

Small Models, Big Insights: Leveraging Slim Proxy Models To Decide When and What to Retrieve for LLMs

Anonymous ACL submission

Abstract

The integration of large language models (LLMs) and search engines represents a significant evolution in knowledge acquisition methodologies. However, determining the knowledge that an LLM already possesses and the knowledge that requires the help of a search engine remains an unresolved issue. Most existing methods solve this problem through the results of preliminary answers or reasoning done by the LLM itself, but this incurs excessively high computational costs. This paper introduces a novel collaborative approach, namely SlimPLM, that detects missing knowledge in LLMs with a slim proxy model, to enhance the LLM’s knowledge acquisition process. We employ a proxy model which has far fewer parameters, and take its answers as heuristic answers. Heuristic answers are then utilized to predict the knowledge required to answer the user question, as well as the known and unknown knowledge within the LLM. We only conduct retrieval for the missing knowledge in questions that the LLM does not know. Extensive experimental results on five datasets with two LLMs demonstrate a notable improvement in the end-to-end performance of LLMs in question-answering tasks, achieving or surpassing current state-of-the-art models with lower LLM inference costs.¹

1 Introduction

Large language models (LLMs) have demonstrated significant prowess in various natural language processing (NLP) tasks (Achiam et al., 2023), attributed to their advanced language comprehension and generation capabilities. Despite being trained on extensive text corpora, these models occasionally produce hallucinated content (Zhou et al., 2020; Maynez et al., 2020). To tackle this problem, the integration of retrieval systems with

LLMs has been proposed, enabling access to external knowledge bases for more accurate and reliable text generation.

Retrieval-augmented generation (RAG) involves using a retrieval system to supplement LLMs with relevant external information, thereby improving text generation quality (Peng et al., 2023; He et al., 2022). Yet, recent studies have suggested that retrieval may not always be beneficial. In cases where LLMs can adequately respond without external knowledge, retrieval may introduce irrelevant information, potentially degrading performance (Kadavath et al., 2022; Wang et al., 2023b; Shi et al., 2023a; Petroni et al., 2020a). Therefore, it is critical to determine when retrieval is necessary for user questions (Shuster et al., 2021). The challenge lies in identifying questions that exceed the LLMs’ intrinsic knowledge and require external retrieval, due to the prevalence of content hallucination. Efforts to address this challenge can be categorized into two groups: (1) The first group of methods involves fine-tuning LLMs for RAG scenarios, allowing them to autonomously signal the need for external knowledge (Nakano et al., 2021; Liu et al., 2023b; Qin et al., 2023b). This method, while effective, demands substantial computational resources and risks diminishing the LLMs’ general capabilities due to potential catastrophic forgetting (Kotha et al., 2023; Zhai et al., 2023). (2) The second category avoids direct tuning of LLMs, assessing the necessity for retrieval based on the quality of the generated content or specific indicators within it (Ram et al., 2023; Min et al., 2022). However, this approach still has its drawbacks, as it requires multiple inferences, thereby increasing both the inference costs and the latency of responses to user questions.

In light of this, we put forward a question: *Is it feasible to employ a proxy model with a relatively smaller parameter size to facilitate effective retrieval results for an LLM?* Theoretically, existing

¹Anonymous code and datasets are available at <https://anonymous.4open.science/r/SlimPlm-2434>

decoder-only language models share similar Transformer structures, and they are pre-trained on some common text corpora, such as Common Crawl web pages, books, and Wikipedia pages (Touvron et al., 2023; Bai et al., 2023; Scao et al., 2022; Almazrouei et al., 2023; Zhang et al., 2024). Therefore, it is possible for them to reach a consensus on relative mastery over different knowledge and the necessity of retrieval. Our preliminary quantitative analysis, shown in Section 4.6, also supports this hypothesis. The experimental results show that on questions well understood by the LLM, the relatively smaller language model also has considerable knowledge. The gap between larger and smaller LLMs mainly manifests in questions they do not understand. This further validates the possibility of employing a proxy model to help determine the necessity of retrieval.

Based on our analysis, in this paper, we introduce a novel approach, called **SlimPLM** (**Slim Proxy Language Model**), which leverages a relatively smaller language model as a “proxy model” to help determine when and how to perform retrieval for LLMs. Specifically, for a user question, SlimPLM first uses the proxy model to generate a preliminary “heuristic answer”. This heuristic answer serves two purposes. First, it is evaluated by a lightweight model designed to assess the necessity for retrieval. If this evaluation shows that the heuristic answer is of high quality, it implies that the question may be addressed directly by LLMs without additional information retrieval. In contrast, a lower-quality answer triggers the retrieval process to identify and supplement missing knowledge. To facilitate this, SlimPLM utilizes the heuristic answer again to generate multiple queries, each reflecting a specific aspect of the initial response. These queries are then individually assessed for their need for retrieval, filtering out queries that do not require retrieval. By this means, the remaining queries can retrieve more relevant knowledge that is lacking in LLMs. The integration of SlimPLM into existing RAG frameworks offers a flexible and effective enhancement without notably increasing computational costs or response latency. Experimental results across five commonly used question-answering datasets validate SlimPLM’s effectiveness in determining the necessity for retrieval and improving retrieval results.

Our contributions are threefold: (1) We propose a novel approach that leverages a small proxy model to generate heuristic answers, helping deter-

mine when and how to perform retrieval for LLMs. (2) We devise a retrieval necessity judgment model based on the heuristic answer. It is capable of accurately identifying which queries necessitate further information retrieval. (3) We formulate a query rewriting strategy that decomposes the heuristic answer into distinct claims. This is complemented by a claim-based filtering mechanism to enhance the relevance of the retrieval results for LLMs’ text generation.

2 Related Work

2.1 Retrieval-augmented Generation (RAG)

Before the rise of LLMs, small language models (Raffel et al., 2019; Lewis et al., 2019) possess limited knowledge, so retrieval-augmentation is used to assist knowledge-incentive QA tasks. Previous works often design a fixed module for retrieval-augmentation (Sun et al., 2022; Lewis et al., 2020; Izacard and Grave, 2020). Since the release of ChatGPT, LLMs have been developing rapidly (Zhao et al., 2023). LLMs possess abundant intrinsic knowledge, but hallucination issue means they also have to rely on retrieval-augmentation to improve factuality (Shuster et al., 2021; White, 2023).

Since LLMs have a rather fixed architecture, retrieval-augmented LLMs generally leave LLMs’ architecture unchanged while pre-pending grounding documents to the input (Ram et al., 2023; Yu et al., 2023b). Generally, a typical RAG pipeline (Zhu et al., 2023; Liu et al., 2023a; Shi et al., 2023b) consists of a query rewriter (Wang et al., 2023a; Gao et al., 2022), a retriever (Guu et al., 2020; Neelakantan et al., 2022), a filter or reranker (Yoran et al., 2023; Yu et al., 2023a; Xu et al., 2023), and an LLM as reader. Some methods attempt to enhance preposition components for overall performance improvement. Others fine-tune the LLM to fit into RAG tasks (Asai et al., 2023; Kadavath et al., 2022).

2.2 Retrieval Necessity Judgment

Except for modifications to RAG pipeline components, retrieval necessity judgment is a more closely related topic. There are several streams to solve this issue.

(1) Fine-tuning LLMs is effective, but incurs significant computational costs (Qin et al., 2023a; Lin et al., 2022). Some work fine-tune the LLM to imitate a human-like web-browsing behavior (Schick

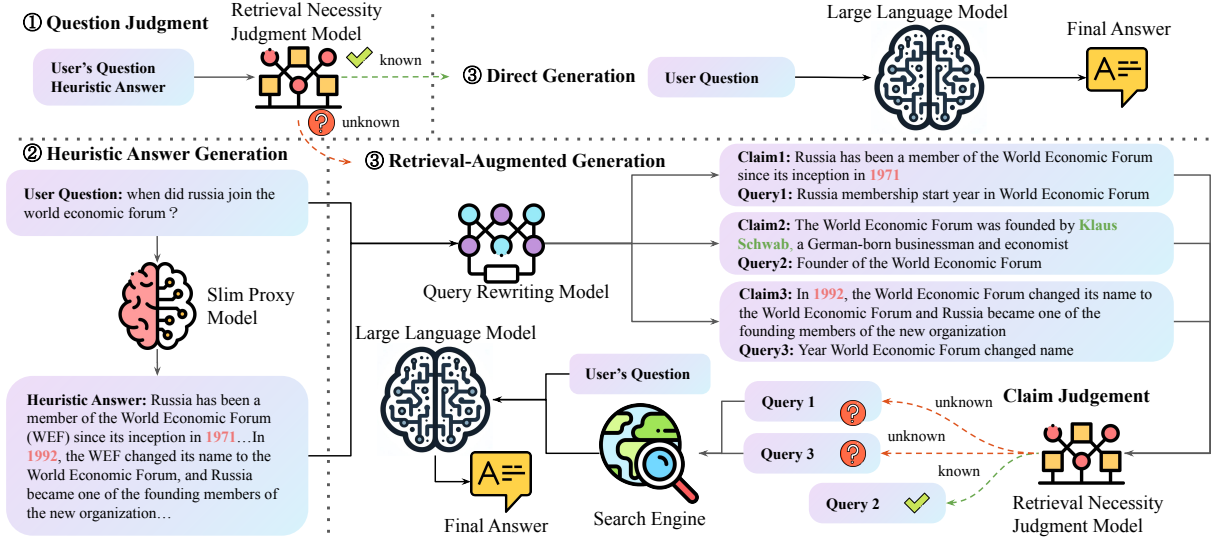


Figure 1: A display of the main process of SlimPLM. Solid lines with arrows represent the flow of data, while dashed lines with arrows signify control signals from the retrieval necessity judgment model. Step 1 and step 2 are mandatory in the pipeline, but step 3 involves choosing between direct generation and RAG.

et al., 2023; Nakano et al., 2021). Self-RAG (Asai et al., 2023) introduces special tokens called reflection tokens to control retrieval behavior.

(2) Another intuitive approach involves assessing the LLM’s confidence based on logits produced by the model (Jiang et al., 2020; Guo et al., 2017). FLARE (Jiang et al., 2023) dynamically triggers RAG if logits fall below a certain threshold.

(3) Other work use iterative prompts to decide if extra information is required (Wei et al., 2022; Liu et al., 2021; Rubin et al., 2021), or synergy Chain-of-Thought prompting (Wei et al., 2022) with RAG (Press et al., 2022; Khattab et al., 2022). For example, ReAct (Yao et al., 2022) alternates between generating thoughts and actions, creating a series of thought-action-observation steps.

(4) Assessing the difficulty or popularity of user questions to judge retrieval necessity is also a viable approach (Mallen et al., 2022). SKR (Wang et al., 2023b) refers to similar questions they have previously encountered to judge retrieval necessity.

By contrast, SlimPLM conducts retrieval necessity judgment by assessing the answer generated by a smaller LLM. SlimPLM does not increase LLM inference times while improving judgment accuracy.

2.3 Query Formulation

Many methods conduct retrieval necessity judgment and also involve query formulation. FLARE (Jiang et al., 2023) masks the low logits tokens in the LLM output sequence and concate-

nates the rest to formulate the query. ReAct (Yao et al., 2022) and Self-Ask (Press et al., 2022) prompt the LLM to generate queries itself. By contrast, SlimPLM conducts query formulation by scrutinizing the answer generated by a smaller LLM providing an accurate perception of the necessary knowledge.

3 Methodology

In this paper, we aim to leverage a relatively smaller model as the *proxy model* to determine whether the user-issued question requires supplementary retrieval results and further provide clues for retrieving relevant knowledge. Our method, SlimPLM, can be flexibly used as a plug-in to various retrieval-augmented generation scenarios, without additional training requirements. The illustration of our method is shown in Figure 1.

3.1 Problem Formulation & Framework

Before diving into the details of our method, we first formulate the concept and notations involved in this paper.

Given a user input x and a text corpus (e.g., a Wikipedia dump) $D = \{d_i\}_{i=1}^N$ of size N , models are expected to generate the annotated answer y . To obtain the information in D that is relevant to x , a retriever (R) is employed. This retriever takes a query q as input and returns a relevant text list $D_{\text{ref}} = R(q)$. Typically, the user input x is used as the query, namely $q = x$. However, existing studies have demonstrated that using refined queries

for retrieval can improve the final generation quality (Gao et al., 2022; Wang et al., 2023a). Therefore, we denote the refined queries as $\{q_1, \dots, q_n\}$. With these refined queries, a collection of relevant retrieval results $D_{\text{ref}} = R(q_1), \dots, R(q_n)$ is assembled to support the generation process, formalized as $\hat{y} = \text{LLM}(D_{\text{ref}}, x)$. Note that, when $D_{\text{ref}} = \emptyset$, the process degenerates to normal generation without retrieval.

We define a proxy model (PM), which is implemented by a relatively smaller LLM. The proxy model generates an answer for the input x as:

$$\hat{a} = \text{PM}(x), \quad (1)$$

where \hat{a} is called a *heuristic answer* in this paper. This heuristic answer serves two purposes: (1) It is used for determining whether the retrieval is necessary for the current input x . The determination is made by a retrieval necessity judgment model (introduced in Section 3.2). (2) It also provides clues for query rewriting. The query rewriting results will help identify knowledge gaps within the LLM that necessitate further retrieval (introduced in Section 3.3).

3.2 Retrieval Necessity Judgment

Because existing LLMs are typically trained on common corpora (such as CommonCrawl and Wikipedia (Touvron et al., 2023; Penedo et al., 2023; Gao et al., 2020)) and employ a similar Transformer decoder-based architecture, it is promising to leverage a smaller LLM for judging the knowledge mastered by larger LLMs and determining the need for additional retrieval. Thus, we propose a retrieval necessity judgment component.

Judgment Model Given the heuristic answer \hat{a} generated by the proxy model (Equation 1), we fine-tune a judgment model RJ (implemented by Llama2-7B in our experiments) by using both the user input x and the heuristic answer \hat{a} . We use the following instructions for fine-tuning:

```
Input:
<SYS> You are a helpful assistant. Your
task is to parse user input into structured
formats and accomplish the task according to
the heuristic answer. </SYS>
Heuristic answer: {Heuristic Answer}
Question: {user question}
Retrieval Necessity Judgment Output:
Output:
Known (True / False)
```

After fine-tuning, the RJ model can predict whether a user question needs further retrieval with the help of the heuristic answer.

Judgment Label Collection To fine-tune the RJ model, we need to collect training samples with reliable labels. Existing studies (Wang et al., 2023b) have proposed an annotation strategy that compares the models’ outputs generated with and without retrieval. In our preliminary study, we find that this strategy is highly influenced by the capability of the retriever and the completeness of the corpus, leading to annotations that cannot accurately reflect the model’s necessity for search. To tackle this problem, we propose to leverage the quality of our heuristic answers, *i.e.*, if the quality of the heuristic answer is higher than a predefined threshold, we infer that the question can be well answered without retrieval; otherwise, we consider retrieval necessary.

Specifically, we collect samples with short answers from existing question-answering datasets and employ the matching ratio between the heuristic answers and the ground-truth answers as the metric. Compared to rouge scores (Lin, 2004) or perplexity (Huyen, 2019), this metric can better align with the evaluation and reflect the generation quality. Notably, while we only use short answers for label collection, the obtained model can well generalize to different datasets, such as long-form QA datasets. Formally, for a question with multiple short answers $Y = \{y_1, y_2, \dots, y_n\}$, we compute the matching ratio r between \hat{a} and Y as:

$$r = \frac{|\{y \mid y \in \hat{a} \wedge y \in Y\}|}{|Y|}. \quad (2)$$

Then, we set a threshold θ and obtain the label as:

$$\text{Label}(\hat{a}, x) = \begin{cases} \text{Known (True)}, & \text{if } r > \theta; \\ \text{Known (False)}, & \text{otherwise.} \end{cases}$$

3.3 Retrieval Target Determination

After determining the necessity of retrieval, the next question is how to perform effective retrieval. A straightforward method is using user input x as the query to retrieve relevant information from the corpora D . However, many studies have reported that the information retrieved by x may lose details and introduce redundant content (Wang et al., 2023a). To address this issue, we propose a query rewrite method based on the heuristic answers and a query filter method to refine these rewritten queries.

Heuristic Answer-Driven Query Rewrite Restricted by parameter scale, the proxy model often hallucinates during the process of answering questions, but the direction in which they answer questions is heuristic (Dhuliawala et al., 2023; Gao et al., 2022). They can extend related aspects and sub-topics of thought when analyzing questions. Inspired by claim decomposition operation intended for factual evaluation (Min et al., 2023; Kryscinski et al., 2019), we perform query rewriting based on each fact mentioned in the heuristic answer given by the proxy models. The specific operations are as follows: we decompose the heuristic answer \hat{a} into multiple claims related to the question, $\{c_1, c_2, \dots, c_n\}$, where each claim related to the question can lead to a query, $\{q_{c_1}, q_{c_2}, \dots, q_{c_n}\}$. In addition, we combine the query rewrites directly derived from the user’s input $\{q_{x_1}, q_{x_2}, \dots, q_{x_n}\}$. Our query rewriting model QR takes the user question and the heuristic answer as input and outputs all query rewrite results, $QR(x, \hat{a}) = \{q_{x_1}, \dots, q_{x_n}, q_{c_1}, \dots, q_{c_n}\}$.

To train the query rewriting model, we collect and annotate a dataset with the help of GPT-4 (Achiam et al., 2023). In each dataset used in our experiments, we sample 1,000 user questions. We utilize the method of instruction fine-tuning (Ouyang et al., 2022; Chung et al., 2022) to fine-tune a decoder-only generative model, accomplishing the task of claim extraction and query rewriting in a single round. Our instructions and the model output are displayed as follows.

```
Input:
<SYS> You are a helpful assistant. Your
task is to parse user input into structured
formats and accomplish the task according to
the heuristic answer. </SYS>
Heuristic answer: {Heuristic Answer}
Question: {User Question}
Query Rewrite Output:
Output:
<Claim> Claim 1 <Query> Query 1 <Claim> Claim
2 <Query> Query 2, ...
```

Claim-based Query Filter In the previous step, our method generates several rewritten queries $QR(x, \hat{a})$, which correspond to the claims in the heuristic answers.

To achieve this, we reuse the judgment model RJ trained in Section 3.2. Specifically, we replace the input of the heuristic answer by the extracted claim and the input of user questions by the rewritten query. Then, the model can predict whether the

rewritten query requires external knowledge from retrieval. We only perform retrieval when the result is Known (False), namely, we have:

$$D_{\text{ref}} = \{R(q_{c_i}) | RJ(c_i, q_{c_i}) = \text{Known (False)}\}.$$

By this means, we can obtain the retrieved result set D_{ref} that only contains the knowledge missing by the LLM.

4 Experiments

We conduct experiments on five widely used question-answering (QA) datasets and compare the performance of our method with several baselines.

4.1 Datasets

We use the following five QA datasets: (1) Natural Questions (NQ) (Kwiatkowski et al., 2019): a dataset consisting of real user questions from Google search. (2) Trivia-QA (Joshi et al., 2017): a realistic text-based question answering dataset. (3) ASQA (Stelmakh et al., 2022): a dataset targeting ambiguous questions requiring answers that integrate factual information from various sources. (4) MuSiQue (Trivedi et al., 2021): a synthetic multi-hop question-answering dataset. (5) ELI5 (Fan et al., 2019): a long-form question answering dataset originated from the Reddit forum. Due to our limited resources, we randomly sample 400 questions from the test set (if any) or validation set of each dataset as the test set for evaluation.

4.2 Evaluation Metrics

For all QA tasks, LLMs can freely generate any answers. For datasets annotated with long-form answers, we employ the Rouge Score (Lin, 2004) (ROUGE) to evaluate the quality of the generated answers by comparing them with the ground-truth ones. For datasets with short answers, we use the Exact Match (EM) metric to compare the generated answer with the golden one. If the dataset provides multiple optional short answers, we also report the proportion of instances where at least one short answer matches (Hit@1).

4.3 Baselines

We first select two baselines without retrieval:

(1) **Vanilla Chat**: This method directly inputs the user question into LLMs to get the answer. (2) **CoT Prompting** (Wei et al., 2022): This method introduces a prompt method that lets LLMs think step-by-step to derive the final answer.

Method	#API	ASQA		NQ		Trivia-QA		MuSiQue	ELI5		
		EM	Hit@1	EM	Hit@1	EM	Hit@1	EM	ROUGE-1	ROUGE-2	ROUGE-L
Llama2-70B-Chat without Retrieval											
Vanilla	1	29.68	62.50	40.49	55.00	27.44	90.75	11.50	28.66	4.88	14.27
CoT	1	26.21	54.50	35.36	48.75	23.50	79.00	11.50	28.12	4.73	14.06
Llama2-70B-Chat with Retrieval											
Direct	1	27.63	58.00	42.40	56.00	28.07	92.25	10.50	28.61	4.76	15.76
FLARE	2.10	30.08	63.50	41.36	55.75	27.41	89.50	11.25	27.95	4.72	13.91
Self-Eval	2	29.45	60.75	42.15	55.75	27.58	91.50	10.25	28.70	4.83	15.39
Self-Ask	2.67	26.37	60.25	38.56	53.00	26.56	89.50	9.50	-	-	-
ITER-RETGEN	3	30.15	60.50	42.85	55.50	28.31	91.00	13.00	28.44	4.74	15.72
SKR-KNN	1	29.38	61.75	41.90	55.75	28.16	92.25	10.25	28.71	4.80	15.73
SlimPLM (Ours)	1	30.73	65.00	47.43	62.25	28.35	92.00	13.00	29.97	5.61	15.13
Qwen-72B-Chat without Retrieval											
Vanilla	1	26.65	58.50	40.38	53.75	27.82	90.25	11.75	30.61	5.21	15.90
CoT	1	27.74	59.50	40.49	53.75	27.62	91.75	12.75	29.94	4.94	14.75
Qwen-72B-Chat with Retrieval											
Direct	1	25.85	57.00	41.27	52.75	26.39	87.75	7.75	25.93	4.55	16.74
FLARE	2.29	27.68	59.00	40.89	54.50	27.10	88.50	12.75	30.31	5.20	15.77
Self-Eval	2	27.64	60.00	42.43	56.00	27.13	90.50	7.75	29.19	5.14	16.05
Self-Ask	2.76	22.82	52.25	36.16	49.25	25.29	87.50	9.75	-	-	-
ITER-RETGEN	3	29.38	61.50	43.51	57.50	27.16	89.75	12.25	26.15	4.41	16.52
SKR-KNN	1	28.08	61.50	43.08	56.00	26.38	88.50	11.25	27.29	4.75	16.31
SlimPLM (Ours)	1	27.97	62.25	44.07	57.75	28.03	92.25	9.75	29.56	5.91	16.36

Table 1: Evaluation results of SlimPLM and baselines on five QA benchmarks. #API is the average LLM inference times. Hit@1 is the proportion of instances where at least one short answer matches.

We also consider several retrieval-augmented generation methods. They differ in time and approach for retrieval necessity judgment and construction of retrieval queries. We include more baseline implementation details in Appendix B.

(1) **Direct Search**: This approach uses search for all questions and directly utilizes the user question as the search query. (2) **FLARE** (Jiang et al., 2023): This method examines the content of each sentence generated by the LLM, and uses retrieval if the generation logits are below a threshold. FLARE uses the masked sentence as a query, wherein tokens associated with low logits are masked. (3) **Self-Eval** (Kadavath et al., 2022): This method uses prompts and few-shot learning to let LLM itself decide whether it needs retrieval or not. (4) **Self-Ask** (Press et al., 2022): This method iteratively prompts the LLM to decide whether to generate follow-up questions as queries or generate the final answer directly. (5) **SKR-KNN** (Wang et al., 2023b). It uses a dense retriever to retrieve top- k nearest neighbor questions from the training set. The necessity of retrieval is determined by the number of neighboring questions that require or do not require retrieval.

4.4 Implementation Details

We conduct experiments on two open-source LLMs, Llama2-70B-Chat (Touvron et al., 2023) and Qwen-72b-Chat (Bai et al., 2023). The default proxy model, fine-tuned query rewriting model, and retrieval necessity judgment model are built on Llama2-7B-Chat. We build a search engine on the KILT dataset’s document library, which is based on the 2019 Wikipedia mirror (Petroni et al., 2020b). BM25 (Robertson and Zaragoza, 2009) is used as the retriever and E5_{base} (Wang et al., 2022) is employed as the reranker. More implementation details are provided in Appendix A.

4.5 Experimental Results

The evaluation results are shown in Table 1, where we uniformly chose Llama2-7B-Chat as the proxy model, a fine-tuned query rewriting model, and a fine-tuned retrieval necessity judgment model. Generally, our SlimPLM achieves superior or competitive performance on all datasets. This clearly demonstrates the effectiveness of our method. Besides, we have the following observations:

(1) On most datasets, retrieval-augmented generation methods can outperform the methods with-




Figure 2: The rouge performance for with and without retrieval on ELI5 dataset

out using retrieval. This clearly demonstrates the benefit of incorporating external knowledge into open-domain QA tasks. (2) Compared to methods that initiate retrieval based on the results or logits generated by LLMs (*i.e.*, Self-Eval, Self-Ask, and FALRE), our method yields better results. This validates the superiority of our method, which employs a proxy model to determine when and what the LLM needs to retrieve. Notably, our method requires the LLM to infer only once, significantly reducing the cost of inference. (3) Comparing methods that judge retrieval necessity merely based on user questions (SKR-KNN), our method also has advantages. By using heuristic answers, it can more accurately assess the LLM’s knowledge capability and formulate queries that are more precisely tailored to the question, thereby improving overall performance. (4) Intriguingly, we notice that retrieval does not uniformly benefit all user questions. For example, in the ELI5 dataset, approximately 66.4% of samples show improvement with retrieval, as shown in Figure 2. This observation highlights the critical need to judge the necessity of retrieval. More cases where retrieval has negative impact are shown in Appendix C.

4.6 Further Analysis

We further conduct a series of experiments to investigate the impact of different settings in our method.

Ablation Study We first examine the effectiveness of different modules in our method by an ablation study. This experiment is conducted by removing the heuristic-answer-driven query rewriting (*w/o* QR), question-Level retrieval necessity judgment (*w/o* RJ), and Claim-based Query Filter (*w/o* QF), respectively. From the results are shown in Table 2, we can see: (1) If query rewriting is removed,

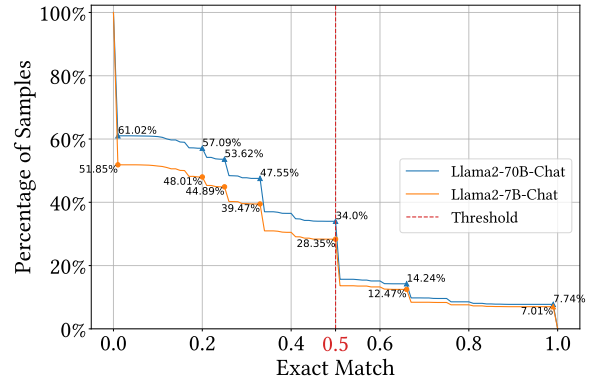


Figure 3: The proportion of samples (*y*-axis) with EM scores higher than certain values (*x*-axis).

then retrieval necessity judgment between vanilla chat and direct search is applied. Performing query rewriting can both enhance the comprehensiveness and relevance of retrieved references. (2) When retrieval necessity judgment is removed, all questions will use retrieval results for generation. LLMs will be led astray on questions that they can perform well on their own knowledge. (3) If claim-based query filter is removed, then retrieval is applied to every query derived from the heuristic answer. Not filtering queries which contain contents that do not require retrieval worsens the search results.

Knowledge Ability Consensus between Proxy Models and LLMs

In this experiment, we compared the knowledge capabilities of LLMs and proxy models, and confirmed their consensus. Our findings can be summarized as follows: (1) The difference in capabilities between the proxy model and the LLM is primarily manifested in the knowledge of lower mastery levels. As illustrated in Figure 3, on the ASQA dataset, the difference between the 70B and 7B language models is very slight for samples with an EM score greater than 0.5. Their differences are primarily evident in samples with an EM score less than 0.5. (2) The higher the level of some knowledge mastered by the proxy model, the higher the level of mastery by the LLM. Further experiments on ASQA shows over 82.19% of the samples with an EM score greater than 0.5 for the 7B model overlaps with those of the 70B model.

The experimental results above offer a theoretical basis for our method. If the proxy model can correctly answer the question, then the LLM is very likely to answer it correctly as well. Applying vanilla chat for them can better leverage the inherent knowledge capabilities of LLMs.

Method	ASQA		NQ		Trivia-QA		MuSiQue
	EM	Hit@1	EM	Hit@1	EM	Hit@1	EM
SlimPLM	30.73	65.00	47.43	62.25	28.35	92.00	13.00
w/o QR	29.19 (5.0%↓)	61.75 (5.0%↓)	45.16 (4.8%↓)	59.75 (4.0%↓)	28.07 (1.0%↓)	92.50 (0.5%↑)	11.50 (1.2%↓)
w/o QJ	29.43 (4.2%↓)	61.75 (5.0%↓)	43.03 (9.3%↓)	57.25 (8.0%↓)	27.91 (1.6%↓)	90.25 (1.9%↓)	12.75 (1.9%↓)
w/o QF	30.73 (0.0%)	64.75 (0.4%↓)	46.62 (1.7%↓)	61.25 (1.6%↓)	28.27 (0.3%↓)	91.75 (0.3%↓)	12.50 (3.9%↓)

Table 2: Ablation study on Llama2-70B-Chat. “QR”, “QJ”, and “QF” denote the query rewriting, question-level retrieval necessity judgment, and claim-based query filter, respectively.

Method	ASQA		NQ		Trivia-QA		MuSiQue	ELI5		
	EM	Hit@1	EM	Hit@1	EM	Hit@1	EM	ROUGE-1	ROUGE-2	ROUGE-L
Llama2-70B-Chat										
Llama2-7B-Chat	30.73	65.00	47.43	62.25	28.35	92.75	13.00	29.97	5.61	15.13
Qwen-7B-Chat	29.62	60.25	42.53	56.25	27.93	92.25	13.00	29.95	5.57	16.16
TinyLlama-1.1B-Chat	30.47	60.50	44.24	56.75	28.02	91.00	11.50	30.05	5.56	15.37
Phi-2 (2.7B)	28.96	60.50	43.33	57.50	27.99	91.50	13.75	30.34	5.82	15.48
Qwen-72B-Chat										
Llama2-7B-Chat	27.97	62.25	44.07	57.75	28.03	92.25	9.75	29.56	5.91	16.36
Qwen-7B-Chat	27.76	59.75	42.54	55.75	27.22	90.25	8.75	29.44	5.74	16.33
TinyLlama-1.1B-Chat	27.61	58.25	42.36	55.25	27.67	91.25	9.25	28.80	5.64	16.12
Phi-2 (2.7B)	26.95	59.50	42.22	54.25	27.10	89.00	10.75	29.17	5.83	16.32

Table 3: Performance Comparison of Various Proxy Methods.

Dataset	Chat	Proxy	Rewrite	Judge	Total
ASQA	192.86	24.42	35.27	3.38	63.07
NQ	249.74	29.61	38.65	3.65	71.91
TQA	114.07	15.73	28.13	2.47	46.33
MuSiQ	168.47	19.00	30.84	2.36	52.20
ELI5	471.22	47.82	46.82	4.23	98.87

Table 4: The number of tokens used by LLM, and the additional tokens brought by SlimPLM.

Impact of Various Proxy Models We also explore the impact of using different proxy models in our method. This experiment is conducted by using four open-source LLMs with different sizes as the proxy model, including Llama2-7B-Chat (Touvron et al., 2023), Qwen-7B-Chat (Bai et al., 2023), TinyLlama-1.1B-Chat (Zhang et al., 2024), and Phi-2 (Li et al., 2023). Experimental results are shown in Table 3. We can see that in most datasets, Llama2-7B can provide the best results. Furthermore, Llama2-7B contributes a greater improvement to Llama2-70B than to Qwen-72B, we attribute this to the better knowledge alignment between Llama models.

Computational Cost Analysis In our method, we use a proxy model, a query rewriting model, and a retrieval necessity judgment model based on relatively smaller LLMs (7B). To investigate their

computational efficiency, we analyze the average number of tokens generated by each model and calculate the associated costs. This calculation is based on the assumption that the computational expense per token for a 7B model is roughly 1/10 that of a 70B model—a conservative estimate, given that the actual cost differential is likely to exceed this ratio (Kaplan et al., 2020). Table 4 lists the additional computational costs required by each component and the total cost. The analysis reveals that the additional costs are substantially lower (1/4 to 1/3) compared to the costs of a single inference by an LLM. This observation validates the economic advantages of our method.

5 Conclusion

In conclusion, our research proposes a new paradigm for RAG, utilizing a smaller LLM as proxy model. Based on the heuristic answer by proxy model, we conduct query rewriting, retrieval necessity judgment, and claim-based query filtering. This approach enables accurate perception for when and what to retrieve for LLMs. Experiments across various datasets show a marked improvement in the end-to-end performance of LLM question-answering, achieving or exceeding state-of-the-art results. Moreover, this enhancement is attained with little additional computational cost.

Limitations

In scenarios where almost all user questions are primarily outside the scope of the LLM’s pre-training corpus, or where almost all the questions do not require external knowledge, our method proves challenging to utilize. In these situations, opting either for a full retrieval or without retrieval at all may be a more suitable approach. Additionally, we acknowledge a gap in the knowledge capabilities between proxy models and LLMs. Heuristic answers are unable to fully reflect the true knowledge capability of the LLMs. Moreover, our current method employs three models: a proxy model, a query rewriting model, and a retrieval necessity judgment model. The pipeline appears somewhat complex; integrating these functions into a single generative framework would be preferable.

References

OpenAI Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madeleine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Benjamin Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Sim'on Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Raphael Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Lukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Hendrik Kirchner, Jamie Ryan Kiros, Matthew Knight, Daniel Kokotajlo, Lukasz Kondraciuk, Andrew Kondrich, Aris Konstantini-

dis, Kyle Kopic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Adeola Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel P. Mossing, Tong Mu, Mira Murati, Oleg Murk, David M'ely, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Ouyang Long, Cullen O'Keefe, Jakub W. Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alexandre Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Pondé de Oliveira Pinto, Michael Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario D. Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin D. Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas A. Tezak, Madeleine Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cer'on Uribe, Andrea Valone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2023. [Gpt-4 technical report](#).

Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra-Aimée Cojocaru, Daniel Hesslow, Julien Launay, Quentin Malartic, Daniele Mazzotta, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. [The falcon series of open language models](#). *ArXiv*, abs/2311.16867.

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). *ArXiv*, abs/2310.11511.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang,

699	Xiaodong Deng, Yang Fan, Wenhong Ge, Yu Han,	Gautier Izacard and Edouard Grave. 2020. Leveraging	755
700	Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang	passage retrieval with generative models for open	756
701	Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang	domain question answering . <i>ArXiv</i> , abs/2007.01282.	757
702	Lu, K. Lu, Jianxin Ma, Rui Men, Xingzhang Ren,		
703	Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong	Zhengbao Jiang, J. Araki, Haibo Ding, and Graham	758
704	Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang	Neubig. 2020. How can we know when language	759
705	Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian	models know? on the calibration of language	760
706	Yang, Jian Yang, Shusheng Yang, Shusheng Yang,	for question answering . <i>Transactions of the Association</i>	761
707	Bowen Yu, Yu Bowen, Hongyi Yuan, Zheng Yuan,	<i>for Computational Linguistics</i> , 9:962–977.	762
708	Jianwei Zhang, Xing Zhang, Yichang Zhang, Zhenru		
709	Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou,	Zhengbao Jiang, Frank F. Xu, Luyu Gao, Zhiqing Sun,	763
710	and Tianhang Zhu. 2023. Qwen technical report .	Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie	764
711	<i>ArXiv</i> , abs/2309.16609.	Callan, and Graham Neubig. 2023. Active retrieval	765
		augmented generation . <i>ArXiv</i> , abs/2305.06983.	766
712	Hyung Won Chung, Le Hou, S. Longpre, Barret Zoph,		
713	Yi Tay, William Fedus, Eric Li, Xuezhi Wang,	Mandar Joshi, Eunsol Choi, Daniel S. Weld, and Luke	767
714	Mostafa Dehghani, Siddhartha Brahma, Albert Web-	Zettlemoyer. 2017. Triviaqa: A large scale distantly	768
715	son, Shixiang Shane Gu, Zhuyun Dai, Mirac Suz-	supervised challenge dataset for reading comprehen-	769
716	gun, Xinyun Chen, Aakanksha Chowdhery, Dasha	sion . <i>ArXiv</i> , abs/1705.03551.	770
717	Valter, Sharan Narang, Gaurav Mishra, Adams Wei		
718	Yu, Vincent Zhao, Yanping Huang, Andrew M.	Saurav Kadavath, Tom Conerly, Amanda Askell, T. J.	771
719	Dai, Hongkun Yu, Slav Petrov, Ed Huai hsin Chi,	Henighan, Dawn Drain, Ethan Perez, Nicholas	772
720	Jeff Dean, Jacob Devlin, Adam Roberts, Denny	Schiefer, Zachary Dodds, Nova DasSarma, Eli Tran-	773
721	Zhou, Quoc V. Le, and Jason Wei. 2022. Scal-	Johnson, Scott Johnston, Sheer El-Showk, Andy	774
722	ing instruction-finetuned language models . <i>ArXiv</i> ,	Jones, Nelson Elhage, Tristan Hume, Anna Chen,	775
723	abs/2210.11416.	Yuntao Bai, Sam Bowman, Stanislav Fort, Deep	776
		Ganguli, Danny Hernandez, Josh Jacobson, John	777
724	Shehzaad Dhuliawala, Mojtaba Komeili, Jing Xu,	Kernion, Shauna Kravec, Liane Lovitt, Kamal	778
725	Roberta Raileanu, Xian Li, Asli Celikyilmaz, and	Ndousse, Catherine Olsson, Sam Ringer, Dario	779
726	Jason Weston. 2023. Chain-of-verification reduces	Amodei, Tom B. Brown, Jack Clark, Nicholas Joseph,	780
727	hallucination in large language models . <i>ArXiv</i> ,	Benjamin Mann, Sam McCandlish, Christopher Olah,	781
728	abs/2309.11495.	and Jared Kaplan. 2022. Language models (mostly)	782
		know what they know . <i>ArXiv</i> , abs/2207.05221.	783
729	Angela Fan, Yacine Jernite, Ethan Perez, David		
730	Grangier, Jason Weston, and Michael Auli. 2019.	Jared Kaplan, Sam McCandlish, T. J. Henighan, Tom B.	784
731	Eli5: Long form question answering . <i>ArXiv</i> ,	Brown, Benjamin Chess, Rewon Child, Scott Gray,	785
732	abs/1907.09190.	Alec Radford, Jeff Wu, and Dario Amodei. 2020.	786
		Scaling laws for neural language models . <i>ArXiv</i> ,	787
733	Leo Gao, Stella Biderman, Sid Black, Laurence Gold-	abs/2001.08361.	788
734	ing, Travis Hoppe, Charles Foster, Jason Phang,		
735	Horace He, Anish Thite, Noa Nabeshima, Shawn	O. Khattab, Keshav Santhanam, Xiang Lisa Li, David	789
736	Presser, and Connor Leahy. 2020. The pile: An	Leo Wright Hall, Percy Liang, Christopher Potts,	790
737	800gb dataset of diverse text for language modeling .	and Matei A. Zaharia. 2022. Demonstrate-search-	791
738	<i>ArXiv</i> , abs/2101.00027.	predict: Composing retrieval and language models	792
		for knowledge-intensive nlp . <i>ArXiv</i> , abs/2212.14024.	793
739	Luyu Gao, Xueguang Ma, Jimmy J. Lin, and Jamie		
740	Callan. 2022. Precise zero-shot dense retrieval with-	Suhas Kotha, Jacob M. Springer, and Aditi Raghu-	794
741	out relevance labels . <i>ArXiv</i> , abs/2212.10496.	nathan. 2023. Understanding catastrophic forgetting	795
		in language models via implicit inference . <i>ArXiv</i> ,	796
742	Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Wein-	abs/2309.10105.	797
743	berger. 2017. On calibration of modern neural net-		
744	works . In <i>International Conference on Machine</i>	Wojciech Kryscinski, Bryan McCann, Caiming Xiong,	798
745	<i>Learning</i> .	and Richard Socher. 2019. Evaluating the factual	799
		consistency of abstractive text summarization . In	800
746	Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasu-	<i>Conference on Empirical Methods in Natural Lan-</i>	801
747	pat, and Ming-Wei Chang. 2020. Realm: Retrieval-	<i>guage Processing</i> .	802
748	augmented language model pre-training . <i>ArXiv</i> ,		
749	abs/2002.08909.	Tom Kwiatkowski, Jennimaria Palomaki, Olivia Red-	803
		field, Michael Collins, Ankur P. Parikh, Chris Alberti,	804
750	Hangfeng He, Hongming Zhang, and Dan Roth. 2022.	Danielle Epstein, Illia Polosukhin, Jacob Devlin, Ken-	805
751	Rethinking with retrieval: Faithful large language	ton Lee, Kristina Toutanova, Llion Jones, Matthew	806
752	model inference . <i>ArXiv</i> , abs/2301.00303.	Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob	807
		Uszkoreit, Quoc V. Le, and Slav Petrov. 2019. Natu-	808
753	Chip Huyen. 2019. Evaluation metrics for language	ral questions: A benchmark for question answering	809
754	modeling. <i>The Gradient</i> .	research . <i>Transactions of the Association for Compu-</i>	810
		<i>tational Linguistics</i> , 7:453–466.	811

812	Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdel rahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	867
813		868
814		869
815		870
816		871
817		
818		
819	Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks . <i>ArXiv</i> , abs/2005.11401.	872
820		873
821		874
822		875
823		876
824		877
		878
		879
825	Yuan-Fang Li, Sébastien Bubeck, Ronen Eldan, Allison Del Giorno, Suriya Gunasekar, and Yin Tat Lee. 2023. Textbooks are all you need ii: phi-1.5 technical report . <i>ArXiv</i> , abs/2309.05463.	880
826		881
827		882
828		883
829	Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	884
830		885
831		886
832	Stephanie C. Lin, Jacob Hilton, and Owain Evans. 2022. Teaching models to express their uncertainty in words . <i>Trans. Mach. Learn. Res.</i> , 2022.	887
833		888
834		889
835	Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2021. What makes good in-context examples for gpt-3? In <i>Workshop on Knowledge Extraction and Integration for Deep Learning Architectures; Deep Learning Inside Out</i> .	890
836		891
837		892
838		893
839		894
840		895
841	Jiongnan Liu, Jiajie Jin, Zihan Wang, Jiehan Cheng, Zhicheng Dou, and Ji rong Wen. 2023a. Reta-llm: A retrieval-augmented large language model toolkit . <i>ArXiv</i> , abs/2306.05212.	896
842		897
843		898
844		
845	Xiao Liu, Hanyu Lai, Hao Yu, Yifan Xu, Aohan Zeng, Zhengxiao Du, P. Zhang, Yuxiao Dong, and Jie Tang. 2023b. Webglm: Towards an efficient web-enhanced question answering system with human preferences . <i>Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining</i> .	899
846		900
847		901
848		902
849		903
850		904
		905
851	Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Hannaneh Hajishirzi, and Daniel Khashabi. 2022. When not to trust language models: Investigating effectiveness of parametric and non-parametric memories . In <i>Annual Meeting of the Association for Computational Linguistics</i> .	906
852		907
853		908
854		909
855		910
856		911
857	Joshua Maynez, Shashi Narayan, Bernd Bohnet, and Ryan T. McDonald. 2020. On faithfulness and factuality in abstractive summarization . <i>ArXiv</i> , abs/2005.00661.	912
858		913
859		914
860		915
		916
861	Sewon Min, Kalpesh Krishna, Xinxu 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 . <i>ArXiv</i> , abs/2305.14251.	917
862		918
863		919
864		920
865		921
866		922
		923
	Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. Rethinking the role of demonstrations: What makes in-context learning work? <i>ArXiv</i> , abs/2202.12837.	
	Reiichiro Nakano, Jacob Hilton, Suchir Balaji, Jeff Wu, Ouyang Long, 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 . <i>ArXiv</i> , abs/2112.09332.	
	Arvind Neelakantan, Tao Xu, Raul Puri, Alec Radford, Jesse Michael Han, Jerry Tworek, Qiming Yuan, Nikolas A. Tezak, Jong Wook Kim, Chris Hallacy, Johannes Heidecke, Pranav Shyam, Boris Power, Tyna Eloundou, Niekou, Girish Sastry, Gretchen Krueger, David P. Schnurr, Felipe Petroski Such, Kenny Sai-Kin Hsu, Madeleine Thompson, Tabarak Khan, Toki Sherbakov, Joanne Jang, Peter Welinder, and Lilian Weng. 2022. Text and code embeddings by contrastive pre-training . <i>ArXiv</i> , abs/2201.10005.	
	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 . <i>ArXiv</i> , abs/2203.02155.	
	Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra-Aimée Cojocaru, Alessandro Cappelli, Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. 2023. The refined-web dataset for falcon llm: Outperforming curated corpora with web data, and web data only . <i>ArXiv</i> , abs/2306.01116.	
	Baolin Peng, Michel Galley, Pengcheng He, Hao Cheng, Yujia Xie, Yu Hu, Qiuyuan Huang, Lars Lidén, Zhou Yu, Weizhu Chen, and Jianfeng Gao. 2023. Check your facts and try again: Improving large language models with external knowledge and automated feedback . <i>ArXiv</i> , abs/2302.12813.	
	Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020a. How context affects language models' factual predictions . <i>ArXiv</i> , abs/2005.04611.	
	Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vassilis Plachouras, Tim Rocktaschel, and Sebastian Riedel. 2020b. Kilt: a benchmark for knowledge intensive language tasks . In <i>North American Chapter of the Association for Computational Linguistics</i> .	

924	Ofir Press, Muru Zhang, Sewon Min, Ludwig Schmidt,	Natan, Francesco De Toni, Gérard Dupont, Germán	982
925	Noah A. Smith, and Mike Lewis. 2022. Measuring	Kruszewski, Giada Pistilli, Hady ElSahar, Hamza	983
926	and narrowing the compositionality gap in language	Benyamina, Hieu Trung Tran, Ian Yu, Idris Abdul-	984
927	models . <i>ArXiv</i> , abs/2210.03350.	mumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier	985
		de la Rosa, Jenny Chim, Jesse Dodge, Jian Zhu,	986
928	Yujia Qin, Zihan Cai, Di Jin, Lan Yan, Shi Liang, Kun-	Jonathan Chang, Jorg Frohberg, Josephine L. To-	987
929	lun Zhu, Yankai Lin, Xu Han, Ning Ding, Huadong	bing, Joydeep Bhattacharjee, Khalid Almubarak,	988
930	Wang, Ruobing Xie, Fanchao Qi, Zhiyuan Liu,	Kimbo Chen, Kyle Lo, Leandro von Werra, Leon	989
931	Maosong Sun, and Jie Zhou. 2023a. Webcpm: In-	Weber, Long Phan, Loubna Ben Allal, Ludovic Tan-	990
932	teractive web search for chinese long-form question	guy, Manan Dey, Manuel Romero Muñoz, Maraim	991
933	answering . In <i>Annual Meeting of the Association for</i>	Masoud, María Grandury, Mario vSavsko, Max	992
934	<i>Computational Linguistics</i> .	Huang, Maximin Coavoux, Mayank Singh, Mike	993
		Tian-Jian Jiang, Minh Chien Vu, Mohammad Ali	994
935	Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen,	Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora	995
936	Ning Ding, Ganqu Cui, Zheni Zeng, Yufei Huang,	Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar	996
937	Chaojun Xiao, Chi Han, Yi Ren Fung, Yusheng Su,	Espejel, Ona de Gibert, Paulo Villegas, Peter Hen-	997
938	Huadong Wang, Cheng Qian, Runchu Tian, Kunlun	derson, Pierre Colombo, Priscilla Amuok, Quentin	998
939	Zhu, Shi Liang, Xingyu Shen, Bokai Xu, Zhen Zhang,	Lhoest, Rheza Harliman, Rishi Bommasani, Roberto	999
940	Yining Ye, Bo Li, Ziwei Tang, Jing Yi, Yu Zhu, Zhen-	L’opez, Rui Ribeiro, Salomey Osei, Sampo Pyysalo,	1000
941	ning Dai, Lan Yan, Xin Cong, Ya-Ting Lu, Weilin	Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan	1001
942	Zhao, Yuxiang Huang, Jun-Han Yan, Xu Han, Xian	Muhammad, Shanya Sharma Sharma, S. Longpre,	1002
943	Sun, Dahai Li, Jason Phang, Cheng Yang, Tong-	So maieh Nikpoor, S. Silberberg, Suhas Pai, Syd-	1003
944	shuang Wu, Heng Ji, Zhiyuan Liu, and Maosong	ney Zink, Tiago Timponi Torrent, Timo Schick, Tris-	1004
945	Sun. 2023b. Tool learning with foundation models .	tan Thrush, Valentin Danchev, Vassilina Nikoulina,	1005
946	<i>ArXiv</i> , abs/2304.08354.	Veronika Laippala, Violette Lepercq, Vrinda Prabhu,	1006
		Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin	1007
947	Colin Raffel, Noam M. Shazeer, Adam Roberts, Kather-	Heinzerling, Chenglei Si, Elizabeth Salesky, Sab-	1008
948	ine Lee, Sharan Narang, Michael Matena, Yanqi	rina J. Mielke, Wilson Y. Lee, Abheesht Sharma, An-	1009
949	Zhou, Wei Li, and Peter J. Liu. 2019. Exploring the	ndrea Santilli, Antoine Chaffin, Arnaud Stiegler, Deba-	1010
950	limits of transfer learning with a unified text-to-text	jyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han	1011
951	transformer . <i>J. Mach. Learn. Res.</i> , 21:140:1–140:67.	Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan	1012
		Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Sai-	1013
952	Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay,	ful Bari, Maged S. Al-Shaibani, Matteo Manica, Ni-	1014
953	Amnon Shashua, Kevin Leyton-Brown, and Yoav	hal V. Nayak, Ryan Teehan, Samuel Albanie, Sheng	1015
954	Shoham. 2023. In-context retrieval-augmented lan-	Shen, Srulik Ben-David, Stephen H. Bach, Taewoon	1016
955	guage models . <i>Transactions of the Association for</i>	Kim, Tali Bers, Thibault Févry, Trishala Neeraj, Ur-	1017
956	<i>Computational Linguistics</i> , 11:1316–1331.	mish Thakker, Vikas Raunak, Xiang Tang, Zheng-	1018
		Xin Yong, Zhiqing Sun, Shaked Brody, Y Uri, Hadar	1019
957	Stephen E. Robertson and Hugo Zaragoza. 2009. The	Tojarieh, Adam Roberts, Hyung Won Chung, Jae-	1020
958	probabilistic relevance framework: Bm25 and be-	sung Tae, Jason Phang, Ofir Press, Conglong Li,	1021
959	yond . <i>Found. Trends Inf. Retr.</i> , 3:333–389.	Deepak Narayanan, Hatim Bourfoune, Jared Casper,	1022
		Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia	1023
960	Ohad Rubin, Jonathan Herzig, and Jonathan Berant.	Zhang, Mohammad Shoeybi, Myriam Peyrounette,	1024
961	2021. Learning to retrieve prompts for in-context	Nicolas Patry, Nouamane Tazi, Omar Sanseviero,	1025
962	learning . <i>ArXiv</i> , abs/2112.08633.	Patrick von Platen, Pierre Cornette, Pierre Franc-	1026
		cois Lavall’ee, Rémi Lacroix, Samyam Rajbhan-	1027
963	Teven Le Scao, Angela Fan, Christopher Akiki, El-	dari, Sanchit Gandhi, Shaden Smith, Stéphane Re-	1028
964	lie Pavlick, Suzana Ili’c, Daniel Hesslow, Ro-	quena, Suraj Patil, Tim Dettmers, Ahmed Baruwa,	1029
965	man Castagn’e, Alexandra Sasha Luccioni, Franc-	Amanpreet Singh, Anastasia Cheveleva, Anne-Laure	1030
966	cois Yvon, Matthias Gallé, Jonathan Tow, Alexan-	Ligozat, Arjun Subramonian, Aur’elie N’ev’eol,	1031
967	der M. Rush, Stella Biderman, Albert Webson,	Charles Lovering, Daniel H Garrette, Deepak R.	1032
968	Pawan Sasanka Ammanamanchi, Thomas Wang,	Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Eka-	1033
969	Benoît Sagot, Niklas Muennighoff, Albert Villanova	terina Voloshina, Eli Bogdanov, Genta Indra Winata,	1034
970	del Moral, Olatunji Ruwase, Rachel Bawden, Stas	Hailey Schoelkopf, Jan-Christoph Kalo, Jekate-	1035
971	Bekman, Angelina McMillan-Major, Iz Beltagy, Huu	rina Novikova, Jessica Zosa Forde, Xiangru Tang,	1036
972	Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz	Jungo Kasai, Ken Kawamura, Liam Hazan, Ma-	1037
973	Suarez, Victor Sanh, Hugo Laurencon, Yacine Jer-	rine Carpuat, Miruna Clinciu, Najoung Kim, New-	1038
974	nite, Julien Launay, Margaret Mitchell, Colin Raf-	ton Cheng, Oleg Serikov, Omer Antverg, Oskar	1039
975	fel, Aaron Gokaslan, Adi Simhi, Aitor Soroa Etx-	van der Wal, Rui Zhang, Ruochen Zhang, Sebastian	1040
976	abe, Alham Fikri Aji, Amit Alfassy, Anna Rogers,	Gehrmann, Shachar Mirkin, S. Osher Pais, Tatiana	1041
977	Ariel Kreisberg Nitzav, Canwen Xu, Chenghao	Shavrina, Thomas Scialom, Tian Yun, Tomasz Lim-	1042
978	Mou, Chris C. Emezue, Christopher Klamm, Colin	isiewicz, Verena Rieser, Vitaly Protasov, Vladislav	1043
979	Leong, Daniel Alexander van Strien, David Ifeoluwa	Mikhailov, Yada Pruksachatkun, Yonatan Belinkov,	1044
980	Adelani, Dragomir R. Radev, Eduardo Gonz’alez	Zachary Bamberger, Zdenvek Kasner, Zdeněk Kas-	1045
981	Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal		

1046	ner, Amanda Pestana, Amir Feizpour, Ammar Khan,	with iterative retrieval-generation synergy. <i>ArXiv</i> ,	1107
1047	Amy Faranak, Ananda Santa Rosa Santos, Anthony	abs/2305.15294.	1108
1048	Hevia, Antigona Unldreaj, Arash Aghagol, Are-		
1049	zoo Abdollahi, Aycha Tammour, Azadeh HajiHos-	Freda Shi, Xinyun Chen, Kanishka Misra, Nathan	1109
1050	seini, Bahareh Behroozi, Benjamin Ayoade Ajibade,	Scales, David Dohan, Ed Huai hsin Chi, Nathanael	1110
1051	Bharat Kumar Saxena, Carlos Muñoz Ferrandis,	Scharli, and Denny Zhou. 2023a. Large language	1111
1052	Danish Contractor, David M. Lansky, Davis David,	models can be easily distracted by irrelevant context .	1112
1053	Douwe Kiela, Duong Anh Nguyen, Edward Tan, Emi	In <i>International Conference on Machine Learning</i> .	1113
1054	Baylor, Ezinwanne Ozoani, Fatim Tahirah Mirza,		
1055	Frankline Ononiwu, Habib Rezanejad, H.A. Jones,	Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon	1114
1056	Indrani Bhattacharya, Irene Solaiman, Irina Sedenko,	Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and	1115
1057	Isar Nejadgholi, Jan Passmore, Joshua Seltzer,	Wen tau Yih. 2023b. Replug: Retrieval-augmented	1116
1058	Julio Bonis Sanz, Karen Fort, Livia Dutra, Mairon	black-box language models . <i>ArXiv</i> , abs/2301.12652.	1117
1059	Samagaio, Maraim Elbadri, Margot Mieskes, Marissa		
1060	Gerchick, Martha Akinlolu, Michael McKenna, Mike	Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela,	1118
1061	Qiu, Muhammed Ghauri, Mykola Burynek, Nafis	and Jason Weston. 2021. Retrieval augmentation re-	1119
1062	Abrar, Nazneen Rajani, Nour Elkott, Nourhan Fahmy,	duces hallucination in conversation . In <i>Conference</i>	1120
1063	Olanrewaju Samuel, Ran An, R. P. Kromann, Ryan	<i>on Empirical Methods in Natural Language Process-</i>	1121
1064	Hao, Samira Alizadeh, Sarmad Shubber, Silas L.	<i>ing</i> .	1122
1065	Wang, Sourav Roy, Sylvain Viguier, Thanh-Cong		
1066	Le, Tobi Oyebade, Trieu Nguyen Hai Le, Yoyo	Ivan Stelmakh, Yi Luan, Bhuwan Dhingra, and Ming-	1123
1067	Yang, Zachary Kyle Nguyen, Abhinav Ramesh	Wei Chang. 2022. Asqa: Factoid questions meet	1124
1068	Kashyap, Alfredo Palasciano, Alison Callahan, An-	long-form answers . <i>ArXiv</i> , abs/2204.06092.	1125
1069	ima Shukla, Antonio Miranda-Escalada, Ayush Ku-		
1070	mar Singh, Benjamin Beilharz, Bo Wang, Caio	Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and	1126
1071	Matheus Fonseca de Brito, Chenxi Zhou, Chirag	Denny Zhou. 2022. Recitation-augmented language	1127
1072	Jain, Chuxin Xu, Clémentine Fourrier, Daniel Le'on	models . <i>ArXiv</i> , abs/2210.01296.	1128
1073	Perin'an, Daniel Molano, Dian Yu, Enrique Man-		
1074	javacas, Fabio Barth, Florian Fuhrmann, Gabriel	Hugo Touvron, Louis Martin, Kevin R. Stone, Peter	1129
1075	Altay, Giyaseddin Bayrak, Gully Burns, Helena U.	Albert, Amjad Almahairi, Yasmine Babaei, Niko-	1130
1076	Vrabec, Iman I.B. Bello, Isha Dash, Ji Soo Kang,	lay Bashlykov, Soumya Batra, Prajjwal Bhargava,	1131
1077	John Giorgi, Jonas Golde, Jose David Posada, Karthi	Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cris-	1132
1078	Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa	tian Cantón Ferrer, Moya Chen, Guillem Cucurull,	1133
1079	Shinzato, Madeleine Hahn de Bykhovetz, Maiko	David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin	1134
1080	Takeuchi, Marc Pàmies, María Andrea Castillo, Mar-	Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami,	1135
1081	ianna Nezhurina, Mario Sanger, Matthias Samwald,	Naman Goyal, Anthony S. Hartshorn, Saghar Hos-	1136
1082	Michael Cullan, Michael Weinberg, M Wolf, Mina	seini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor	1137
1083	Mihaljcic, Minna Liu, Moritz Freidank, Myung-	Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V.	1138
1084	sun Kang, Natasha Seelam, Nathan Dahlberg,	Korenev, Punit Singh Koura, Marie-Anne Lachaux,	1139
1085	Nicholas Michio Broad, Nikolaus Muellner, Pas-	Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai	1140
1086	cale Fung, Patricia Haller, R. Chandrasekhar, Renata	Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov,	1141
1087	Eisenberg, Robert Martin, Rodrigo L. Canalli, Ros-	Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew	1142
1088	aline Su, Ruisi Su, Samuel Cahyawijaya, Samuele	Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan	1143
1089	Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid	Saladi, Alan Schelten, Ruan Silva, Eric Michael	1144
1090	Kiblawi, Simon Ott, Sinee Sang-aaroonsiri, Srishti	Smith, R. Subramanian, Xia Tan, Binh Tang, Ross	1145
1091	Kumar, Stefan Schweter, Sushil Pratap Bharati, Tan-	Taylor, Adina Williams, Jian Xiang Kuan, Puxin	1146
1092	may Laud, Théo Gigant, Tomoya Kainuma, Wojciech	Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, An-	1147
1093	Kusa, Yanis Labrak, Yashasvi Bajaj, Y. Venkatraman,	gela Fan, Melanie Kambadur, Sharan Narang, Aure-	1148
1094	Yifan Xu, Ying Xu, Yu Xu, Zhee Xiao Tan, Zhongli	lien Rodriguez, Robert Stojnic, Sergey Edunov, and	1149
1095	Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and	Thomas Scialom. 2023. Llama 2: Open foundation	1150
1096	Thomas Wolf. 2022. Bloom: A 176b-parameter	and fine-tuned chat models . <i>ArXiv</i> , abs/2307.09288.	1151
1097	open-access multilingual language model . <i>ArXiv</i> ,		
1098	abs/2211.05100.	H. Trivedi, Niranjan Balasubramanian, Tushar Khot,	1152
		and Ashish Sabharwal. 2021. Musique: Multi-	1153
1099	Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta	hop questions via single-hop question composition .	1154
1100	Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola	<i>Transactions of the Association for Computational</i>	1155
1101	Cancedda, and Thomas Scialom. 2023. Toolformer:	<i>Linguistics</i> , 10:539–554.	1156
1102	Language models can teach themselves to use tools .		
1103	<i>ArXiv</i> , abs/2302.04761.	Liang Wang, Nan Yang, Xiaolong Huang, Binx-	1157
		ing Jiao, Linjun Yang, Daxin Jiang, Rangan Ma-	1158
		jumder, and Furu Wei. 2022. Text embeddings by	1159
		weakly-supervised contrastive pre-training . <i>ArXiv</i> ,	1160
		abs/2212.03533.	1161
1104	Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie		
1105	Huang, Nan Duan, and Weizhu Chen. 2023. En-	Liang Wang, Nan Yang, and Furu Wei. 2023a.	1162
1106	hancing retrieval-augmented large language models	Query2doc: Query expansion with large language	1163

models. In *Conference on Empirical Methods in Natural Language Processing*.

Yile Wang, Peng Li, Maosong Sun, and Yang Liu. 2023b. [Self-knowledge guided retrieval augmentation for large language models](#). *ArXiv*, abs/2310.05002.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed Huai hsin Chi, F. Xia, Quoc Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). *ArXiv*, abs/2201.11903.

Ryen W. White. 2023. [Navigating complex search tasks with ai copilots](#). *ArXiv*, abs/2311.01235.

Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023. [Recomp: Improving retrieval-augmented lms with compression and selective augmentation](#). *ArXiv*, abs/2310.04408.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2022. [React: Synergizing reasoning and acting in language models](#). *ArXiv*, abs/2210.03629.

Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2023. [Making retrieval-augmented language models robust to irrelevant context](#). *ArXiv*, abs/2310.01558.

W. Yu, Hongming Zhang, Xiaoman Pan, Kaixin Ma, Hongwei Wang, and Dong Yu. 2023a. [Chain-of-note: Enhancing robustness in retrieval-augmented language models](#). *ArXiv*, abs/2311.09210.

W. Yu, Zhihan Zhang, Zhenwen Liang, Meng Jiang, and Ashish Sabharwal. 2023b. [Improving language models via plug-and-play retrieval feedback](#). *ArXiv*, abs/2305.14002.

Yuexiang Zhai, Shengbang Tong, Xiao Li, Mu Cai, Qing Qu, Yong Jae Lee, and Y. Ma. 2023. [Investigating the catastrophic forgetting in multimodal large language models](#). *ArXiv*, abs/2309.10313.

Peiyuan Zhang, Guangtao Zeng, Tianduo Wang, and Wei Lu. 2024. [Tinyllama: An open-source small language model](#). *ArXiv*, abs/2401.02385.

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, Z. Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jianyun Nie, and Ji rong Wen. 2023. [A survey of large language models](#). *ArXiv*, abs/2303.18223.

Chunting Zhou, Jiatao Gu, Mona T. Diab, Paco Guzmán, Luke Zettlemoyer, and Marjan Ghazvininejad. 2020. [Detecting hallucinated content in conditional neural sequence generation](#). *ArXiv*, abs/2011.02593.

Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Zhicheng Dou, and Ji rong Wen. 2023. [Large language models for information retrieval: A survey](#). *ArXiv*, abs/2308.07107.

A SlimPLM Implementation Details

Model Fine-tuning Our query rewriting model and the retrieval necessity judgment model are both obtained by instruction fine-tuning from Llama2-7B-Chat. We find that models fine-tuned with data collected from datasets annotated with multiple short answers possess better generalization abilities. They can adapt to various tasks including ambiguous QA, natural questions, long-form QA, and rewritten queries. We collect 5000 samples each from the training sets of ASQA, Natural Questions, and Trivia-QA. Through rule-based filtering, we formed the fine-tuning data for the retrieval necessity judgment model, as shown in Table 5. Because the number of unknown samples significantly exceeds that of known samples, we downsample the unknown samples to make their proportions roughly equal. For the query rewriting model, we collect 1000 samples each from ASQA, Natural Questions, Trivia-QA, MuSiQue, ELI5, and then use GPT-4 for auxiliary annotation. The prompt we use to induce GPT-4 annotation is displayed in Table 8.

RAG Prompts RAG prompts concatenate the reference document in front of the question for enhanced retrieval generation. For datasets annotated with short-form answers and long-form answers, we use different RAG prompts. This is because short-form QA requires the completeness of answers, while long-answer QA demands the fluency of answers. Prompts we use are demonstrated in Table 9. We apply the same prompt strategy across all baselines unless some methods have very strict requirements for prompts, such as Self-Ask (Press et al., 2022) and FLARE (Jiang et al., 2023).

B Baseline Implementation Details

For methods that require multiple rounds of large language model reasoning, we observe that three rounds of reasoning can already solve most of the problems in our dataset. Methods with an indefinite number of reasoning rounds (Self-Ask (Press et al., 2022), FLARE (Jiang et al., 2023)) mostly stop iterating after three rounds. Considering the limitations of computational resources, we set the maximum number of iterations to three rounds. We also set the iteration count of 3 for ITER-RETGEN (Shao et al., 2023).

The results of Self-Ask (Press et al., 2022) on the ELI5 dataset are not compared is that Self-Ask can only output short answers due to prompt limi-

Dataset	Known	Unknown	Dropped
ASQA	593	592	3168
NQ	1873	1874	1253
TQA	588	588	3824

Table 5: The number of unknown, known and dropped samples for retrieval necessity judgment model.

RAG Prompt for FLARE
Search results:
[1] doc 1
[2] doc 2
...
question

Table 6: RAG prompt for FLARE.

tations, which does not meet the ELI5 setting for long text annotations.

The special prompt for FLARE is demonstrated in Table 6.

The special prompt for Self-Ask is demonstrated in Table 10. Specifically, we use the LLM itself as a reader to extract concise answers as intermediate answers from the documents found in search. This was implemented using the Google API in the original paper, but we use our own Wiki document search library, hence the need for this approach.

C Case Study

We provide some cases of misleading references in Table 7. There are mainly two scenarios where searching can have adverse effects: (1) The references retrieved is misleading, the LLM is provided incorrect references; (2) The references retrieved is incomplete, causing the language model to focus on the answer found and overlook other possible answers.

<p>Question: Where are the winter Olympics and when do they start?</p> <p>reference: Åre and Östersund, Sweden will host the next World Winter Games between February 2 to 13, 2021. It will mark the first time that Sweden has ever hosted the Special Olympics World Games.</p> <p>Error Type: irrelevant reference</p> <p>reference: The EOC launched the bid process on 20 September 2018 after a meeting of the constituent National Olympic Committees in Stockholm.</p> <p>Error Type: irrelevant reference</p>
<p>Question: When did the golden state warriors win the finals?</p> <p>reference: The 2017 NBA playoffs began on April 15, 2017. It concluded with the Golden State Warriors defeating the Cleveland Cavaliers 4 games to 1 in the NBA Finals, their third consecutive meeting at the Finals.</p> <p>Error Type: incomplete reference</p> <p>reference: This Finals was the first time in NBA history the same two teams had met for a third consecutive year. The Cavaliers sought to repeat as champions after winning the championship in 2016, while the Warriors won the first meeting in 2015.</p> <p>Error Type: incomplete reference</p>

Table 7: Cases of misleading and incomplete references.

GPT-4 Prompt for Annotating Query Rewrite from User Question

Your task is to perform text analysis on user conversations, and complete the last json item. You need to follow the following rules:

1. Classify user conversations into the following categories: text rewriting, mathematical problems, knowledge questions, text creation, table processing, translation, summarization, logical reasoning, open qa, coding, text classification, information extraction, brainstorming, exams, role-playing, others. The format should be a string and stored in the task field.
2. Determine whether the answer of user input is closely related to current datetime, and store it in the timeliness field in boolean format.
3. If the user's request involves reasoning, each reasoning process should be described as questions and split into as many sub-questions as possible.
4. The sub-questions after splitting should be placed in the question field in questions, and the sub-questions should be fully described without using pronouns such as "he", "this", or "that".
5. If the sub-question involves very strict factual information such as personal relationships, time, location, policies, regulations, etc., which requires the use of a search engine to answer, then it needs to be marked as needSearch=true, and the generated search term should be placed in searchWord.
6. If the sub-question is a chit-chat question such as "how are you" or a pure mathematical problem, coding, logical reasoning, creative thinking, or common sense problem, then no search is needed.
7. Extract the entities and events involved in the user's request and store them in the entities and events fields respectively. The format is a list of strings. Note that the entities and events should be highly informative, and should not be a user instruction or a question.

GPT-4 Prompt for Annotating Query Rewrite from User Question

«SYS»You are asked to first separate a given text by claims and then provide a search query to verify each claim if needed. Here are some requirements: 1. The separation is conducted according to the meaning and each claim should be brief and contain as one key claim. 2. Do not add any hallucinated information or miss any information. 3. The claims should be independent and self-contained, and the claims should be fully described without using pronouns such as "he", "this", or "that". 4. The query is derived from it's corresponding claim and the original user question, and should be useful to check the factuality of the claim. 5. If the claim does not contain any fact relevant with the original user question, or only contains simple common senses, then search is not required. 6. The final return should strictly follow the given format. Like this: <Claims> <Claim(claim1)> <Search(True/False)> <Query(query1)> <Claim(claim2)> <Search(True/False)> <Query(query2)> <Claim(claim3)><Search(True/False)><Query(query3)>.....</Claims> «/SYS»

Table 8: The prompt to induce GPT-4 auxiliary annotation for query rewriting model.

<p>RAG Prompt for Short-Form QA</p> <hr/> <p>«SYS» Now, based on the following reference and your knowledge, please answer the question more succinctly and professionally. The reference is delimited by triple brackets [[[reference]]]. The question is delimited by triple parentheses (((question))). You should include as many possible answers as you can. «/SYS» Reference: [[[reference]]], question: (((question)))</p> <hr/> <p>RAG Prompt for Long-form QA</p> <hr/> <p>«SYS» Now, based on the following reference and your knowledge, please answer the question more succinctly and professionally. The reference is delimited by triple brackets [[[reference]]]. The question is delimited by triple parentheses (((question))). You are not allowed to add fabrications or hallucinations. «/SYS» Reference: [[[reference]]], question: (((question)))</p> <hr/>

Table 9: RAG prompt for different tasks.

RAG Prompt for Self-Ask

«SYS»

Given the following question, answer it by providing follow up questions and intermediate answers. If no follow up questions are necessary, answer the question directly.

«SYS»

Question: Who lived longer, Muhammad Ali or Alan Turing?

Are follow up questions needed here: Yes.

Follow up: How old was Muhammad Ali when he died?

Intermediate answer: Muhammad Ali was 74 years old when he died.

Follow up: How old was Alan Turing when he died?

Intermediate answer: Alan Turing was 41 years old when he died.

So the final answer is: Muhammad Ali

Question: When was the founder of craigslist born?

Are follow up questions needed here: Yes.

Follow up: Who was the founder of craigslist?

Intermediate answer: Craigslist was founded by Craig Newmark.

Follow up: When was Craig Newmark born?

Intermediate answer: Craig Newmark was born on December 6, 1952.

So the final answer is: December 6, 1952

Question: **question**

RAG Prompt for Self-Ask Reference Reader

Given the following reference, answer it by a brief sentence. You are not allowed to add fabrications or hallucinations.

reference

Question: How old was Muhammad Ali when he died?

Answer: Muhammad Ali was 74 years old when he died.

Question: Who was the founder of craigslist?

Answer: Craigslist was founded by Craig Newmark.

Question: Who was the father of Mary Ball Washington?

Answer: The father of Mary Ball Washington was Joseph Ball.

Question: Who is the director of Casino Royale?

Answer: The director of Casino Royale is Martin Campbell.

Question: **question**

Answer:

Table 10: RAG prompt for Self-Ask.