

MRKE: The Multi-hop Reasoning Evaluation of LLMs by Knowledge Edition

Anonymous ACL submission

Abstract

Although Large Language Models (LLMs) have shown strong performance in Multi-hop Question Answering (MHQA) tasks, their real reasoning ability remains exploration. Current LLM QA evaluation benchmarks have shown limitations, including 1) data contamination, the evaluation data are potentially exposed to LLMs during the pretraining stage; and 2) ignorance of the reasoning chain evaluation. Thus we introduce an LLM MHQA evaluation benchmark, the first QA benchmark based on the new, unprecedented knowledge by editing the off-the-shelf HotpotQA dataset; Besides, we also annotate and evaluate the reasoning chain in the form of sub-questions and intermediate answers corresponding to the multi-hop questions. Specifically, based on the observation, 1) LLMs show a performance gap between the original HotpotQA and our edited data, deeming that current MHQA benchmarks have the potential risk of data contamination that hard to evaluate LLMs' performance objectively and scientifically; 2) LLMs only get a small percentage of the right reasoning chain, e.g. GPT-4 only gets 36.3% right reasoning chain. We believe this new Multi-hop QA evaluation benchmark and novel evaluation methods will facilitate the development of trustworthy LLM evaluation on the MHQA task.

1 Introduction

Recent advancements in large language models (LLMs) (Zhao et al., 2023) have dramatically transformed the landscape of natural language processing(NLP), showcasing remarkable capabilities across a spectrum of tasks such as natural language understanding (Bai et al., 2023; Hendrycks et al., 2020; Zhong et al., 2023a) and question answering (QA) (Wang et al., 2023; Chen et al., 2023; Zhang et al., 2023; Zhuang et al., 2023). Notably, Retrieval-Augmented Generation (RAG) could help LLM integrate external data retrieval into the generative process, thereby enhancing the

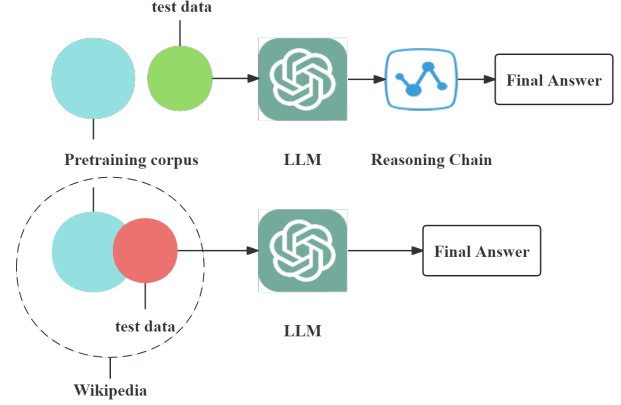


Figure 1: The differences between our benchmark(top) and previous benchmark(bottom). Our benchmark eliminates the risk of data contamination and evaluates the LLMs' reasoning chain.

model's ability to provide accurate and relevant responses (Gao et al., 2023b).

Multi-step reasoning, which necessitates intricate step-by-step reasoning across multiple texts, highlights the sophisticated inferential abilities of LLM RAG systems and the importance of the reasoning chain. For instance, in Figure 2, the left part of the figure is an example from HotpotQA dataset(Yang et al., 2018) that QA systems are required to extract evidence from multiple paragraphs and reasoning the final answer. To answer the question "The Oberoi family is part of a hotel company that has a head office in what city?", LLMs are required to answer the first sub-question "The Oberoi family is part of what hotel company?" by extracting the first evidence "Oberoi Group" and then answering the second sub-question "The head office of Oberoi Group is based in what city?" through retrieving the second evidence "Delhi" and finally reach the correct final answer.

Although traditional MHQA datasets e.g. HotpotQA (Yang et al., 2018) and 2WikiMulti-hopQA(Ho et al., 2020) could be applied to evaluate the multi-step reasoning ability of LLM RAG

systems, the bottom of Figure 1 shows the drawbacks of traditional MHQA benchmarks: 1) retrieving evidence is not sufficient for us to interpret the reasoning ability of LLMs. It is unclear whether the LLMs have performed the desired reasoning chain to reach the correct answer. 2) traditional MHQA benchmarks such as HotpotQA, 2Wiki-MultihopQA, and MuSiQue (Trivedi et al., 2022) collected and annotated passages from Wikipedia, facing the challenge of data contamination with LLMs’ pretraining corpus.

To address these issues, we propose MRKE (Multi-hop Reasoning chain Evaluation by Knowledge Editing), a novel benchmark designed for a more complete evaluation of LLM reasoning ability based on multi-hop QA tasks. The top of Figure 1 contrasts our MRKE with the HotpotQA dataset, which eliminates the data contamination risk and evaluates LLMs’ reasoning ability by including reasoning chain evaluation as well as final answer evaluation. We also propose a new evaluation metric in a fair setting that jointly considers intermediate answers equally important to the final answer.

To our knowledge, we are the first to introduce the multifaceted performance of LLMs on both the original HotpotQA and MRKE and evaluate the performance gap between original data and re-annotated data. We show that GPT-4 gets 69.3 EM and 82.2 F_1 scores on the original HotpotQA dataset respectively but only gets 53.2 EM and 67.7 F_1 scores on MRKE. Moreover, we jointly evaluate the intermediate reasoning chain performance of LLMs. GPT-4 gets 0.7 joint F_1 RC and 1.5 joint EM RC scores respectively, indicating that LLMs’ reasoning ability remains to be improved. LLMs although show strong performance on QA tasks, they usually bypass the right reasoning chain. Meanwhile, we reveal that LLMs also get an inflated performance due to the low proportion of right reasoning chains especially GPT-4 only gets 36.3% right reasoning chains among the whole dataset. Finally, we find that, combining the sub-question into a prompt as the reasoning chain is a more efficient approach for improving model performance, highlighting its significance as a future research direction for improving LLMs’ reasoning ability.

The original HotpotQA data and MRKE are currently available at https://anonymous.4open.science/r/LLM_multihop_eval-055C.

2 Related Work

RAG improves LLM’s response (Borgeaud et al., 2021) and also mitigates the occurrence of hallucinations, thereby enhancing the models’ credibility (Gao et al., 2023a). As demonstrated by Khattab et al. (2021), designs a RAG system for multi-hop question answering and claim verification tasks. These tasks require the extraction of evidence from two or more documents to produce a correct answer. Tang and Yang (2024) proposes a Multihop-RAG benchmark, which consists of a large collection of multi-hop queries, ground-truth answers, and the corresponding supporting evidence. Multihop-RAG requires LLM to reason and answer multi-hop queries given the evidence.

Multi-hop QA requires more than one reasoning step in multiple paragraphs to answer a question (Dua et al., 2019; Yang et al., 2018; Ho et al., 2020; Trivedi et al., 2022). Notably, Tang et al. (2021) introduced a human-validated sub-question dataset derived from the HotpotQA corpus (Yang et al., 2018), undertaking a detailed investigation into models’ capabilities to reason through sub-questions. Their findings revealed that notable models like DFGN (Xiao et al., 2019), DecomprC (Min et al., 2019), and CogQA (Ding et al., 2019) exhibit deficiencies in resolving sub-questions, even though they may successfully address the overarching multi-hop question.

The traditional MHQA datasets face the challenge of data contamination that hard to objectively and truthfully evaluate the reasoning ability of LLMs. Data contamination, i.e., the presence of test data from downstream tasks in the training data of large language models (LLMs), is a major issue in measuring LLMs’ real performance on other tasks. A handful of recent studies have provided several strategies, methods, and benchmarks for detecting contamination without needing to access pre-training data (Shi et al., 2023; Roberts et al., 2023; Golchin and Surdeanu, 2023). However, these data contamination detection benchmarks are required to dynamically update with the development of LLMs and the expansion of pertaining data. Dynamical maintenance is time-consuming and effortless, while our proposed benchmark MRKE, based on the knowledge edition, is fixed and maintains the cleanness of the test data.

The knowledge edition methods focus on editing the input knowledge e.g. prompt, data (Zheng et al., 2023). COUNTERFACT (Meng et al., 2022)

Multihop Question:

Who was the captain of Argentine team that was born in 1987?

Raw Context

Lionel Andrés Messi (born 24 June 1987) also known as Leo Messi, is an Argentine professional footballer who plays as a forward for and captains both Major League Soccer club Inter Miami and the Argentina national team.

Sub Question 1:

Who was the captain of Argentine?
Answer: Lionel Messi, Maradona ...

Sub Question 2:

Which captain of Argentine was born in 1987?
Answer: Lionel Messi

Final Answer:

Lionel Messi

Multihop Question:

Who was coach of the Brazil national team that was born in 1997?

Replaced & Paraphrased Context

Alice Brown (born 12 April 1997), also known as Brown is a Brazil vocational basketballer whose position was point guard and coaches both Houston Rockets and the Brazil national team.

Sub Question 1:

Who was the coach of Brazil national team?
Answer: Alice Brown

Sub Question 2:

Which Brazil coach was born in 1997?
Answer: Alice Brown

Final Answer:

Alice Brown

Figure 2: Example of the annotation on raw context, multi-hop questions, sub-questions, and intermediate answers. The words in red in the right part of the figure are the replacement results and the passage is the new, unprecedented knowledge that has never appeared before.

is a dataset specifically designed for editing a variety of counterfactual knowledge. The target of the dataset is to assess whether the model can provide counterfactual answers when asked about the corresponding factual knowledge. MQAKE (Zhong et al., 2023b) particularly focuses on evaluating whether edited models can answer multi-hop questions where the answer should change necessarily. Gu et al. (2023) proposed a programmable knowledge editing MHQA benchmark with decomposed knowledge-augmented multihop questions, to check the conflict signal via comparing with edited facts.

3 Proposed Benchmark and Evaluation method

In this section, we describe our data statistics, and process for annotating passages, multi-hop questions, sub-questions, intermediate answers, final answers, and evaluation methods for MRKE.

3.1 Data Construction

We randomly collect 300 data from the HotpotQA dataset (Yang et al., 2018) where the passages are from Wikipedia¹. Then, inspired by recent studies on LLM’s ability to aid human annotation (Bartolo

et al., 2021; Törnberg, 2023). we design a framework to efficiently replace the named entities, noun phrases, and synonyms and then do the back translation by introducing LLM and human evaluation. After getting the paraphrased passages, we again generate multi-step complex reasoning questions by incorporating LLMs and human feedback. After obtaining the passages, questions, sub-questions, intermediate answers, and final answers, we ask human experts to do the data review and refinement.

As shown in Figure 3, the framework has two main steps: 1) passage rewriting followed by human evaluation and feedback to enhance them based on guidelines (No grammar issues and not in the real world); 2) Multihop Question Generation.

Understanding and answering multi-hop questions following the right reasoning chain to offer a trustworthy reasoning process is a key part of the QA evaluation task. Figure 2 shows an example of original data and annotated new, unprecedented data. Given a raw context and corresponding multi-hop question, we first do the named entity, noun phrase, and synonym replacement. Then we translate the replaced text into Chinese and finally back translation into English. The words in red of the figure are the replaced named entities, noun phrases, and synonyms. Given the new context, we also generate a new multi-hop question to fit the new

¹<https://www.wikipedia.org/>

The Framework of Automatic LLM annotation and human evaluation

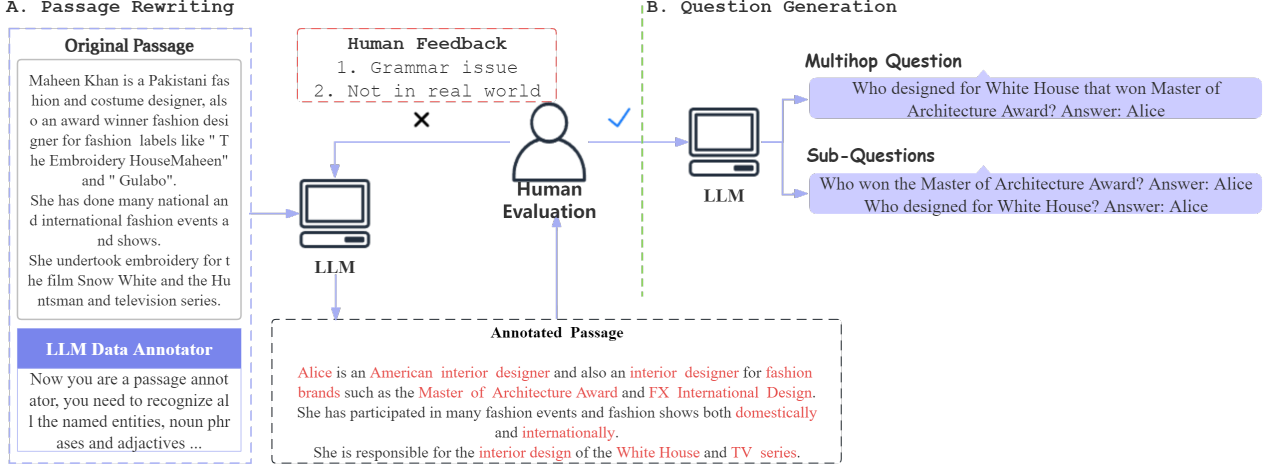


Figure 3: The framework of our LLM automatic data annotation. From left to right, we first ask LLM as a passage rewriter to do the replacement and paraphrasing. Then we do the human review to check the grammar issue and whether the knowledge has appeared or not. Finally, we send the reviewed high-quality data to GPT-4 to generate QA pairs.

context. In the real world, "*Lionel Messi*" is a football player who was born in "*1987*" and has also been the "*captain of the Argentina team*". Such knowledge is published in the real world before January 2022 and ChatGPT and other LLMs are prone to be exposed to the knowledge.

We also manually decompose the newly generated multi-hop question into several single-hop questions. E.g. in the right part of the figure 2 are the re-annotated context, newly generated multi-hop question, decomposed sub-questions, and intermediate answers. The newly generated multi-hop question is "*Who was the coach of the Brazil team that was born in 1997?*" and it is a two-hop question. The decomposed sub-questions are "*Who was the coach of the Brazil national team?*" and "*Which Brazil coach was born in 1997?*". To answer the multi-hop question correctly, LLMs are required to not only give the right final answer but also give the right intermediate answers to the related sub-questions. Otherwise, LLMs are following a wrong reasoning chain. we asked the two questions mentioned before on the ChatGPT and the figure is shown in Appendix A which also offers evidence that LLMs sometimes get answers relying on memory rather than reasoning.

Passage Rewriting Inspired by MQuAKE (Zhong et al., 2023b), which edits the reasoning chain by editing one or a few facts in a reasoning chain. we use LLM (i.e., GPT-3.5, GPT-4) as the passage rewriter, which largely reduces the cost

of data annotations. Specifically, we design an instructional prompt containing requirements for replacing named entities, noun phrases, synonyms, and back translation.

However, it is difficult for LLM to continuously output high-quality annotation results. After automatic passage rewriting, we do the human evaluation and feedback.

Human Evaluation and Feedback To make the rewritten passages clean and new, after LLM automatic annotation, we designed 2 rules as evaluation guidelines: 1) No grammar issues, make sure the annotated passage is right and can be understood by LLM; 2) The annotated passages haven't appeared on the Internet or other data source before.

The first rule is relatively easy and we ask English major students to do the check. While the second rule is hard, we randomly select and search the re-annotated passages on the Internet especially in Wikipedia to make sure they haven't appeared before.

Question Generation We leverage GPT-4 to automatically generate questions given a rewritten passage, because (1) this provides us with more diverse information of good quality; (2) it is challenging to manually write question templates for all the different types. We prompt GPT-4 to generate three questions (2, 3, 4-hop questions) and examples of generated multi-hop questions in table 2. Finally, our MRKE contains 900 data, 300 2-hop QA pairs, 300 3-hop QA pairs, and 300 4-hop QA

Property	Value
Total Data	900
Unique Passages	300
2 hop Questions	300
3 hop Questions	300
4 hop Questions	300
Sentences per data (Avg/Median)	38.42
Inter-annotator Agreement	94%

Table 1: The statistics of MRKE.

pairs respectively. The prompts of passage rewriting and question generation are shown in Appendix B.

3.2 Human Agreement on MRKE

We sample 300 instances randomly from MRKE (100 for 2-hop, 100 for 3-hop, and 100 for 4-hop) and assign each instance to two paper authors (as expert human annotators). We ask them to select the annotated data that badly follows the given guidelines. The experts first checked the grammar issues of the annotated questions, sub-questions, sub-answers, and answers. Then search the knowledge of the annotated passages on the Internet to see whether they have appeared or not. The agreement rate between expert annotators on the sampled MRKE set is 94%. The human agreement rates are 96%, 92%, and 95% on the 2-hop, 3-hop, and 4-hop datasets respectively. This suggests that MRKE instances reflect a good data quality on annotation guideline following and achieve high human agreement among expert annotators.

3.3 Dataset Analysis and Statistics

Followed by benchmarks such as HotpotQA (Yang et al., 2018), 2WikiMulthopQA (Ho et al., 2020), we propose a taxonomy on fine-grained question types and examples commonly used in multi-hop QA illustrated in table 2.

3.4 Multi-hop QA Evaluation

The task of multi-hop QA evaluation is referred to as finding answers for complex questions that require reasoning multiple times from given passages.

We employ three representative QA evaluation methods to assess the correctness of LLM-generated MHQA responses: sub-question answering evaluation, reasoning chain evaluation, and the joint performance of sub-qa and MHQA.

Sub-Question Answering Evaluation This part is the basis of the whole experiment and all evaluation results. Following reading comprehension (Rajpurkar et al., 2016), evaluation is conducted through lexical matching using two widely used metrics to assess the performance of models. In this section, we employ F_1 and EM scores to evaluate the answers to sub-questions, similar to the single-hop QA task.

Reasoning chain evaluation To interpret the behavior of existing LLMs on each hop of the reasoning process required for multi-hop questions and to determine their reasoning ability to answer simple questions. We followed the experiment setting proposed by Tang et al. (2021). For example, in the 2-hop dataset, we have 300 2-hop questions, each 2-hop question has 2 sub-questions, 2 intermediate answers, and a final answer. In order to understand whether LLMs can correct answers by following the right reasoning chain, we calculate the proportion of right and wrong reasoning chains to compare LLMs’ reasoning performance.

Each question or sub-question has two results, correct or wrong, thus an N-hop question with its N sub-questions has $2^{(N+1)}$ different reasoning chains. Due to the space limitation, we measure and collect correctness statistics for the 2-hop question dataset, q_{sub1} , q_{sub2} , and q , and show the percentage of 8 reasoning chains given by LLMs.

The joint performance of Sub-QA and Multi-hop QA The previous MHQA benchmarks were traditionally evaluated on the EM or F_1 score on the final answer (Rajpurkar et al., 2016; Yang et al., 2018; Ho et al., 2020), which is partially correct. The previous MHQA systems and LLMs are treated as a black box and we can not figure out how they find the final answer. Hence, the final answer evaluation shows limitations as it does not consider whether previous MHQA systems could answer sub-questions correctly or not. To understand the impact of sub-qa on MHQA, we introduce a joint performance that combines the evaluation of Sub-QA performance and MHQA performance. For example, for a N-hop question and its N sub-questions, given their precisions and recalls on the MHQA ($P^{(MHQA)}$, $R^{(MHQA)}$) and the Sub-QA ($P^{(sub_qa^1)}$, $R^{(sub_qa^1)}$), ... ($P^{(sub_qa^N)}$, $R^{(sub_qa^N)}$), respectively, we calculate joint performance as:

$$P^{(joint)} = P^{(MHQA)} P^{(sub_qa^1)} \dots P^{(sub_qa^N)},$$

$$R^{(joint)} = R^{(MHQA)} R^{(sub_qa^1)} \dots R^{(sub_qa^N)},$$

Question Type	Hop	Multi-hop Question
Bridge	2 hop	<i>When</i> was the actor <i>who</i> played Helen in FBC series The Murder born?
	3 hop	<i>Who</i> were the learners of the people <i>that</i> was the principal violist in the Fioba Symphony Band <i>and</i> instructed music to Michard Rokney?
	4 hop	<i>Which is later</i> , the birthday of Zephyr Bolt-Anderson <i>or</i> the time <i>that</i> 2060 Kingdom of Azkaban ATP Conqueror occurred in Gleeful Peak, Atlantis?
Comparison	2 hop	Where is the Blue Falls Empire located <i>and</i> what products are it responsible for importing?
	3 hop	<i>Which is later</i> , the opening time of Gold <i>or</i> the opening time of the Mad Book in 2006?
	4 hop	<i>Was</i> the release of the movie Ocean Secrets <i>before or after</i> Echoes of Tomorrow & Victoria Wright?

Table 2: Examples of MRKE with different question types and question hops. We emphasize keywords for their respective categories.

$$\text{Joint } F_1 \text{ } RC = -\log \frac{2P^{(\text{joint})} R^{(\text{joint})}}{P^{(\text{joint})} + R^{(\text{joint})}}.$$

where the Joint F_1 RC means the joint F_1 performance of the reasoning chain.

Given their EM scores on the MHQA ($EM^{(\text{MHQA})}$) and the Sub-QA ($EM^{(\text{sub_qa}^1)}$), ... $EM^{(\text{sub_qa}^N)}$.

$$\text{Joint EM } RC = -\log \frac{2EM^{(\text{MHQA})}, \dots EM^{(\text{sub_qa}^N)}}{EM^{(\text{MHQA})} + \dots EM^{(\text{sub_qa}^N)}}.$$

where the Joint EM RC means the joint EM performance of the reasoning chain.

4 Experiments

We conduct comprehensive experiments and evaluate different LLMs on MRKE to answer the following questions: 1) Do LLMs show a performance gap between the original HotpotQA dataset and our MRKE? 2) When separating the memory of LLMs, how do LLMs show their reasoning ability? 3) How do sub-questions affect the performance of LLMs? (4) How do LLMs perform on reasoning chain evaluation?

4.1 Experiment Setup

Datasets We evaluate LLMs on the selected 300 original HotpotQA data and our 900 MRKE dataset (divided into 2-hop, 3-hop and 4-hop subsets). We employ the proprietary LLMs in our experiments and to enhance reproducibility, we set the temperature to 0 for proprietary models.

Baselines We adopt the GPT-4 (Achiam et al., 2023), GPT-3.5 (Ouyang et al., 2022), text-davinci-003, Bing Chat and GEMINI-pro (Team et al., 2023). To decouple LLMs’ internal knowledge and reasoning ability and let LLMs retrieve answers from the given passage as much as possible, we design a prompt that requires LLMs to only retrieve answers based on the given context. The prompt of QA is also shown in Appendix B.

4.2 Analysis

Reasoning VS Memorization The results of the comparison between the original data and our MRKE can be found in Table 3. LLMs show a performance gap between the original HotpotQA dataset and MRKE. E.g. GPT-4 achieves 69.3 EM and 82.2 F_1 scores respectively. While for the MRKE 2-hop dataset, GPT-4 only gets 53.2 EM and 67.7 F_1 scores. For 3-hop and 4-hop datasets, GPT-4 even performs worse, with only 41.6 EM, 61.5 F_1 scores, and 37.9 EM, 48.7 F_1 scores respectively. Since our MRKE is new, unprecedented knowledge, it illustrates the LLMs’ reasoning ability objectively.

In light of the results, we can find that LLMs achieve an inflated high performance on Wikipedia passages of the MHQA dataset possibly because of the data contamination that leads to utilizing LLMs’ memory ability rather than reasoning ability.

Joint Performance The joint F_1 RC and joint EM RC scores in table 3 are the whole reasoning chain evaluation results. We found that with the increases in the reasoning chain, the performances of LLMs dropped swiftly. E.g., the Bing Chat could get comparable performance with GPT-4 (0.7 joint F_1) on answering 2 hop questions, and get a 0.9

	Original		2 hop				3 hop				4 hop			
	EM	F ₁	EM	F ₁	Joint F ₁	RC	Joint EM	RC	EM	F ₁	Joint F ₁	RC	Joint EM	RC
GPT-4	69.3	82.2	53.2	67.7	0.7	1.5	41.6	61.5	1.1	1.6	37.9	48.7	2.3	4.2
GPT-3.5	60.1	70.4	40.4	56.9	1.7	2.4	32.9	48.3	2.7	3.2	31.2	41.7	3.6	5.8
GEMINI-pro	62.4	71.5	35.0	65.1	2.1	3.9	21.3	39.3	4.6	8.70	14.1	31.0	5.4	9.5
text-davinci-003	53.1	64.1	32.3	53.2	2.4	2.9	22.3	45.8	3.9	5.2	20.1	42.3	5.5	7.4
Bing Chat	68.7	79.0	41.7	66.8	0.9	1.9	33.6	54.3	4.2	8.4	29.6	48.5	4.7	8.9

Table 3: EM and F₁ score of LLM models on HotpotQA original data and our new annotated data. Our dataset is divided into 3 parts, 2-hop, 3-hop, and 4-hop.

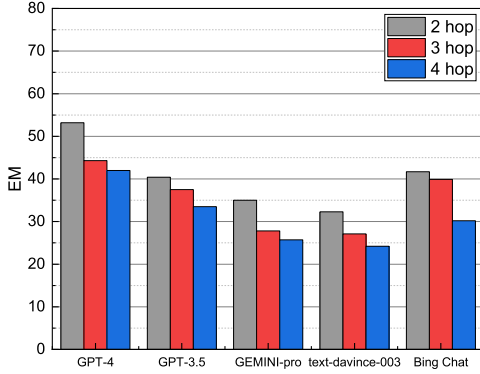


Figure 4: The performance change of EM score when answering 2 hop questions to 4 hop questions.

joint F₁ score. However, in the 3-hop question, the joint F₁ RC and joint EM RC scores of Bing Chat are 4.2 and 8.4. In the 4-hop dataset, Bing Chat gets 4.7 joint F₁ RC and 8.9 joint EM RC scores respectively. Since the joint performance is a negative log, the larger scores mean the worse performance on the reasoning chain. We can conclude that LLMs’ reasoning ability decreases with the increases in reasoning chain length.

	2 hop		3 hop		4 hop	
	EM	F ₁	EM	F ₁	EM	F ₁
GPT-4	43.7	65.4	41.6	61.5	37.9	48.7
GPT-3.5	34.3	51.3	32.9	48.3	31.2	41.7
GEMINI-pro	25.7	55.9	21.3	39.3	14.1	31
text-davinci-003	23.7	48.9	22.3	45.8	20.1	42.3
Bing Chat	37.3	62.8	33.6	54.3	29.6	48.5

Table 4: In the ablation study of the MHQA task, we remove the sub-question information from the prompt and only ask LLMs to get the final answer.

Sub QA evaluation Figures 4 and 5 show the performance of LLMs on the different hops of questions. According to the observation of the three figures, we find that with the hop increases, the complexity of multi-hop questions also increases, leading the LLMs’ performance decrease.

Figure 6 shows that LLMs also suffer the error

propagation. When wrongly answering the previous sub-question, the latter one will also be influenced. Consequently, the performance of Sub_Q2 is worse than that of Sub_Q1. Tables 7 and 8 also illustrate the sub-qa performance of LLMs on the 3-hop and 4-hop datasets in appendix C.

Reasoning Chain Evaluation In this part, due to the limitation of space, we calculate the proportion of the reasoning chain on the 2-hop dataset and present the table. We follow the setting of Tang et al. (2021) on calculating the percentage of right or wrong answers and record the results. Table 5 shows the reasoning chain evaluation results. The green row shows the percentage of examples whose multi-hop questions can be correctly answered from the right reasoning chain. The red rows show the percentage of examples whose multi-hop questions can be correctly answered but through a wrong reasoning chain. Among these examples, we notice that there is a low percentage of the LLMs successfully getting the correct final answer based on the right reasoning chain. There is also a large proportion of wrong final answers as shown in rows 2,4,6 and 8. We take the results of GPT-3.5 as an example, the right reasoning chain only accounts for 13.3% although it shows a relatively high QA performance in previous tables. The percentage of wrong reasoning chain of GPT-3.5 is 17.7% (sum of the three red rows). However, total failure cases account for 69% (sum of rows 2, 4, 6, and 8) which is substantial for the whole dataset.

We conclude that LLMs only get a small proportion of the right reasoning chain and their high performance is relatively inflated due to the considerable proportion of wrong reasoning chain.

4.3 Ablation Study

To evaluate the impact of sub-questions for LLMs, we conduct an ablation study testing the performance of answering the final answer and removing the sub-questions from prompts. The results, shown in Table 4, indicate that when directly asking LLM a multi-hop question and corresponding

q_{sub1}	q_{sub2}	q	GPT-4	GPT-3.5	GEMINI-pro	text-davinci-003	Bing Chat
c	c	c	36.3	13.3	15.0	17.3	28.3
c	c	w	12.3	9.3	9.0	10.7	7.7
c	w	c	2.0	6.7	5.3	7.7	6.0
c	w	w	25.3	24.3	14.7	25.0	16.3
w	c	c	5.7	3.7	5.3	6.7	2.3
w	c	w	3.7	3.7	5.3	3.7	3.0
w	w	c	0.3	7.3	13.3	8.7	5.0
w	w	w	14.3	31.7	32.3	30.3	31.3

Table 5: Categorical EM statistics (%) of sub-question evaluation for the five LLMs on our 2-hop dataset. Under the first three columns, *c* stands for *correct* and *w* stands for *wrong*. For example, the third row shows the percentage of questions where models correctly answer both 2-hop questions and the first sub-question but wrongly answer the second sub-question.

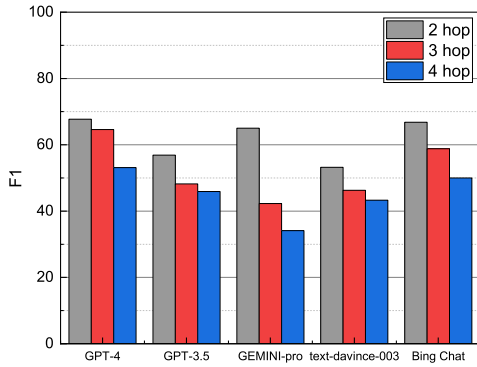


Figure 5: The performance change of F_1 score when answering 2 hop questions to 4 hop questions.

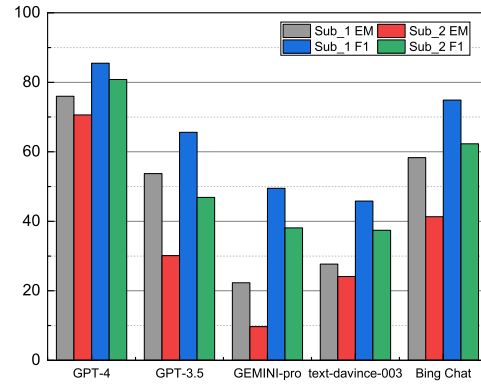


Figure 6: The performance change of F_1 score and EM scores when answering 2 sub-questions on the 2-hop dataset.

passage, the performance is much lower than that of adding sub-questions to require LLMs reasoning step-by-step. For example, computed from table 3 and 4 the performance of GPT-4 on the 2-hop dataset decreased the F_1 score and EM by 2.3 and 9.5 respectively. The results show the sub-questions could help LLMs improve the performance of final QA.

5 Conclusion

In this work, we present a benchmark that assesses the reasoning ability of LLMs via multi-hop QA and Sub QA. Although LLMs performed relatively well on QA tasks, the performance dropped on multi-hop questions that were based on new, unseen knowledge. In addition, their high performances are inflated, and benefit from the high proportion of wrong reasoning chains. We also propose a new evaluation method that jointly measures the whole reasoning chain rather than the only final answer and finds that LLMs' reasoning chain ability remains improved. We hope our work can facilitate

future research on developing faithful knowledge editing methods.

6 Limitations

The limitations of our work are as follows. 1) We only select and evaluate a small number of LLMs, The reasoning chain and data contamination risk evaluation on other LLMs remains less explored. We leave the evaluation on other models as future work. 2) We only use the EM and F_1 scores as the metrics, but such metrics did not consider the variant or abbreviation of the answers, sometimes correct answers are recognized as wrong, which remains explored in the future. 3) The multi-hop qa pair, sub-qa pairs, and related passages in MRKE are automatically generated by GPT-4, rather than being crafted by humans. Although we manually reviewed the 300 data, it is hard to make sure the data quality of other data.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Simran Arora, Avani Narayan, Mayee F Chen, Laurel J Orr, Neel Guha, Kush Bhatia, Ines Chami, Frederic Sala, and Christopher Ré. 2022. Ask me anything: A simple strategy for prompting language models. *arXiv preprint arXiv:2210.02441*.
- Yushi Bai, Xin Lv, Jiajie Zhang, Hong Lyu, Jiankai Tang, Zhidian Huang, Zhengxiao Du, Xiao Liu, Aohan Zeng, Lei Hou, Yuxiao Dong, Jie Tang, and Juanzi Li. 2023. [Longbench: A bilingual, multitask benchmark for long context understanding](#). *ArXiv*, abs/2308.14508.
- Max Bartolo, Tristan Thrush, Sebastian Riedel, Pontus Stenetorp, Robin Jia, and Douwe Kiela. 2021. [Models in the loop: Aiding crowdworkers with generative annotation assistants](#). In *North American Chapter of the Association for Computational Linguistics*.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, T. W. Hennigan, Saffron Huang, Lorenzo Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and L. Sifre. 2021. [Improving language models by retrieving from trillions of tokens](#). In *International Conference on Machine Learning*.
- Zichen Chen, Jianda Chen, Mitali Gaidhani, Ambuj Singh, and Misha Sra. 2023. [Xplainllm: A qa explanation dataset for understanding llm decision-making](#). *ArXiv*, abs/2311.08614.
- Shizhe Diao, Pengcheng Wang, Yong Lin, and Tong Zhang. 2023. Active prompting with chain-of-thought for large language models. *arXiv preprint arXiv:2302.12246*.
- Ming Ding, Chang Zhou, Qibin Chen, Hongxia Yang, and Jie Tang. 2019. [Cognitive graph for multi-hop reading comprehension at scale](#). *ArXiv*, abs/1905.05460.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *North American Chapter of the Association for Computational Linguistics*.
- Tianyu Gao, Ho-Ching Yen, Jiatong Yu, and Danqi Chen. 2023a. [Enabling large language models to generate text with citations](#). *ArXiv*, abs/2305.14627.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu Guo, Meng Wang, and Haofen Wang. 2023b. [Retrieval-augmented generation for large language models: A survey](#). *ArXiv*, abs/2312.10997.
- Shahriar Golchin and Mihai Surdeanu. 2023. [Data contamination quiz: A tool to detect and estimate contamination in large language models](#). *ArXiv*, abs/2311.06233.
- Hengrui Gu, Kaixiong Zhou, Xiaotian Han, Ninghao Liu, Ruobing Wang, and Xin Wang. 2023. [Pokemqa: Programmable knowledge editing for multi-hop question answering](#). *ArXiv*, abs/2312.15194.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Xiaodong Song, and Jacob Steinhardt. 2020. [Measuring massive multitask language understanding](#). *ArXiv*, abs/2009.03300.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. [Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- O. Khattab, Christopher Potts, and Matei A. Zaharia. 2021. [Baleen: Robust multi-hop reasoning at scale via condensed retrieval](#). *ArXiv*, abs/2101.00436.
- Kevin Meng, David Bau, Alex Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in gpt](#). In *Neural Information Processing Systems*.
- Sewon Min, Victor Zhong, Luke Zettlemoyer, and Han-naneh Hajishirzi. 2019. Multi-hop reading comprehension through question decomposition and rescoring. In *ACL*.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke E. Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Francis Christiano, Jan Leike, and Ryan J. Lowe. 2022. [Training language models to follow instructions with human feedback](#). *ArXiv*, abs/2203.02155.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Manley Roberts, Himanshu Thakur, Christine Herlihy, Colin White, and Samuel Dooley. 2023. [Data contamination through the lens of time](#). *ArXiv*, abs/2310.10628.
- Weijia Shi, Anirudh Ajith, Mengzhou Xia, Yangsibo Huang, Daogao Liu, Terra Blevins, Danqi Chen, and Luke Zettlemoyer. 2023. [Detecting pretraining data from large language models](#). *ArXiv*, abs/2310.16789.

639	Yixuan Tang, Hwee Tou Ng, and Anthony Tung. 2021.	human-centric benchmark for evaluating foundation	694
640	Do multi-hop question answering systems know how	models. <i>ArXiv</i> , abs/2304.06364.	695
641	to answer the single-hop sub-questions? In <i>Proceed-</i>		
642	ings of the 16th Conference of the European Chap-	Zexuan Zhong, Zhengxuan Wu, Christopher D. Man-	696
643	ter of the Association for Computational Linguistics:	ning, Christopher Potts, and Danqi Chen. 2023b.	697
644	Main Volume, pages 3244–3249, Online. Association	Mquake: Assessing knowledge editing in language	698
645	for Computational Linguistics.	models via multi-hop questions. In <i>Conference on</i>	699
		<i>Empirical Methods in Natural Language Processing</i> .	700
646	Yixuan Tang and Yi Yang. 2024. Multihop-rag: Bench-		
647	marking retrieval-augmented generation for multi-	Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun,	701
648	hop queries. <i>ArXiv</i> , abs/2401.15391.	and Chao Zhang. 2023. Toolqa: A dataset for	702
		llm question answering with external tools. <i>ArXiv</i> ,	703
649	Gemini Team, Rohan Anil, Sebastian Borgeaud,	abs/2306.13304.	704
650	Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu,		
651	Radu Soricut, Johan Schalkwyk, Andrew M Dai,	A Example Appendix	705
652	Anja Hauth, et al. 2023. Gemini: a family of		
653	highly capable multimodal models. <i>arXiv preprint</i>	We test the two questions on ChatGPT, one is the	706
654	<i>arXiv:2312.11805</i> .	original question for the raw context and the other	707
655	Petter Törnberg. 2023. Chatgpt-4 outperforms ex-	is the new question for the re-annotated context.	708
656	perts and crowd workers in annotating political	We found that when we input the original ques-	709
657	twitter messages with zero-shot learning. <i>ArXiv</i> ,	tion, ChatGPT could give the right answer without	710
658	abs/2304.06588.	any background information but failed with the	711
659	Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot,	new question, demonstrating that ChatGPT gives	712
660	and Ashish Sabharwal. 2022. MuSiQue: Multi-	answer based on memory, rather than reasoning	713
661	hop questions via single-hop question composition.	ability.	714
662	<i>Transactions of the Association for Computational</i>		
663	<i>Linguistics</i> .	B Prompts	715
664	Cunxiang Wang, Sirui Cheng, Qipeng Guo, Zhikun Xu,		
665	Bowen Ding, Yidong Wang, Xiangkun Hu, Zheng	When evaluating large language models, prompting	716
666	Zhang, and Yue Zhang. 2023. Evaluating open-qa	is a brittle process wherein small modifications to	717
667	evaluation.	the prompt can cause large variations in the model	718
668	Yunxuan Xiao, Yanru Qu, Lin Qiu, Hao Zhou, Lei Li,	predictions, and therefore significant effort should	719
669	Weinan Zhang, and Yong Yu. 2019. Dynamically	be dedicated to designing a painstakingly crafted	720
670	fused graph network for multi-hop reasoning. <i>ArXiv</i> ,	perfect prompt for the given task (Arora et al., 2022;	721
671	abs/1905.06933.	Diao et al., 2023). In this study, We investigate the	722
672	Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Ben-	performance of zero-shot on our benchmark. To	723
673	gio, William W. Cohen, Ruslan Salakhutdinov, and	eliminate the randomness, we manually select one	724
674	Christopher D. Manning. 2018. HotpotQA: A dataset	demonstration for each task, ensuring that all tasks	725
675	for diverse, explainable multi-hop question answer-	are covered.	726
676	ing. In <i>Conference on Empirical Methods in Natural</i>	We give our designed input examples for three	727
677	<i>Language Processing (EMNLP)</i> .	different tasks to help readers understand our im-	728
678	Yunxiang Zhang, Muhammad Khalifa, Lajanugen Lo-	plementation, as shown in Table 6, respectively.	729
679	geswaran, Moontae Lee, Honglak Lee, and Lu Wang.		730
680	2023. Merging generated and retrieved knowledge		
681	for open-domain qa. <i>ArXiv</i> , abs/2310.14393.	C Sub-QA evaluation	731
682	Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang,		
683	Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen	As the quantitative complementary of Sub-QA eval-	732
684	Zhang, Junjie Zhang, Zican Dong, et al. 2023. A	uation, we here list the results of LLMs’ perfor-	733
685	survey of large language models. <i>arXiv preprint</i>	mance on 3-hop and 4-hop datasets. The LLMs’	734
686	<i>arXiv:2303.18223</i> .	reasoning performance dropped dramatically, e.g.	735
687	Ce Zheng, Lei Li, Qingxiu Dong, Yuxuan Fan, Zhiyong	in table 7, GPT-4 achieves 70.9 EM and 80.8 F ₁	736
688	Wu, Jingjing Xu, and Baobao Chang. 2023. Can	scores on sub-question1 but only gets 59.7 EM,	737
689	we edit factual knowledge by in-context learning?	74.9 F ₁ , and 58.1 EM, 68.8 F ₁ scores on sub-	738
690	<i>ArXiv</i> , abs/2305.12740.	question2 and sub-question3 respectively. In ta-	739
691	Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang,	ble 8, we further find that when answering 4 hop	740
692	Shuai Lu, Yanlin Wang, Amin Saied Sanosi Saied,	questions, the results show a cliff-like descent from	741
693	Weizhu Chen, and Nan Duan. 2023a. Agieval: A		

Answers from ChatGPT

who is the captain of Argentina football national team that was born in 1987

ChatGPT

As of my last knowledge update in January 2022, Lionel Messi was born in 1987 and has been the captain of the Argentina national football team. However, please note that team rosters and captaincies can change, so it's advisable to check the latest information from reliable sources for the most up-to-date details.



Who was the coach of the Brazil team that was born in 1997?

ChatGPT

As of my last knowledge update in January 2022, I don't have specific information about the coach of the Brazil national football team born in 1997. Coaches and staff can change over time, and it's advisable to check the latest sports news or official team announcements for the most up-to-date information on the coaching staff of the Brazil national football team.

Figure 7: Example answers generated from ChatGPT

Prompts of NER, Noun Phrase and Adjective replacement

Prompt Now you are a passage annotator, you need to recognize all the named entities, noun phrases, and adjectives from the given [CONTEXT], then translate the passage into Chinese and translate to English. Please output the response in JSON format {Passage: String}
[CONTEXT] The given context.

Prompts of Question Generation

Example One-shot example with multihop QA pairs, Sub-QA pairs, and passage.

Prompt Now you are a multihop question generation machine, given an example of 2 hop question and its sub-questions, sub-answers, and final answer is [2 hop question],[Sub-Questions],[Sub-Answers] and [Final Answer], you need to generate a new 2 hop multihop question same with the given example and its sub-questions, sub-answers and final answer from the given [Context]. Please follow the sentence structure of give examples and output the response in JSON format {2 hop question: String, sub-questions: List, sub-answers:List, final answer:String}:

[2 hop question] The given example of 2 hop question.

[Sub-Questions] The given example of sub-questions.

[Sub-Answers] The given example of sub-answers.

[Final Answer] The given example of final answer.

[CONTEXT] The given passage

Prompts of QA

Prompt You are a QA test machine, you need to answer the [Question] from given the [Context], and you only need to come out with the correct answer without other words. Let's think step by step, and please output the answer to the [Question] in the format of: {Final Answer: String}.

[QUESTION] The given question.

[CONTEXT] The given passage.

Table 6: The prompt template of passage rewriting and question generation. [WORDS] denotes the information we should give.

sub-question2 to sub-question3, especially GPT-3.5 gets 46.9 F₁ in sub-question2 but drop to 36.3 F₁ score in sub-question3.

	3 hop					
	Sub_Q1 EM	Sub_Q1 F ₁	Sub_2 EM	Sub_Q2 F ₁	Sub_Q3 EM	Sub_Q3 F ₁
GPT-4	70.9	80.8	59.7	74.9	58.1	68.8
GPT-3.5	43.0	56.4	38.6	49.3	29	40.6
GEMINI-pro	5.8	33.8	4.4	30.8	4.1	31.5
text-davinci-003	23.3	42.4	20.5	33.7	19.5	29.6
Bing Chat	7.2	34.0	5.8	31.5	3.1	32.3

Table 7: The LLM evaluation on MRKE 3 hop dataset. We here measure the sub-qa task and compare the performance between each hop.

	4 hop							
	Sub_Q1 EM	Sub_Q1 F ₁	Sub_Q2 EM	Sub_Q2 F ₁	Sub_Q3 EM	Sub_Q3 F ₁	Sub_Q4 EM	Sub_Q4 F ₁
GPT-4	60.9	66.7	56.4	62.6	28.4	58.7	23.1	56.3
GPT-3.5	40.7	46.9	30.1	36.3	20.2	47.2	14.7	44.8
GEMINI-pro	14.9	39.2	10.4	38.3	9.1	34.9	7.2	29.5
text-davinci-003	19.8	39.2	19.2	30.7	18.8	28.6	18.5	27.8
Bing Chat	20.8	39.4	16.9	37.1	6.2	35.8	5.5	35.1

Table 8: The LLM performance on MRKE 4 hop dataset. We here measure the sub-qa task and compare the performance between each hop.