answer	Paraphrased Question
1-hop	[Gaozong, date of death, 1187]	What is the date of death of Gaozong?	1187	When did Gaozong die?
2-hop	[Lin'an, located in the administrative territorial entity, China], [China, contains administrative territorial entity, Bianjing]	What administrative territorial entity are contained in the administrative territorial entity where Lin'an is located?	Bianjing	Which administrative territorial entities are encompassed within the administrative territorial entity that includes Lin'an?
3-hop	[Huizong, child, Qinzong], [Qinzong, sibling, Gaozong], [Gaozong, date of birth, 12 June 1107]	What is the date of birth which is the sibling of the Huizong's child?	12 June 1107	When was Huizong's child's sibling born?

Table 7: 1-hop, 2-hop, and 3-hop questions generated from a document.

947 titles and attributes, thereby providing more
948 targeted answers. For instance, if I know a
949 question pertains to a specific attribute of a
950 particular entity, I can directly extract relevant
951 information to answer the question. However,
952 textual input has its own advantages. Text can
953 contain more contextual information and se-
954 mantic details, which might be more helpful
955 for open-ended questions and those requiring
956 a deeper understanding of the context. More-
957 over, textual input can also include narratives,
958 discourses, and descriptions, all of which are
959 crucial for providing more comprehensive and
960 detailed answers. Ultimately, the choice be-
961 tween unordered structured triples and fluent
962 text depends on the specific question and the
963 type of information required.

964 We gave ChatGPT the question and the related
965 knowledge and asked it to choose one of the knowl-
966 edge formats as the answer.

- 967 • **Q:** I can provide you with two kinds of knowl-
968 edge: unordered triple structured knowledge
969 or more fluent textual knowledge. In order to
970 accurately answer the question "Who lives in
971 New York City and has ISNI 0000 0001 1475
972 6720?" you must choose one of them and tell
973 me why. (**Knowledge...**)
- 974 • **LLM Response:** To accurately answer the
975 question "Who lives in New York City and
976 has ISNI 0000 0001 1475 6720?" the most
977 suitable choice would be the unordered triple
978 structured knowledge. This is because the
979 question requires specific information about
980 an individual's residence in New York City
981 and their associated ISNI (International Stan-
982 dard Name Identifier) number. Unordered
983 triple structured knowledge allows for precise
984 retrieval of relevant facts, such as matching
985 the ISNI with the individual's location. Flu-
986 ent textual knowledge, on the other hand, is

987 more adept at generating coherent and con-
988 textually relevant responses based on a given
989 prompt. It excels at understanding and gen-
990 erating NL text but may not be as effective
991 in retrieving specific structured information
992 like the one required in this case. By utiliz-
993 ing unordered triple structured knowledge, we
994 can directly access the required information
995 about individuals residing in New York City
996 and cross-reference it with their associated
997 ISNI numbers to provide an accurate answer
998 to the question.

999 A.4 Attention Distribution Level Analysis

1000 The last token of the input sequence is considered
1001 the model's predicted label. As illustrated in fig-
1002 ure 4, we calculate the average attention proportion
1003 between the last token (i.e., the colon ":") and the
1004 answer words (highlighted in gray). We observe
1005 that, in both input modes, unordered triples exhibit
1006 a higher proportion of attention weight. In our ex-
1007 periments, we evaluate all data in the dataset and
1008 compute the average values, finding that the overall
1009 observations are consistent with this example.

1010 A.5 Scorer Training and Prompt Strategies

1011 We utilize the BERT-Base (Devlin et al., 2018) as
1012 the foundation for our model. The training data
1013 comprises triples and corresponding questions. We
1014 partition the dataset into training and test sets in
1015 an 8:2 ratio. Triples in the reasoning path linked
1016 to the question are labeled as positive examples;
1017 otherwise, they are designated as negative exam-
1018 ples. For the cross-encoder, the batch size is set
1019 at 50, we experimented with initial learning rates
1020 of {5e-4, 2e-5, 5e-5, 2e-5}, and the learning rate
1021 decays every 3 epochs. We set the multiplicative
1022 factor, gamma, for updating the learning rate to 0.2.
1023 Upon training the model, it exhibits an accuracy
1024 of 98.89% in determining whether triples are per-
1025 tinent to the question, i.e., whether they are part
1026 of the crucial reasoning path. This cross-encoder
1027 is employed to assign scores to the questions and

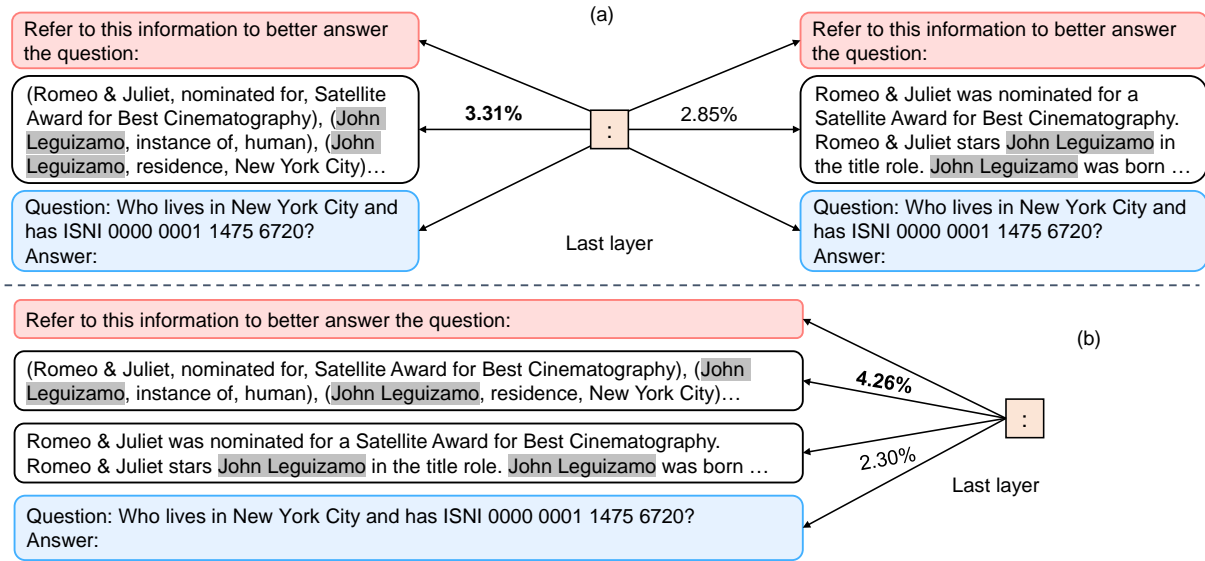


Figure 4: There are two input modes: in (a), we eliminate the interference between the two formats of knowledge and provide the LLM with only one type of knowledge separately. In (b), we provide the LLM with both types of knowledge simultaneously to observe the extent to which the LLM focuses on the answer words from different types of knowledge.

1028 their associated triples. By setting thresholds of
 1029 0.3 and 0.8 at either end to segment the triples, we
 1030 can categorize them into high, medium, and low
 1031 correlation groups relative to a question.

1032 After the assignment of scores to triples by the
 1033 cross-encoder, Table 8 illustrates the organizational
 1034 format of the structured knowledge to the ques-
 1035 tion "What trade structure did Straight to the point
 1036 Gehry design?". This part is incorporated into the
 1037 prompt as knowledge. For details on the incorpora-
 1038 tion format, refer to Appendix A.6.

1039 A.6 LLM Input and Metric Setting

1040 As illustrated in Figure 5, the input to the LLM is
 1041 primarily divided into four components: task in-
 1042 struction, in-context learning examples, external
 1043 knowledge, and questions. Within the instruction
 1044 section, we impose constraints on the model’s out-
 1045 put mode, mandating that different questions must
 1046 be answered according to the prescribed format.
 1047 For example, if the answer pertains to an entity,
 1048 we stipulate that the answer should be presented
 1049 as a list of entities devoid of any explanations. For
 1050 counting questions, only numerical values are per-
 1051 missible. Unanswerable questions are indicated by
 1052 returning *None*. This specification ensures that the
 1053 format of the model’s responses consistently aligns
 1054 with our unified evaluation process.

1055 During the experiment, we observed some cases
 1056 where the LLM’s replies still failed to meet our

Injection Method	Knowledge Prompt
Grouping	<p>Here are some triples that are highly relevant to the question: (DZ Bank building, architect, Frank Gehry), (Gehry Tower, instance of, office building), ... Here are some triples that are likely relevant to the question: (IAC Building, architect, Frank Gehry), (Gehry Tower, architect, Frank Gehry) ... Here are some triples that are less relevant to the question: (Toledo Museum of Art, architect, Frank Gehry), (Vlado Miluni, notable work, Dancing House), ...</p>
Ranking	<p>The triples are sorted from high to low according to their relevance score to the question for your reference: (DZ Bank building, architect, Frank Gehry), (Dancing House, instance of, office building), (Gehry Tower, architect, Frank Gehry), (Dancing House, architect, Frank Gehry), (IAC Building, instance of, office building), ...</p>
Scoring	<p>You can refer to these information to better answer the question. Each triple is followed by a confidence score of its relevance to the question, which helps in solving the question: {(DZ Bank building, architect, Frank Gehry) 0.9981}, {(Toledo Museum of Art, architect, Frank Gehry) 0.0019}, {(Gehry Tower, instance of, office building) 0.998}, {(Vlado Miluni, notable work, Dancing House) 0.0023}...</p>

Table 8: Prompt organization.

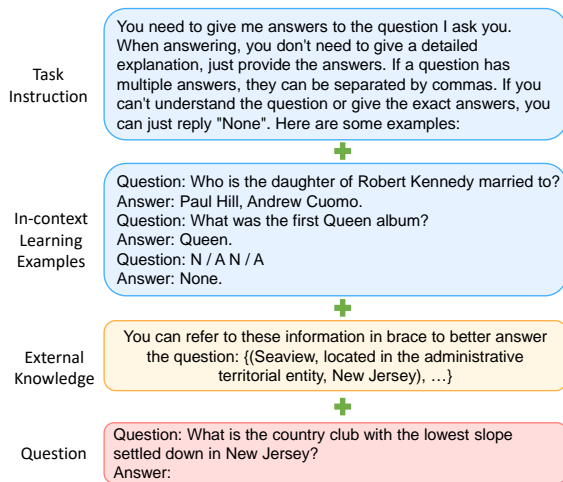


Figure 5: The structure of the input. The knowledge component is provided in triple format, and alternative formats (such as meta path or sentence) can be used to represent the knowledge.

from similar knowledge, even though the external knowledge does not proffer direct answers. As illustrated in rows three and four of Table 9, 90% of the triples in the knowledge we gave have been replaced by wrong entities, which contain a lot of noisy knowledge. The model can provide correct answers when responding directly, however, the introduction of erroneous external knowledge leads to incorrect responses from the model. This indicates that the model lacks robustness against noisy information and is significantly influenced by the introduction of external inaccuracies.

1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099

criteria, occasionally returning sentences or blocks of text. To address this issue, we refer to the concept of in-context learning, wherein we provided multiple examples resembling the question to guide the model towards producing responses in the desired format. Regarding the knowledge injection method, we augmented the prompt with text derived from the KG, encompassing various formats, and supplied it to the LLM as additional knowledge. Finally, we appended the NL question that requires answers at the end of the model input.

In evaluation, our approach employs in-context learning within prompts, mandating that the LLM responses adhere to a strict format, such as providing only entity names. This requirement significantly streamlines our evaluation process. Therefore, we re-evaluated the answers of some samples based on experience and found that using the `difflib.SequenceMatcher()` function in the Python library under the threshold of 0.7 can better complete the answer validation.

A.7 Qualitative Example Study

We employ ChatGPT as our baseline model to elucidate the influence of external knowledge on the reasoning capability of the model. As shown in the first and second rows of Table 9, the model tends to commit errors when directly responding to the questions. Nevertheless, upon integrating a sub-graph devoid of inference paths, the model succeeds in providing correct responses. This is attributable to the model's capacity to draw analogies

Question	Knowledge	Without Knowledge Answer	With Knowledge Answer	Correct Answer
Which is the calling for the field of occupation of manga?	without reason path: (Jack Bauer, occupation, field agent), (field agent, instance of, profession), (Comic Gum, intended public, seinen), (manga, subclass of, manga), (seinen, subclass of, manga), (comedian, occupation, comedian), (Son Gokb, creator, Akira Toriyama), (field agent, occupation, field agent), (Nozomu Tamaki, occupation, mangaka), (Dragon Ball GT, after a work by, Akira Toriyama), (Akira Toriyama, occupation, mangaka), (seinen, intended public, seinen), (Douglas Adams, occupation, comedian), (comedian, instance of, profession)	Illustration, storytelling, comic art.	mangaka	mangaka
Is it true that David Koresh's given name was David or Wayne?	without reason path: (Waco siege, instance of, religious persecution), (Wayne Shorter, given name, Wayne), (David Bowie, given name, David), (Dave Arneson, given name, David), (religious persecution, statement supported by, David Koresh)	False.	True.	True.
Which is the island country for the nation of pound sterling?	90% replace: (Germany, diplomatic relation, Bahrain), (South Holland, contains the administrative territorial entity, Nieuw-Lekkerland), (Antwerp, twinned administrative body, Rotterdam), (Nieuw-Lekkerland, contains the: administrative territorial entity, Nieuw-Lekkerland), (Nieuw-Lekkerland, contains the administrative territorial entity, Nieuw-Lekkerland), (Nieuw-Lekkerland, instance of, village), (Antwerp, twinned administrative body, Rotterdam), (South Holland, contains the administrative territorial entity, Spijkenisse), (Nieuw-Lekkerland, instance of, village), (South Holland, contains the administrative territorial entity, Rijnsburg), (South Holland, contains the administrative territorial entity, Nieuw-Lekkerland), (South Holland, contains the administrative territorial entity, Rijnsburg), (Rijnsburg, contains the administrative territorial entity, Rijnsburg), (Nieuw-Lekkerland, instance of, village), (Rijnsburg, instance of, village), (European Netherlands, has part(s), South Holland)	United Kingdom.	None.	United Kingdom.
What is the inverse class for fiction?	90% replace: (The Night Watch, genre, historical fiction), (Gerry Adams, position held, Mary Lou McDonald), (Sinn Fin, chairperson, Mary Lou McDonald), (Sinn Fin, chairperson, Mary Lou McDonald), (Gerry Adams, position held, Mary Lou McDonald), (Martin McGuinness, member of political party, Sinn Fin), (Gerry Adams, position held, Mary Lou McDonald), (Sinn Fin, chairperson, Mary Lou McDonald), (Lynn Boylan, member of political party, Sinn Fin), (2001 United Kingdom general election, followed by, 2005 United Kingdom general election), (Martin McGuinness, candidacy in election, 2005 United Kingdom general election), (Martin McGuinness, member of political party, Sinn Fin), (2005 United Kingdom general election, candidate, Sinn Fin), (Mary Lou McDonald, replaced by, Mary Lou McDonald), (Martin McGuinness, member of political party, Sinn Fin), (Lynn Boylan, member of political party, Sinn Fin), (Martin McGuinness, member of political party, Sinn Fin), (Martin McGuinness, member of political party, Sinn Fin), (Gerry Adams, position held, Mary Lou McDonald), (Martin McGuinness, member of political party, Sinn Fin)	nonfiction.	historical fiction.	Non-fiction.

Table 9: The impact of external knowledge on LLM inference performance.