

LLatrieval: LLM-Verified Retrieval for Verifiable Generation

Xiaonan Li*, Changtai Zhu*, Linyang Li, Zhangyue Yin, Tianxiang Sun, Xipeng Qiu

School of Computer Science, Fudan University

Shanghai Key Laboratory of Intelligent Information Processing, Fudan University

{lixn20, xpqiu}@fudan.edu.cn, ctzhu23@m.fudan.edu.cn

Abstract

Verifiable generation aims to let the large language model (LLM) generate text with supporting documents, which enables the user to flexibly verify the answer and makes the LLM’s output more reliable. Retrieval plays a crucial role in verifiable generation. Specifically, the retrieved documents not only supplement knowledge to help the LLM generate correct answers, but also serve as supporting evidence for the user to verify the LLM’s output. However, the widely used retrievers become the bottleneck of the entire pipeline and limit the overall performance. Their capabilities are usually inferior to LLMs since they often have much fewer parameters than the large language model and have not been demonstrated to scale well to the size of LLMs. If the retriever does not correctly find the supporting documents, the LLM can not generate the correct and verifiable answer, which overshadows the LLM’s remarkable abilities. To address these limitations, we propose **LLatrieval** (Large Language Model Verified Retrieval), where the LLM updates the retrieval result until it verifies that the retrieved documents can sufficiently support answering the question. Thus, the LLM can iteratively provide feedback to retrieval and facilitate the retrieval result to fully support verifiable generation. Experiments on ALCE show that LLatrieval significantly outperforms extensive baselines and achieves state-of-the-art results.

1 Introduction

Large language models (LLMs) have shown remarkable abilities over various downstream tasks (OpenAI, 2023; Ouyang et al., 2022; Chowdhery et al., 2023; Touvron et al., 2023; Zhao et al., 2023). However, LLMs struggle with factual errors, and often produce non-factual and fabricated content (Zhang et al., 2023; Rawte et al., 2023; Wang et al., 2023a; Huang et al., 2023b; Min et al., 2023;

*Equal Contribution

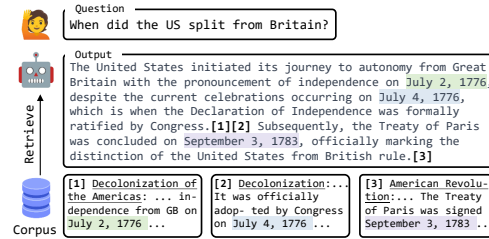


Figure 1: Verifiable Generation (Gao et al., 2023c)

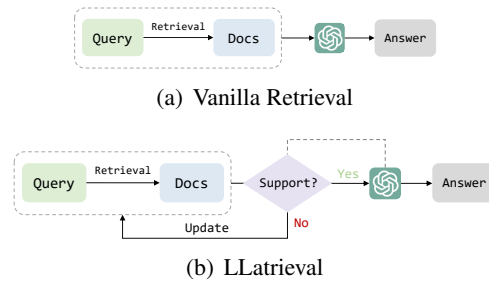


Figure 2: When the vanilla retrieval overshadows the LLM’s remarkable abilities in the pipeline, LLatrieval can fully harness the LLM’s abilities to the retrieval by verify-update iterations.

Chern et al., 2023), which is usually referred to as “hallucination” and makes the LLM’s response not trustworthy. To address these challenges, researchers propose a new generation paradigm, Verifiable Generation (Gao et al., 2023a; Bohnet et al., 2022; Gao et al., 2023c; Li et al., 2023a), shown in Figure 1. For a given question, it requires the LLM to generate the answer with corresponding supporting documents. In this way, these documents can serve as evidence and enable users to flexibly verify the answer, which makes the LLM’s response more reliable and facilitates its application in various important scenarios, medical diagnosis (Schuster et al., 2021), scientific writing (Salvagno et al., 2023), situation reports (Reddy et al., 2023), etc.

Retrieval plays a crucial role in verifiable generation. Take retrieval-read, a typical pipeline for verifiable generation (Gao et al., 2023c; Bohnet

et al., 2022), as an example, shown in Figure 2(a). It is divided into two stages: 1) retrieve the relevant documents for the given question, 2) generate the answer according to the retrieved documents. On the one hand, the retrieved documents serve as knowledge supplement to help the LLM generate the answer. On the other hand, the answer’s verifiability relies on the retrieved documents. Thus, both of the answer’s correctness and verifiability highly rely on the quality of retrieval result.

However, the widely used retrievers become the bottleneck of the entire pipeline and limit the overall performance (Wang et al., 2023c; Yoran et al., 2023). Specifically, although the dense retrievers have achieved a series of state-of-the-art results in various scenarios (Karpukhin et al., 2020; Zhang et al., 2022a; Su et al., 2023), their capabilities are substantially inferior to LLMs (Brown et al., 2020; Chowdhery et al., 2023), since they usually have much fewer parameters, and have not been demonstrated to scale well to the size of LLMs (Brown et al., 2020) and achieve comparable capabilities. If the retriever does not correctly find the supporting documents, it is challenging for the LLM to output the answer which is both correct and verifiable. In the typical pipeline, the LLM usually receives the retrieval result in a passive manner and can not provide feedback to the low-quality retrieval even if it is capable of identifying that the retrieved documents are irrelevant to the question. These make the retrieval overshadow the LLM’s remarkable abilities and limit the overall performance.

To address these limitations, we propose the framework of **LLattribution** (Large Language Model Verified Retrieval), shown in Figure 2(b), where the LLM fully provides feedback to the retrieval and augments it by the verify-update iteration: 1) Retrieval Verification: The LLM verifies whether the retrieval result can support answering the give question. For questions whose documents fail the verification, we will update their documents and thus avoid low-quality retrieval. 2) Retrieval Update: The LLM updates and improves the retrieval result. Specifically, we propose Progressive Selection and Missing-Info Query to let the LLM progressively scan document candidates and supplement missing information to facilitate better verifiable generation. Through the verify-update iteration, the LLM can iteratively refine the retrieval result until it verifies that the retrieved documents can support answering the given question, and thus generate both correct and verifiable answers.

We summarize our contributions as follows:

- To the best of our knowledge, LLattribution is the first framework to let the LLM fully provide feedback to retrieval and thus augments it through verify-update iterations.
- We conduct comparison with extensive baselines and results show that LLattribution significantly outperforms baselines and achieves new state-of-the-art results on verifiable generation.
- Further analyses show that each component contributes critically to improvements and our method shows the potential to scale retrieval by scaling the LLM.
- We make our code publicly available to facilitate future research. ¹

We hope that LLattribution can inspire researchers of the important design choices for LLM-augmented Retrieval and pave the way for further improvements.

2 Background: Verifiable Generation

In this paper, we follow the setting proposed by Gao et al. (2023c), which requires the LLM to generate the answer with citations to the supporting documents, as shown in Figure 1. This setting is close to the real-world applications of generative search engines, e.g., New Bing, and the citations allow the user to verify the answer more conveniently. We introduce its typical pipeline and the verifiability evaluation as follows.

2.1 Pipeline

Verifiable generation is usually demonstrated by retrieval-read (Gao et al., 2023c; Bohnet et al., 2022; Malaviya et al., 2023), divided into two stages: retrieval and generation. First, given the question q , the relevant documents, $D = \{d_1, d_2, \dots, d_k\}$, are retrieved as follows:

$$D = R(q, \mathbb{D}, k) = \text{Top-}k_{d \in \mathbb{D}} \text{ score}(q, d) \quad (1)$$

where R is the retriever and $\text{score}(\cdot, \cdot)$ is the similarity score of R (Robertson and Zaragoza, 2009; Karpukhin et al., 2020). Then the LLM leverages these documents to generate the answer as:

$$A = \text{LLM}(I_q; \text{Demos}; q; D), \quad (2)$$

¹<https://github.com/BeastyZ/LLM-Verified-Retrieval>

where we follow Gao et al. (2023c) to use the instruction, I_g , and demonstrations, $Demos$, to make the LLM generate answer and citations to D .

2.2 Verifiability Evaluation

Following Gao et al. (2023c); Liu et al. (2023b), we evaluate the answer’s verifiability by measuring citation quality: **Citation Recall** evaluates whether the output is fully supported by cited documents; **Citation Precision** identifies those citations which can not help support the answer. According to them, we calculate **Citation F1** to comprehensively evaluate the answer’s verifiability. Since the retrieved documents determine the citation quality and the answer’s verifiability, the retrieval is crucial for the overall pipeline. We adopt the recently proposed ALCE benchmark for automatic citation evaluation (Gao et al., 2023c).

3 LLattribution: LLM-Verified Retrieval

In the retrieval-read pipeline, the LLM passively receives the retrieval result and can not provide feedback to the low-quality retrieval, which makes the large language model with remarkable abilities overshadowed by the retriever with much fewer parameters. To address these limitations, we propose LLattribution, where the LLM iteratively provides feedback to the retrieval through verify-update iterations, and thus fully harnesses its abilities for verifiable generation. In this section, we introduce the modules of retrieval verification and update, and the iterative verify-update process as follows.

3.1 Retrieval Verification

Since low-quality retrieval overshadows the LLM’s remarkable abilities for verifiable generation, we propose to use the LLM itself to identify the low-quality retrieval and thus avoid it. In this paper, we instantiate the retrieval verification as two simple instruction-based methods, shown in Figure 3, to let the LLM verify whether the retrieved documents can support answering the question, and leave more elaborate methods as future work.

Classification Given the question q and retrieved documents D , we verify whether D can sufficiently support answering q by prompting LLM to give a binary label (*yes* or *no*) as:

$$\text{Verify-Result} = \text{LLM}(I_v^c; q; D) \quad (3)$$

where I_v^c is the corresponding instruction. Through the LLM’s strong natural language understanding

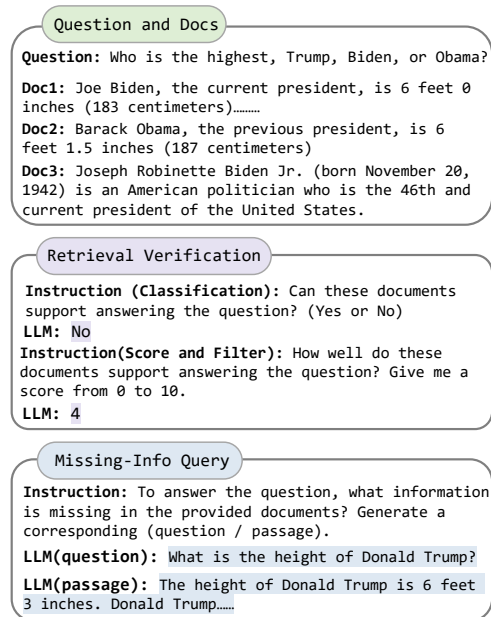


Figure 3: **Retrieval Verification:** We propose two ways to verify whether the documents can support answer the question. **Missing-Info Query:** We propose to let the LLM generate missing-info query in two styles.

ability, we can fully identify the low-quality retrieval and thus avoid it. Meanwhile, LLattribution finishes when the example’s retrieved documents pass the retrieval verification, which enables it to dynamically adjust the LLM cost under varying retrieval demands.

Score-and-Filter We further explore a score-and-filter strategy for retrieval verification. We first score the current documents by how well they can support answering the question and then filter low-quality retrieval by the pre-defined threshold τ as:


$$\text{Score}_D = \text{LLM}(I_v^s; q; D) \quad (4)$$

$$\text{Verify-Result} = \text{Yes if } \text{Score}_D \geq \tau \text{ else No} \quad (5)$$


where I_v^s is the corresponding instruction. The higher τ leads to stricter retrieval verification criterion, better quality of the final retrieval result and more iterations. By tuning τ , we can dynamically adjust the strictness of verification and the corresponding LLM cost.

3.2 Retrieval Update

When the current documents fail the retrieval verification, we will update and refine the low-quality retrieval result. Specifically, we propose the following two LLM-based modules to update the retrieval result in synergy and thus better harness the LLM’s abilities for retrieval.



(a) Progressively selecting the documents, where the blue and purple color represent the current documents D and new candidates D_c^* in each selection.



(b) Using the LLM to select docs from candidates.

Figure 4: Retrieval update by progressive selection.

Progressive Selection Since the dense retrievers (Karpukhin et al., 2020; Su et al., 2023) measure similarity based on the dual-encoder architecture and usually have fewer parameters than the prevailing large language models (Brown et al., 2020; Chowdhery et al., 2023), it is challenging for them to fully demonstrate fine-grained token-level interactions and thus accurately rank the most relevant documents at top-K (Zhang et al., 2022a). Inspired by the retrieval-rerank framework (Nogueira and Cho, 2019), which reranks the document candidates returned by the retriever to get the final retrieval result, we propose Progressive Selection for retrieval update: the LLM progressively scans the document candidates returned by the retriever and selects the supporting documents, shown in Figure 4. For each selection, we obtain a small list of new document candidates, D_c^* , by sliding window on the document candidates (Figure 4(a)). Then we use the LLM to select k documents which can maximally support answering q , from $D \cup D_c^*$, as the *updated* D (Figure 4(b)). Thus, the LLM can remove irrelevant documents in D and integrate the critical information in D_c^* to update and improve D , with D 's size unchanged. Compared to reranking, the selection directly outputs a combination of documents, so the updated D can avoid documents with redundant information and thus get more comprehensive. With iteratively updating D from new D_c^* , the LLM can progressively improve D to make it better support verifiable generation.

Missing-Info Query Since the current documents can not sufficiently support answering the question, we propose to update the current documents D by querying missing information of D for supplement, shown in Figure 3. Specifically, we first use the LLM to generate the query of the missing information in D for answering q as:

$$Q = \text{LLM}(I_q; q; D), \quad (6)$$

where I_q instructs the LLM to identify the missing information in the current documents D for the given question q and generate the corresponding query Q , shown in Figure 3. Then we use Q to retrieve the corresponding documents for information supplement. For the fusion of the newly retrieved documents and D , please refer to section 3.3.

When the current dense retrievers often use the BERT-based architecture (Karpukhin et al., 2020) and produce one vector for calculating similarity, it is challenging for them to comprehensively tackle the user query involving the knowledge from multiple aspects (Chen et al., 2023; Zhang et al., 2022b) and complicated reasoning, which widely exists in the real-world scenarios. During the verify-update iterations, the missing-info query can supplement multiple-aspect information through iteratively identifying and querying the missing information by LLM, which can fully harness the LLM's remarkable language understanding and reasoning abilities. In experiments, we explore two styles of missing-info query: 1) *question*: the corresponding question of missing information; 2) *pseudo passage*: inspired by the query rewriting methods (Gao et al., 2023b; Shen et al., 2023), we generate a pseudo passage of missing information as retrieval query, which neighbors the relevant documents and thus can help better retrieve them. We regard developing stronger missing-info query modules as future work.

3.3 Verify-Update Iteration

We show the overall procedure in Algorithm 1. In each iteration, we first update the missing-info query and retrieve a new list of document candidates D_c that may contain missing information of D (Line 4 ~ 7). Then we update the current D by progressively selecting from D_c (Line 8 ~ 10). Through the synergy between progressive selection and missing-info query, we can identify and retrieve D 's missing information and progressively incorporate them into D . Since the progressive selection is conducted on $D \cup D_c$, we can not only

Algorithm 1 LLattribution

Input: Question q , document pool \mathbb{D} , the large language model LLM, the retriever R , the maximum iteration T , each iteration’s document candidates quantity N

Output: Supporting Documents D

```
1:  $Q \leftarrow q$ 
2:  $D \leftarrow \{\}$ 
3: for  $i$  in  $1, 2 \dots T$  do
4:   if  $D \neq \{\}$  then
5:      $Q \leftarrow$  Use LLM to generate missing-info’s query
6:   end if
7:    $D_c \leftarrow R(Q, \mathbb{D}, N)$ 
8:   for  $D_c^*$  in SlidingWindow( $D_c$ ) do
9:      $D \leftarrow$  Use LLM to select  $k$  docs from  $D \cup D_c^*$ 
10:  end for
11:  if Verify( $q, D$ )  $\rightarrow$  Yes then
12:    #  $D$  can support answering  $q$ , so return  $D$ .
13:    break
14:  end if
15: end for
16: return  $D$ 
```

supplement D ’s missing information for q , but also fully retain the critical documents in D . In this way, we comprehensively integrate the relevant documents with multifaceted information and thus support better verifiable generation. After each update, we use the LLM to verify whether the current documents D can support answering the given question q (Line 11 ~ 13) and continue updating the low-quality retrieval result in the next iteration. Through the verify-update iteration, the retrieval result can be iteratively updated and improved until it can sufficiently support high-quality verifiable generation, and thus help the LLM generate both correct and verifiable answers. Meanwhile, our method can dynamically adjust the number of iterations under varying retrieval demands and thus help save the LLM cost.

4 Experiments

4.1 Experimental Settings

Method Comparison Since we focus on augmenting the retrieval for verifiable generation, we compare LLattribution with various retrieval methods by evaluating the LLM output’s correctness and verifiability under the same retrieval-read pipeline (Gao et al., 2023c), including: **Retriever:** 1) BM25 (Robertson and Zaragoza, 2009) 2) DPR (Karpukhin et al., 2020) 3) Contriever (Izacard et al., 2022): A large-scale unsupervised contrastive learned dense encoder (Izacard et al., 2022). 4) GTR (Ni et al., 2022): A large dense retriever initialized from T5-XXL, pre-trained on a web-mined corpus and fine-tuned on information retrieval datasets. 5) Instruc-

tor (Su et al., 2023): An instruction-tuned dense retriever on about 330 diverse tasks. 6) BGE-Embedding (Xiao et al., 2023): A recent large-scale trained general-purpose embedder which outperforms OpenAI’s “text-embedding-ada-002” API. **Re-ranker:** 1) monoBERT (Nogueira and Cho, 2019)&monoT5 (Nogueira et al., 2020): traditional point-wise re-rankers based on BERT and T5. 2) BGE-Reranker (Xiao et al., 2023): A recent large-scale trained general purpose re-ranker. **LLM-based Retrieval Augmentation:** 1) Query Rewriting: We follow HYDE (Gao et al., 2023b) and LAMER (Shen et al., 2023) to generate the pseudo passage based on the question and its relevant passages, respectively. Then we use the pseudo passage to conduct retrieval. 2) LLM-reranking: RankGPT (Sun et al., 2023), which prompts the LLM to rerank passages by permutation generation. Except the retriever baselines, we use the same retriever as in LLattribution. For more details, please refer Appendix A. We show the correlation between retrieval quality and generation quality in Appendix B.

Dataset and Evaluation We conduct experiments on the ALCE benchmark (Gao et al., 2023c), which consists of three long-form QA datasets, and follow ALCE to evaluate the correctness and verifiability of generation.

ASQA (Stelmakh et al., 2022) is a long-form factoid QA dataset where each question requires multiple short answers to cover the multiple aspects of it. To evaluate the answer’s correctness, we calculate the recall of correct short answers (from the dataset) by identifying whether the short answers match substrings of model output (*exact match recall*; EM-R).

QAMPARI (Amouyal et al., 2022): a factoid QA dataset where the answer is a list of entities from different documents. To evaluate the answer’s correctness, we calculate the F1 of the prediction by identifying the exact match to the gold answer list. Following Gao et al. (2023c), we consider the recall to be 100% when prediction covers at least 5 correct answers.

ELI5 (Fan et al., 2019): a long-form QA dataset where each question requires comprehensive long answers and multiple documents as evidence. To evaluate the answer’s correctness, we follow Gao et al. (2023c) to measure whether the model prediction entails the sub-claims of the gold answer.

Following ALCE (Gao et al., 2023c), on ASQA

Dataset	ASQA				QAMPARI				ELI5				Overall	
Evaluation Metric	Correct		Citation		Correct		Citation		Correct		Citation		Correct	Citation F1
	EM-R	Rec	Prec	F1	F1	Rec	Prec	F1	Claim	Rec	Prec	F1		
BM25	51.7	43.3	48.5	45.7	17.5	21.4	25.7	23.3	11.7	52.9	49.0	50.9	27.0	40.0
DPR	54.4	52.2	54.6	53.3	18.5	20.0	23.2	21.5	-	-	-	-	-	-
Contriever-MS	54.9	54.0	57.1	55.5	17.8	20.7	23.7	22.1	-	-	-	-	-	-
GTR	54.4	52.2	54.6	53.3	18.5	20.0	23.2	21.5	-	-	-	-	-	-
Instructor-large	55.6	52.1	54.1	53.1	17.2	19.4	21.9	20.6	-	-	-	-	-	-
BGE-E-large	55.9	56.4	58.6	57.5	17.3	22.7	25.5	24.0	-	-	-	-	-	-
monoBERT (50)	53.8	55.8	57.4	56.6	20.9	26.6	30.3	28.3	13.8	63.5	58.3	60.8	29.5	48.5
monoBERT (100)	53.0	54.0	56.5	55.2	20.2	25.5	29.1	27.2	14.1	63.3	58.2	60.6	29.1	47.7
monoT5-3B (50)	54.0	55.6	57.7	56.6	20.3	28.1	32.0	30.0	13.9	64.2	59.3	61.7	29.4	49.4
monoT5-3B (100)	53.7	53.0	55.2	54.1	19.7	26.0	29.6	27.7	14.1	63.9	58.8	61.2	29.2	47.7
BGE-R-large (50)	55.3	57.2	59.0	58.1	20.7	26.6	30.8	28.5	13.9	65.9	61.4	63.6	30.0	50.1
BGE-R-large (100)	54.7	57.4	58.7	58.1	19.9	25.9	29.6	27.6	14.4	67.1	62.6	64.8	29.7	50.2
HYDE	56.8	59.1	60.5	59.8	16.2	20.2	27.7	23.4	15.2	70.4	63.4	66.7	29.4	49.9
HYDE*	56.7	57.5	57.4	57.5	16.5	20.1	26.2	22.7	-	-	-	-	-	-
LAMER	55.3	59.1	60.1	59.6	19.0	23.0	31.0	26.4	14.4	70.8	62.7	66.5	29.0	50.0
LAMER*	55.7	58.7	60.2	59.4	16.8	20.6	29.1	24.1	-	-	-	-	-	-
RankGPT (50)	56.4	58.5	58.9	58.7	20.7	27.8	31.6	29.6	13.5	64.7	58.8	61.6	30.2	50.0
RankGPT (100)	57.0	58.8	59.7	59.2	22.0	28.4	32.9	30.5	13.5	63.5	57.6	60.4	30.8	50.1
LLattribution (passage)	57.7	60.9	61.3	61.1	24.3	32.4	36.6	34.4	16.5	75.8	67.4	71.3	32.9	55.6
LLattribution (question)	57.3	60.5	60.6	60.6	23.6	30.9	34.8	32.7	16.7	75.4	68.0	71.5	32.5	54.9

Table 1: The performance comparison on ALCE (Gao et al., 2023c), where “(passage)” and “(question)” represent the styles of missing-info query in LLattribution. For reranking baselines, the number in parentheses is the number of reranked document candidates. For query rewriting baselines, the superscript “*” represents using their original retriever, otherwise using ours.

and QAMPARI, we use aliases of short answers provided by the dataset and normalize the model output and the short answers when measuring exact match. For ASQA, we use its sub-questions as the question to eliminate the original question’s ambiguity, for simplicity.

For details of verifiability evaluation, please refer to Section 2.2.

Implementation Details We use the public OpenAI API of “gpt-3.5-turbo-0301” for retrieval and generation, unless otherwise specified. We use window size=20 and set the number of document candidates per query as 50. We set the number of supporting documents and the maximum of iterations as 5 and 4. We show LLattribution’s performance under different hyper-parameters and retrievers in Appendix C. For the document pool, we follow Gao et al. (2023c) to use the Wikipedia for ASQA and QAMPARI, and Sphere (Piktus et al., 2021), a filtered version of a filtered version of Common Crawl for ELI5, respectively. For ASQA and QAMPARI, we use the recently proposed dense embedder, BGE-large (Xiao et al., 2023), as the retriever. For ELI5, we follow Gao et al. (2023c) to use BM25 (Robertson and Zaragoza, 2009) since dense retrievers are costly and slow for large-scale web

corpus. Since our method may involve multiple queries for one example, the number of its document candidates per example is from 50 to 100 (See Appendix A). For comprehensive comparison, we evaluate the re-ranking baselines with both 50 and 100 document candidates. For the overall implementation details, please refer to Appendix A.

4.2 Main Results

We show the results in Table 1, where we adopt the classification mode in retrieval verification for stable performance since the performance of score-and-filter mode varies with different thresholds. We see that LLattribution significantly outperforms baselines on all three datasets’ correctness and verifiability, which shows its best overall retrieval performance. Specifically, LLattribution outperforms the used retriever by 3.4 and 5.9 points on correctness and citation-F1, respectively, which straightly demonstrates its effectiveness. Notably, our method shows consistent improvements over other LLM-augmented retrieval methods, which only show comparable performance with traditional rerankers, and these indicate that our method can more effectively harness the LLM’s abilities for augmenting retrieval. Meanwhile, various re-




Figure 5: The generation quality of filtered examples over different thresholds.

trievers consistently show worse performance than LLM-augmented methods, and this demonstrates that the traditional retrievers become the bottleneck for verifiable generation and overshadow the LLM’s strong abilities. Additionally, our method demonstrates significant improvements over traditional neural rerankers, which indicates that the bottleneck effect of retriever can not be fully mitigated by traditional rerankers.

4.3 Analyses

In this section, we conduct analyses of LLaTrieval, where we adopt classification-based retrieval verification and “passage” style missing-info query, and report Citation-F1 for verifiability for simplicity.

Ablation Study We evaluate the effectiveness of each component on ASQA and QAMPARI. We show the performance changes after adding each component one-by-one, from the original retriever to the second iteration, in Table 2. First we conduct progressive selection on the original question’s document candidates and find it leads to performance improvements, which directly indicates that it can help the LLM find better relevant documents from the candidates. Then we verify whether each question’s current documents can support answering it. Based on the verification results, we divide them into two groups and separately evaluate their generation quality. We observe that on questions whose documents fail the verification, the LLM shows significantly worse generation quality in correctness and verifiability. And after retrieval update, the updated retrieval result leads to better generation quality. These demonstrate that the retrieval verification and update can identify and improve low-quality retrieval results respectively, and thus lead to better verifiable generation. In general, these results show that each component of LLaTrieval contributes critically to the overall performance.

Performance over Different Thresholds We evaluate the score-and-filter mode of retrieval verification on ASQA and QAMPARI. We first plot the generation quality on examples (question & docs

	ASQA			QAMPARI		
	Num	EM-R	Cite	Num	F1	Cite
BGE-large	948	55.9	57.5	1000	17.3	24.0
+ Progressive Selection	948	56.8	60.2	1000	23.0	32.3
+ Verification ✓	779	60.3	64.5	827	25.8	36.6
+ Verification ✗	169	41.0	40.0	173	10.1	9.1
+ Update ↑	169	44.3	44.3	173	17.7	21.1

Table 2: The effect of each component. “Verification ✓ / ✗” represent the question whose documents pass / fail the retrieval verification, respectively. “Update ↑” represents the questions paired with updated retrieval result. “Num” is the corresponding example quantity.

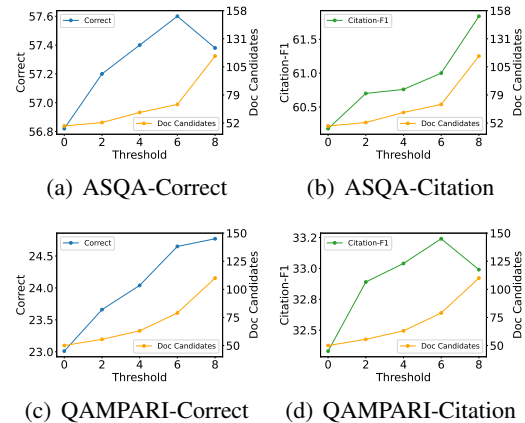


Figure 6: The performance over different thresholds.

pairs) filtered by different thresholds in the first iteration, shown in Figure 5. We see that the examples filtered by the higher threshold have better generation quality, which shows that our verification score can effectively evaluate the retrieval result. Then we plot the performance and the LLM cost of our method over different thresholds in Figure 6, where we use each example’s overall document candidates quantity to measure the LLM cost since their tokens predominate in the LLM’s input. We see that increasing the threshold leads to better performance with higher LLM cost trade-off. This shows that our method’s speed, i.e., average iterations, and LLM cost can be flexibly adjusted to adapt to varying demands. Meanwhile, compared with reranking baselines, which do not show significant improvements with increasing document candidates (see Table 1), we can effectively scale our method’s performance by stricter threshold and more document candidates from missing-info query. Additionally, we observe that a too-strict threshold may not necessarily improve performance, e.g., “8” for ASQA’s correctness. We observe that its failure is because there is no relevant information in the

Dataset	QAMPARI		ELI5	
Metric	F1	Cite	Claim	Cite
Original Retriever	17.3	24.0	15.2	66.7
GPT-3.5-Turbo-0301	24.3	34.4	16.5	71.3
GPT-3.5-Turbo-0613	24.8	32.9	17.1	71.7
GPT-3.5-Turbo-1106	23.9	33.6	17.3	71.7
Llama-2-70B-Chat	21.0	28.9	15.8	61.0
Xwin-LM-70B-V0.1	22.5	30.8	16.1	70.8
Tulu-2-DPO-70B	22.7	32.0	16.9	71.8

Table 3: The results of LLattribution across various LLMs.

document pool or the LLM fails to follow the instruction of missing-info query, and we regard the research of these two aspects as future work.

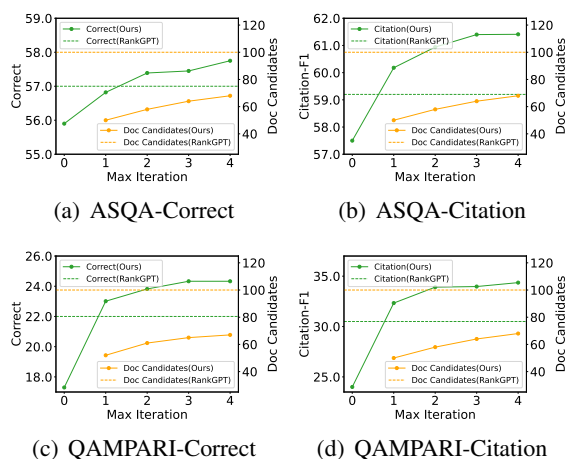


Figure 7: The performance over different max iteration.

Performance during Iteration We observe the performance changes during verify-update iterations on ASQA and QAMPARI. Specifically, we plot the performance curve and each example’s average document candidate quantity over different maximum iterations in Figure 7. We see that the performance of LLattribution gets improved with increasing max iterations, which shows that the document quality can be gradually improved during the multiple verify-update iterations. Meanwhile, we observe that although the average document candidates slightly increase with increasing max iterations, our method can outperform RankGPT with fewer average document candidates, which shows LLattribution’s efficiency that it can dynamically finish when the retrieved documents can sufficiently support answering the question.

The Comparison with Gold Feedback Recent work indicates that the feedback from the LLM can not help improve the prediction of itself (Huang

	ASQA		
	Doc Candidates	Correct	Cite
Our Method	68.9	57.7	61.4
Our Method (with gold)	136.9	58.1	60.8

	QAMPARI		
	Doc Candidates	Correct	Cite
Our Method	64.5	24.3	34.4
Our Method (with gold)	101.0	25.2	33.6

Table 4: The comparison with using gold answers for retrieval verification.

et al., 2023a; Stechly et al., 2023; Valmeekam et al., 2023). Since we use the LLM to verify the retrieval result updated by itself, and such verification can be seen as a type of internal feedback, we compare it with the external feedback from gold answers on ASQA and QAMPARI. Specifically, for external feedback, we take whether gold answers fully appear in the retrieved result as verification criteria. We show the results in Table 4. We see that both internal and external feedbacks achieve comparable performance with each other, when the internal feedback leads to less LLM cost. These show that the LLM is capable of evaluating the retrieval result and has the potential to provide competitive feedback for retrieval as gold answers.

The Generality across Various LLMs We analyze the performance of LLattribution over various closed-source and open-source LLMs (Touvron et al., 2023; Team, 2023; Ivison et al., 2023), shown in Table 3. In general, our method consistently leads to significant improvements over the original retriever, which shows its robustness and generality. Meanwhile, we find that our method with stronger LLM (from Llama2-70B to GPT-3.5) can lead to better performance, which indicates the potential to scale retrieval performance by scaling LLM and our method’s utility in the future where more powerful LLMs are built, while dense retrievers have not been proved to scale well to the size of LLMs.

5 Related Work

Verifiable Generation Verifiable generation, which aims to generate content with supporting documents, has drawn increasing attention in recent years. The previous work mainly focuses on two aspects: 1) **Modeling**: Nakano et al. (2021), Menick et al. (2022), Liu et al. (2023c) and Qin et al. (2023a) explore training the specialized LLM

to browse web pages and answer long-form questions with supporting evidence. Gao et al. (2023a) propose research-and-revision to retrieve supporting evidence for LLM’s output and fix unsupported content. Li et al. (2023b) explore verifiable generation on knowledge graph. 2) **Verifiability Evaluation:** Rashkin et al. (2021) define *Attributable to Identified Sources* (AIS), a framework of human evaluation that considers whether the answer can be supported by the evidence. Gao et al. (2023a) propose auto-AIS, which utilizes a strong NLI model (Honovich et al., 2022) to approximate human AIS judgments. Liu et al. (2023b) propose citation-recall and citation-precision to evaluate the verifiability of prevailing generative search engines, e.g., New Bing. Gao et al. (2023c) propose a framework that automates the evaluation of recall and precision for the citation in the LLM’s output. Since retrieval is important for finding supporting evidence, we propose LLattribution to iteratively update and verify the retrieval result, and thus fully harness the LLM for better verifiable generation.

LLM-based Retrieval Augmentation The existing methods of LLM-augmented Retrieval can be divided into three categories: 1) **Query Rewriting:** Gao et al. (2023b); Wang et al. (2023b) propose to use the LLM to generate a pseudo passage of the question as retrieval query, which neighbors relevant documents and helps retrieve them. Shen et al. (2023) use the retrieved passage to help the LLM generate better pseudo passage as retrieval query. Shao et al. (2023) propose to iteratively improve the retrieval query by synergizing retrieval and generation. Ma et al. (2023) train specialized query rewriters for downstream RAG tasks. 2) **Re-ranking:** Sun et al. (2023) propose RankGPT, which uses ChatGPT to generate the permutation of documents for reranking. Tang et al. (2023) and Qin et al. (2023b) use multiple-sampling and pair-wise ranking to mitigate the positional bias of RankGPT, respectively. 3) **Data Augmentation:** Bonifacio et al. (2022), Dai et al. (2023) and Jeronimo et al. (2023) use the LLM to generate relevant queries for existing documents for training the retriever. Sachan et al. (2022) use the LLM to provide soft labels of unpaired question&documents for the retriever’s training. Although these methods use the LLM to enhance the retrieval from various aspects, the LLM can not fully provide feedback to the retrieval result and the overall performance is highly limited by the retriever. For example, the query rewrit-

ing methods take the pseudo passage as query in a heuristic manner, which relies on the generated passage quality and can not ensure the relevance of retrieved documents. In LLattribution, the LLM can iteratively provide feedback to the retriever until the documents sufficiently support verifiable generation, which can fully harness the LLM’s ability to facilitate the retrieval stage.

6 Conclusion

We propose LLattribution to retrieve supporting documents for verifiable generation, which lets the LLM iteratively refine the retrieval result until the LLM verifies the retrieved documents can support answering the given question. In this manner, LLattribution can fully harness the LLM’s ability and facilitate the retrieval result to help generate both correct and verifiable answers. Experimental results show that our method outperforms extensive baselines and achieves new state-of-the-art results.

Limitations

- Although LLattribution can dynamically adjust the running time according to the retrieval demand and shows better performance-efficiency trade-off than RankGPT (Sun et al., 2023), it relies on real-time LLM inference and thus may not be suitable for tasks that require low latency or high throughput. However, over the years we have seen the model acceleration mechanisms advance and hardware performance increase, which can help improve the efficiency of LLM inference.
- Since existing works have shown LLMs exhibit various types of biases (Bender et al., 2021), LLattribution based on LLMs, may bias the final retrieval result. We are optimistic that this problem can be mitigated as LLattribution is based on ChatGPT, and OpenAI has made substantial efforts to reduce its toxicity and bias (Ouyang et al., 2022). Additionally, we can use more elaborate prompts on LLMs to further reduce the bias in the process of LLattribution.

Acknowledgments

This work was supported by the National Key Research and Development Program of China (No.2022CSJGG0801). The computations in this research were performed using the CFFF platform of Fudan University.

References

- Samuel Joseph Amouyal, Ohad Rubin, Ori Yoran, Tomer Wolfson, Jonathan Herzig, and Jonathan Berant. 2022. [QAMPARI: : An open-domain question answering benchmark for questions with many answers from multiple paragraphs](#). *CoRR*, abs/2205.12665.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *FAccT '21: 2021 ACM Conference on Fairness, Accountability, and Transparency, Virtual Event / Toronto, Canada, March 3-10, 2021*, pages 610–623. ACM.
- Bernd Bohnet, Vinh Q. Tran, Pat Verga, Roei Aharoni, Daniel Andor, Livio Baldini Soares, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, Kai Hui, Tom Kwiatkowski, Ji Ma, Jianmo Ni, Tal Schuster, William W. Cohen, Michael Collins, Dipanjan Das, Donald Metzler, Slav Petrov, and Kellie Webster. 2022. [Attributed question answering: Evaluation and modeling for attributed large language models](#). *CoRR*, abs/2212.08037.
- Luiz Henrique Bonifacio, Hugo Queiroz Abonizio, Marzieh Fadaee, and Rodrigo Frassetto Nogueira. 2022. [Inpars: Data augmentation for information retrieval using large language models](#). *CoRR*, abs/2202.05144.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Sihao Chen, Hongming Zhang, Tong Chen, Ben Zhou, Wenhao Yu, Dian Yu, Baolin Peng, Hongwei Wang, Dan Roth, and Dong Yu. 2023. [Sub-sentence encoder: Contrastive learning of propositional semantic representations](#). *CoRR*, abs/2311.04335.
- I-Chun Chern, Steffi Chern, Shiqi Chen, Weizhe Yuan, Kehua Feng, Chunting Zhou, Junxian He, Graham Neubig, and Pengfei Liu. 2023. [Factool: Factuality detection in generative AI - A tool augmented framework for multi-task and multi-domain scenarios](#). *CoRR*, abs/2307.13528.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2023. [Palm: Scaling language modeling with pathways](#). *J. Mach. Learn. Res.*, 24:240:1–240:113.
- Zhuyun Dai, Vincent Y. Zhao, Ji Ma, Yi Luan, Jianmo Ni, Jing Lu, Anton Bakalov, Kelvin Guu, Keith B. Hall, and Ming-Wei Chang. 2023. [Promptagator: Few-shot dense retrieval from 8 examples](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: long form question answering](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 3558–3567. Association for Computational Linguistics.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Y. Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023a. [RARR: researching and revising what language models say, using language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 16477–16508. Association for Computational Linguistics.
- Luyu Gao, Xueguang Ma, Jimmy Lin, and Jamie Callan. 2023b. [Precise zero-shot dense retrieval without relevance labels](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 1762–1777. Association for Computational Linguistics.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023c. [Enabling large language models to generate text with citations](#). *CoRR*, abs/2305.14627.
- Or Honovich, Roei Aharoni, Jonathan Herzig, Hagai Taitelbaum, Doron Kukliansy, Vered Cohen, Thomas Scialom, Idan Szpektor, Avinatan Hassidim, and Yossi Matias. 2022. [TRUE: re-evaluating factual](#)

- consistency evaluation. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 3905–3920. Association for Computational Linguistics.
- Jie Huang, Xinyun Chen, Swaroop Mishra, Huaixiu Steven Zheng, Adams Wei Yu, Xinying Song, and Denny Zhou. 2023a. [Large language models cannot self-correct reasoning yet](#). *CoRR*, abs/2310.01798.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023b. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *CoRR*, abs/2311.05232.
- Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep Dasigi, Joel Jang, David Wadden, Noah A. Smith, Iz Beltagy, and Hannaneh Hajishirzi. 2023. [Camels in a changing climate: Enhancing lm adaptation with tulu 2](#).
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. [Unsupervised dense information retrieval with contrastive learning](#). *Trans. Mach. Learn. Res.*, 2022.
- Vitor Jeronimo, Luiz Henrique Bonifacio, Hugo Queiroz Abonizio, Marzieh Fadaee, Roberto de Alencar Lotufo, Jakub Zavrel, and Rodrigo Frassetto Nogueira. 2023. [Inpars-v2: Large language models as efficient dataset generators for information retrieval](#). *CoRR*, abs/2301.01820.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6769–6781. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Dongfang Li, Zetian Sun, Xinshuo Hu, Zhenyu Liu, Ziyang Chen, Baotian Hu, Aiguo Wu, and Min Zhang. 2023a. [A survey of large language models attribution](#). *CoRR*, abs/2311.03731.
- Xinze Li, Yixin Cao, Liangming Pan, Yubo Ma, and Aixin Sun. 2023b. [Towards verifiable generation: A benchmark for knowledge-aware language model attribution](#). *CoRR*, abs/2310.05634.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023a. [Lost in the middle: How language models use long contexts](#). *CoRR*, abs/2307.03172.
- Nelson F. Liu, Tianyi Zhang, and Percy Liang. 2023b. [Evaluating verifiability in generative search engines](#). *CoRR*, abs/2304.09848.
- Xiao Liu, Hanyu Lai, Hao Yu, Yifan Xu, Aohan Zeng, Zhengxiao Du, Peng Zhang, Yuxiao Dong, and Jie Tang. 2023c. [Webglm: Towards an efficient web-enhanced question answering system with human preferences](#). In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6-10, 2023*, pages 4549–4560. ACM.
- Xinbei Ma, Yeyun Gong, Pengcheng He, Hai Zhao, and Nan Duan. 2023. [Query rewriting for retrieval-augmented large language models](#). *CoRR*, abs/2305.14283.
- Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2023. [Expertqa: Expert-curated questions and attributed answers](#). *CoRR*, abs/2309.07852.
- Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, H. Francis Song, Martin J. Chadwick, Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, and Nat McAleese. 2022. [Teaching language models to support answers with verified quotes](#). *CoRR*, abs/2203.11147.
- Sewon Min, Kalpesh Krishna, Xixi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. [Factscore: Fine-grained atomic evaluation of factual precision in long form text generation](#). *CoRR*, abs/2305.14251.
- Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Long Ouyang, Christina Kim, Christopher Hesse, Shantanu Jain, Vineet Kosaraju, William Saunders, Xu Jiang, Karl Cobbe, Tyna Eloundou, Gretchen Krueger, Kevin Button, Matthew Knight, Benjamin Chess, and John Schulman. 2021. [Webgpt: Browser-assisted question-answering with human feedback](#). *CoRR*, abs/2112.09332.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, and Yinfei Yang. 2022. [Large dual encoders are generalizable retrievers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 9844–9855, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rodrigo Nogueira, Zhiying Jiang, Ronak Pradeep, and Jimmy Lin. 2020. [Document ranking with a pre-trained sequence-to-sequence model](#). In *Findings*

- of the Association for Computational Linguistics: EMNLP 2020*, pages 708–718, Online. Association for Computational Linguistics.
- Rodrigo Frassetto Nogueira and Kyunghyun Cho. 2019. [Passage re-ranking with BERT](#). *CoRR*, abs/1901.04085.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Long Ouyang, Jeffrey 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 Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *NeurIPS*.
- Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Dmytro Okhonko, Samuel Broscheit, Gautier Izacard, Patrick S. H. Lewis, Barlas Oguz, Edouard Grave, Wen-tau Yih, and Sebastian Riedel. 2021. [The web is your oyster - knowledge-intensive NLP against a very large web corpus](#). *CoRR*, abs/2112.09924.
- Yujia Qin, Zihan Cai, Dian Jin, Lan Yan, Shihao Liang, Kunlun Zhu, Yankai Lin, Xu Han, Ning Ding, Huadong Wang, Ruobing Xie, Fanchao Qi, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2023a. [Webcpm: Interactive web search for chinese long-form question answering](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 8968–8988. Association for Computational Linguistics.
- Zhen Qin, Rolf Jagerman, Kai Hui, Honglei Zhuang, Junru Wu, Jiaming Shen, Tianqi Liu, Jialu Liu, Donald Metzler, Xuanhui Wang, and Michael Bendersky. 2023b. [Large language models are effective text rankers with pairwise ranking prompting](#). *CoRR*, abs/2306.17563.
- Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2021. [Measuring attribution in natural language generation models](#). *CoRR*, abs/2112.12870.
- Vipula Rawte, Amit P. Sheth, and Amitava Das. 2023. [A survey of hallucination in large foundation models](#). *CoRR*, abs/2309.05922.
- Revanth Gangi Reddy, Yi R. Fung, Qi Zeng, Manling Li, Ziqi Wang, Paul Sullivan, and Heng Ji. 2023. [Smartbook: Ai-assisted situation report generation](#). *CoRR*, abs/2303.14337.
- Stephen E. Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Devendra Singh Sachan, Mike Lewis, Dani Yogatama, Luke Zettlemoyer, Joelle Pineau, and Manzil Zaheer. 2022. [Questions are all you need to train a dense passage retriever](#). *CoRR*, abs/2206.10658.
- M. Salvagno, F.S. Taccone, A.G. Gerli, and ChatGPT. 2023. [Can artificial intelligence help for scientific writing?](#) *Critical Care*, 27(1). Publisher: BioMed Central Ltd.
- Tal Schuster, Adam Fisch, and Regina Barzilay. 2021. [Get your vitamin c! robust fact verification with contrastive evidence](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pages 624–643. Association for Computational Linguistics.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. [Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy](#). *CoRR*, abs/2305.15294.
- Tao Shen, Guodong Long, Xiubo Geng, Chongyang Tao, Tianyi Zhou, and Daxin Jiang. 2023. [Large language models are strong zero-shot retriever](#). *CoRR*, abs/2304.14233.
- Kaya Stechly, Matthew Marquez, and Subbarao Kambhampati. 2023. [GPT-4 doesn't know it's wrong: An analysis of iterative prompting for reasoning problems](#). *CoRR*, abs/2310.12397.
- Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-Wei Chang. 2022. [ASQA: factoid questions meet long-form answers](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 8273–8288. Association for Computational Linguistics.
- Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2023. [One embedder, any task: Instruction-finetuned text embeddings](#). In *Findings of the Association for Computational Linguistics: ACL 2023, Toronto, Canada, July 9-14, 2023*, pages 1102–1121. Association for Computational Linguistics.
- Weiwei Sun, Lingyong Yan, Xinyu Ma, Pengjie Ren, Dawei Yin, and Zhaochun Ren. 2023. [Is chatgpt good at search? investigating large language models as re-ranking agent](#). *CoRR*, abs/2304.09542.
- Raphael Tang, Xinyu Zhang, Xueguang Ma, Jimmy Lin, and Ferhan Ture. 2023. [Found in the middle: Permutation self-consistency improves listwise ranking in large language models](#). *CoRR*, abs/2310.07712.
- Xwin-LM Team. 2023. [Xwin-lm](#).

- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and finetuned chat models](#). *CoRR*, abs/2307.09288.
- Karthik Valmeekam, Matthew Marquez, and Subbarao Kambhampati. 2023. [Can large language models really improve by self-critiquing their own plans?](#) *CoRR*, abs/2310.08118.
- Cunxiang Wang, Xiaoze Liu, Yuanhao Yue, Xiangru Tang, Tianhang Zhang, Jiayang Cheng, Yunzhi Yao, Wenyang Gao, Xuming Hu, Zehan Qi, Yidong Wang, Linyi Yang, Jindong Wang, Xing Xie, Zheng Zhang, and Yue Zhang. 2023a. [Survey on factuality in large language models: Knowledge, retrieval and domain-specificity](#). *CoRR*, abs/2310.07521.
- Liang Wang, Nan Yang, and Furu Wei. 2023b. [Query2doc: Query expansion with large language models](#). *CoRR*, abs/2303.07678.
- Zhiruo Wang, Jun Araki, Zhengbao Jiang, Md. Rizwan Parvez, and Graham Neubig. 2023c. [Learning to filter context for retrieval-augmented generation](#). *CoRR*, abs/2311.08377.
- Shitao Xiao, Zheng Liu, Peitian Zhang, and Niklas Muennighof. 2023. [C-pack: Packaged resources to advance general chinese embedding](#). *CoRR*, abs/2309.07597.
- Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2023. [Making retrieval-augmented language models robust to irrelevant context](#). *CoRR*, abs/2310.01558.
- Hang Zhang, Yeyun Gong, Yelong Shen, Jiancheng Lv, Nan Duan, and Weizhu Chen. 2022a. [Adversarial retriever-ranker for dense text retrieval](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Shunyu Zhang, Yaobo Liang, Ming Gong, Daxin Jiang, and Nan Duan. 2022b. [Multi-view document representation learning for open-domain dense retrieval](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 5990–6000. Association for Computational Linguistics.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. [Siren’s song in the AI ocean: A survey on hallucination in large language models](#). *CoRR*, abs/2309.01219.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. [A survey of large language models](#). *CoRR*, abs/2303.18223.

A Implementation Details

A.1 Method Details

We use the public OpenAI language model API of “gpt-3.5-turbo-0301” for retrieval and generation, unless otherwise specified. We set the temperature as 0 to mitigate the random fluctuation. We use window size=20 and set the number of document candidates per query as 50. We set the number of supporting documents and the maximum of iteration as 5 and 4. For the document pool, we follow Gao et al. (2023c) to use the Wikipedia for ASQA and QAMPARI, and Sphere (Piktus et al., 2021), a filtered version of a filtered version of Common Crawl² for ELI5, respectively. For ASQA and QAMPARI, we use the recently proposed dense embedder, BGE-large (Xiao et al., 2023), as the retriever. For ELI5, we follow Gao et al. (2023c) to use BM25 (Robertson and Zaragoza, 2009) since dense retrievers are costly and slow for large-scale web corpus.

In progressive selection, we put the current documents before the new list of candidates in the LLM’s input. Inspired by Gao et al. (2023c), we use the document summary generated by LLM in progressive selection and thus we can significantly reduce the overall input length when the LLM select documents, which alleviates the challenges of LLMs finding relevant content in long context (Liu et al., 2023a). For ELI5, we use the HYDE (Gao et al., 2023b) to generate the first iteration’s query to get the better document candidates.

Compared Method Details For HYDE, LAMER and RankGPT (Sun et al., 2023), we use its original instructions and the LLM used in LLa-trieval for fair comparison. For DPR, we follow Gao et al. (2023c) to use its original checkpoints on NQ (Kwiatkowski et al., 2019). For Contriever, we use its unsupervised version³ and supervised version⁴ fine-tuned on MS MARCO. For GTR (Ni et al., 2022), we follow Gao et al. (2023c) to use its T5-XXL verison⁵. For instructor (Su et al., 2023), we use its large⁶ version. For BGE-Embedding, we use its large⁷ version. Following Sun et al. (2023), we use the 340M version of monoBERT⁸,

²<https://commoncrawl.org>.

³<https://huggingface.co/facebook/contriever>

⁴<https://huggingface.co/facebook/contriever-msmarco>

⁵<https://huggingface.co/sentence-transformers/gtr-t5-xxl>

⁶<https://huggingface.co/hkunlp/instructor-large>

⁷<https://huggingface.co/BAAI/bge-large-en-v1.5>

⁸<https://huggingface.co/castorini/monobert-large-msmarco>

	ASQA	QAMPARI	ELI5
Total Eexample	948	1000	1000
Document Candidates	68.9	67.5	74.5
The number of examples in each iteration			
Iteration 1	948	1000	1000
Iteration 2	169	173	234
Iteration 3	105	70	144
Iteration 4	84	46	111

Table 5: The iteration details and average document candidate quantity for our method (passage-style missing-info query).

	ASQA	QAMPARI	ELI5
Total Eexample	948	1000	1000
Document Candidates	69.2	66.9	76.0
The number of examples in each iteration			
Iteration 1	948	1000	1000
Iteration 2	169	173	234
Iteration 3	112	91	159
Iteration 4	83	73	128

Table 6: The iteration details and average document candidate quantity for our method (question-style missing-info query).

Dataset	ASQA		QAMPARI	
Correctness	Retrieval	Generation	Retrieval	Generation
BM25	47.7	51.7	29.7	17.5
Instructor-large	57.0	55.6	27.8	17.2
BGE-E-large	60.6	55.9	29.1	17.3
monoBERT	60.8	53.8	35.7	20.9
BGE-R-large	62.4	55.3	35.8	20.7
RankGPT	62.1	56.4	34.4	20.7
LLatriveal	62.9	57.7	38.3	24.3

Table 7: The correlation between retrieval quality and generation quality.

and 3B⁹ verison of monoT5. For BGE-Reranker, we use its large¹⁰ version for comparison.

A.2 Instructions

We show the overall instructions in Table 9, 10, 11, 12, 13.

A.3 Average Document Candidates

In Table 5 and 6, we show the average document candidate quantity and the number of examples that pass the retrieval verification in each iteration, for the result in Table 1.

⁹<https://huggingface.co/castorini/monot5-3b-msmarco-10k>

¹⁰<https://huggingface.co/BAAI/bge-reranker-large>

Dataset	ASQA				QAMPARI			
	Correct		Citation		Correct		Citation	
	EM-R	Rec	Prec	F1	F1	Rec	Prec	F1
Original Retriever	55.9	56.4	58.6	57.5	17.3	22.7	25.5	24.0
Window Size=15	56.9	60.2	59.6	59.9	24.3	31.6	35.7	33.5
Window Size=20	57.7	60.9	61.3	61.1	24.3	32.4	36.6	34.4
Window Size=25	57.2	59.6	60.1	59.8	24.0	32.5	37.1	34.6
Window Size=30	57.4	60.3	60.6	60.4	24.7	32.4	36.3	34.2
GTR	54.4	52.2	54.6	53.3	18.5	20.0	23.2	21.5
+ LLattribution	57.6	59.0	60.5	59.7	22.6	27.1	31.9	29.3
Instructor-large	55.6	52.1	54.1	53.1	17.2	19.4	21.9	20.6
+ LLattribution	57.6	60.1	60.2	60.2	21.1	24.7	29.1	26.7
BGE-E-Base	55.3	56.2	58.2	57.2	17.9	22.9	24.8	23.8
+ LLattribution	56.7	58.5	59.3	58.9	25.0	32.6	37.3	34.8
Original Retriever ($ D = 3$)	55.5	52.4	57.0	54.6	15.1	21.3	23.9	22.5
+ LLattribution ($ D = 3$)	56.6	57.9	61.9	59.9	23.1	31.3	35.0	33.0
Original Retriever ($ D = 7$)	56.9	59.0	58.4	58.7	17.5	22.6	25.6	24.0
+ LLattribution ($ D = 7$)	57.1	61.9	61.0	61.5	24.6	31.2	35.7	33.3

Table 8: The performance of LLattribution over various hyper-parameters and retrievers.

B The Correlation between Retrieval Quality and Generation Quality

In exploratory experiments, we observe the correctness of retrieval result and LLM’s generation when using various retrieval methods. Specifically, we conduct experiments on ASQA and QAMPARI, and evaluate the exact match recall of retrieval result and the corresponding LLM’s output, as shown in Table 7. We see that the higher retrieval quality leads to higher generation quality, which shows the positive correlation between retrieval quality and the generation quality.

C Performance across Various Hyper-Parameters

We run LLattribution with different hyper-parameters and retrievers, and show the results in Table 8. We see that LLattribution consistently outperforms the original retriever under the different window sizes and document quantities. Meanwhile, LLattribution can lead to performance improvements based on various retrievers. These show the robustness and generality of LLattribution.

You are JudgeGPT as introduced below.

Role: JudgeGPT

Profile

- Language: English
- Description: You are JudgeGPT, capable of judging whether a specified number (k) of documents can maximally support giving a direct, accurate, clear and engaging answer, similar to the answer of the demonstration, closely related to the core of the user's specific question(s).

Demonstration

{Demo}

Input

- Question: The specific question(s).
- Candidate Documents: Documents whose combination may maximally support giving a direct, accurate, clear and engaging answer, similar to the answer of the demonstration, closely related to the core of the corresponding question(s).

Skill

1. Analyzing the given question(s) and understanding the required information.
2. Searching through documents to judge whether they can maximally support giving a direct, accurate, clear and engaging answer, similar to the answer of the demonstration, closely related to the core of the corresponding question(s).

Output

- Judgment: "[YES]" if provided documents can maximally support giving a direct, accurate, clear, and engaging answer, similar to the answer of the demonstration, closely related to the core of the corresponding question(s), otherwise "[NO]".

Output Format

Judgment: [YES] or [NO]

Output Example

If provided documents can maximally support giving a direct, accurate, clear, and engaging answer, similar to the answer of the demonstration, closely related to the core of the corresponding question(s), the output should be as follows: [YES]

Rules

1. Don't break character.
2. When outputting final verdict, only providing "[YES]" or "[NO]".
3. Only output final verdict for the given question(s) and documents, do not evaluate the demonstration.
4. Strictly follow the specified output format. Do not answer the given question. Just conduct the specified judgment task.

Judgment Criteria (Very Important)

1. Do not allow the length of the documents to influence your evaluation.
2. Be as objective as possible.
3. Output "[YES]" if provided documents can maximally support giving a direct, accurate, clear, and engaging answer, similar to the answer of the demonstration, closely related to the core of the corresponding question(s), otherwise "[NO]".

Workflow

1. Read and understand the questions posed by the user.
2. Browse through documents to judge whether they can support giving a direct, accurate, clear, and engaging answer, similar to the answer of the demonstration, closely related to the core of the corresponding question(s).
3. Output your final verdict.

Reminder

You will always remind yourself of the role settings.

Table 9: The instruction for retrieval verification based on classification.

You are ScoreGPT as introduced below.

Role: ScoreGPT

Profile

- Language: English
- Description: You are ScoreGPT, capable of scoring candidate documents based on their level of support for the corresponding question(s), with a rating range from 0 to 10.

Input

- Question: The specific question(s).
- Candidate Documents: Documents whose combination may maximally support the corresponding question(s).

Skill

1. Analyzing the given question(s) and understanding the required information.
2. Searching through documents to score them based on their level of support for the corresponding question(s), with a rating range from 0 to 10.

Output

- A score ranging from 0 to 10, where a higher score indicates greater support of the candidate documents for the corresponding question(s), and a lower score indicates lesser support.

Output Format

[SCORE]

Rules

1. Don't break character.
2. When outputting final score, only providing "[SCORE]".
3. Strictly follow the specified output format. Do not answer the given question(s). Just conduct the specified scoring task.

Score Criteria (Very Important)

1. Do not allow the length of the documents to influence your evaluation.
2. Be as objective as possible.
3. Output "[SCORE]" ranging from 0 to 10, where a higher score indicates greater support of the candidate documents for the corresponding question(s), and a lower score indicates lesser support.

Workflow

1. Read and understand the question(s) posed by the user.
2. Browse through documents to score them based on their level of support for the corresponding question(s), with a rating range from 0 to 10.
3. Output your final score surrounded by square brackets.

Reminder

You will always remind yourself of the role settings.

Table 10: The instruction for retrieval verification based on score-and-filter.

You are DocSelectorGPT as introduced below.

Role: DocSelectorGPT

Profile

- Language: English
- Description: You are DocSelectorGPT, capable of selecting a specified number (k) of documents for answering the user's specific question(s). k is a value specified by the user.

Input

- Question: The specific question(s)
- Candidate Documents: Documents contain supporting documents which can support answering the given questions. Candidate documents will have their own identifiers for FactRetrieverGPT to cite.

Skill

1. Analyzing the given question(s) and understanding the required information.
2. Searching through candidate documents to select k supporting documents whose combination can maximally support giving a direct, accurate, clear and engaging answer to the question and make the answer and is closely related to the core of the question.

Output

- Selected Documents: The identifiers of selected supporting documents whose combination can maximally support giving an accurate and engaging answer to the question and make the answer and is closely related to the core of the question.

Output Format

Selected Documents: [document identifiers]

Output Example

If the selected documents are 2, 6 and 8, the output should be as follows:

Selected Documents: 2 6 8

Rules

1. Don't break character.
2. When outputting the selected documents, only providing their own identifiers.
3. Strictly follow the specified output format. Do not answer the given question. Just conduct the specified retrieval task.

Selection Criteria (Very Important)

1. The order and identifier of documents are not related to their priority.
2. Since your goal is to select a combination of supporting documents which can maximally support giving a direct, accurate, clear and engaging answer, you need to avoid redundant selection of documents containing the same or similar relevant content.

Workflow

1. Read and understand the questions posed by the user.
2. Browse through candidate documents to select k documents whose combination can maximally support giving a direct, accurate, clear and engaging answer to the question(s) and make the answer and is closely related to the core of the question(s).
3. List all selected documents.

Reminder

You will always remind yourself of the role settings.

Table 11: The instruction for LLM to select documents in progressive selection.

You are a helpful assistant as introduced below.

Profile

- Language: English
- Description: You are a helpful assistant, capable of identifying missing content that answers the given question(s) but does not exist in the given possible answering passages and then using your own knowledge to generate correct answering passages using missing content you identify.

Input

- Question: The specific question(s).
- Answering Passages: Possible answering passages.

Output

- Correct answering passages generated using missing content you identify based on your own knowledge.

Rules

1. Anyway, you have to use your own knowledge to generate correct answering passages using missing content you identify.
2. Only generate the required correct answering passages. Do not output anything else.
3. Directly use your own knowledge to generate correct answering passages if you think the given possible answering passages do not answer to the given question(s).
4. Do not output the given question(s) and possible answering passages.
5. Do not output your analysis statement.

Workflow

1. Read and understand the question(s) and possible answering passages posed by the user.
2. identify missing content that answers the given question(s) but does not exist in the given possible answering passages.
3. Directly use your own knowledge to generate correct answering passages if you think the given possible answering passages do not answer to the given question(s). Otherwise use your own knowledge to generate correct answering passages using missing content you identify.

Reminder

You will always remind yourself of the role settings.

Table 12: The Instruction for Missing-info Query (the passage style).

You are a helpful assistant as introduced below.

Profile

- Language: English
- Description: You are a helpful assistant, capable of identifying missing content that answers the given question(s) but does not exist in the given possible answering passages and then using your own knowledge to generate a new question based on the missing content you identify.

Input

- Question: The specific question(s).
- Answering Passages: Possible answering passages.

Output

- A new question generated using missing content you identify based on your own knowledge.

Rules

1. Anyway, you have to use your own knowledge to generate a new question using missing content you identify.
2. Only generate the required new question. Do not output anything else.
3. Do not output the given question(s) and possible answering passages.
4. Do not output your analysis statement.

Workflow

1. Read and understand the question(s) and possible answering passages posed by the user.
2. Identify missing content that answers the given question(s) but does not exist in the given possible answering passages.
3. Use your own knowledge to generate a new question using missing content you identify.

Reminder

You will always remind yourself of the role settings.

Table 13: The instruction for Missing-info Query (the question style).