

Supportiveness-based Knowledge Rewriting for Retrieval-augmented Language Modeling

Zile Qiao^{1,3}, Wei Ye^{2,†}, Yong Jiang^{3,†}, Tong Mo¹, Pengjun Xie³
Weiping Li¹, Fei Huang³, Shikun Zhang²

¹ School of Software and Microelectronics, Peking University

² National Engineering Research Center for Software Engineering, Peking University

³ Tongyi Lab, Alibaba Group

wye@pku.edu.cn

{qiaozile.qzl, yongjiang.jy}@alibaba-inc.com

Abstract

Retrieval-augmented language models (RALMs) have recently shown great potential in mitigating the limitations of implicit knowledge in LLMs, such as untimely updating of the latest expertise and unreliable retention of long-tail knowledge. However, since the external knowledge base, as well as the retriever, can not guarantee reliability, potentially leading to the knowledge retrieved not being helpful or even misleading for LLM generation. In this paper, we introduce Supportiveness-based Knowledge Rewriting (SKR), a robust and pluggable knowledge rewriter inherently optimized for LLM generation. Specifically, we introduce the novel concept of "supportiveness"—which represents how effectively a knowledge piece facilitates downstream tasks. Based on supportiveness, we first design a training data curation strategy for our rewriter model, effectively identifying and filtering out poor or irrelevant rewrites to improve data efficacy. We then introduce the direct preference optimization (DPO) algorithm to align the generated rewrites to optimal supportiveness, guiding the rewriter model to summarize augmented content that better improves the final response. Comprehensive evaluations across six popular knowledge-intensive tasks and four LLMs have demonstrated the effectiveness and superiority of SKR. With only 7B parameters, SKR has shown better knowledge rewriting capability over GPT-4.

1 Introduction

Recent advances in large language models (LLMs)(OpenAI, 2023a; Wang et al., 2022; Touvron et al., 2023; Anil et al., 2023) have significantly enhanced their performance in various natural language processing tasks. Pre-training on large-scale unsupervised corpora enables LLMs

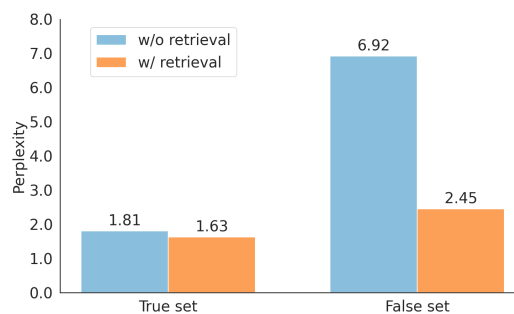


Figure 1: To measure the change in perplexity induced by correct knowledge, we sampled 500 correctly answered questions (True set) and 500 incorrectly answered questions (False set) from TriviaQA. We then calculate the perplexity before and after introducing retrieved knowledge, where the retrieved knowledge included the correct answers.

to store extensive knowledge within their parameters. Updating LLMs with recent information is challenging due to high training costs (Lin et al., 2023b), and sensitive data is often excluded to prevent leakage (Huang et al., 2022; Carlini et al., 2020). Retaining long-tail knowledge is also a bottleneck (Wang et al., 2023a). Retrieval-Augmented Language Modeling (RALM) addresses these issues by integrating non-parametric knowledge with LLMs (Borgeaud et al., 2021; Lin et al., 2023b; Shi et al., 2023a; Ram et al., 2023a), achieving remarkable results in various tasks.

However, the significant noise and misleading information in the retrieved content can impair the performance of RALM. Recent efforts have aimed to mitigate this issue by rewriting the retrieved content. CoN (Yu et al., 2023) and RECOMP (Xu et al., 2023a) leverage off-the-shelf strong LLMs as supervision signal providers to train the rewrite. More recently, Wang et al. (2023b); Jiang et al. (2024) and Jin et al. (2024) have attempted to measure the quality of retrieved knowledge by assessing the perplexity (or probability) changes in the outputs of LLMs induced by retrieved knowledge,

[†]Corresponding authors.

providing supervision signals for model training. However, these methods overlook the influence of the internal knowledge of LLMs. When an LLM’s internal knowledge is sufficient to respond to a particular query accurately, the incorporation of retrieved knowledge, even if entirely accurate, may not significantly alter the output probabilities of the LLM, as illustrated in Figure 1. This phenomenon leads to feedback from the LLM that does not accurately reflect the quality of the retrieved knowledge.

To address this issue, we introduce the concept of "supportiveness". Beyond analyzing the variations in LLM output perplexity induced by retrieval knowledge, we further examined how well the internal knowledge of the LLM covers various queries. This approach enabled us to attain a more precise evaluation of the quality of the retrieved knowledge. Building on supportiveness, we present Supportive Knowledge Rewriting (SKR), an abstractive rewriter inherently optimized for LLM generation. This rewriter is trained using two supportiveness-based mechanisms.

The first one is Supportiveness-based Rewrite Data Generation. Though powerful general-purpose LLMs are widely known for their capability of serving as data annotators (Yu et al., 2023; Wang et al., 2022), directly using LLM-generated rewrite data to train the rewriter may yield sub-optimal performance due to the bias and potential inaccuracy in LLM responses. Therefore, we use supportiveness to assess the quality of automatically generated rewrite data and then filter and refine it. In particular, we initially utilize a powerful LLM to create multiple different rewrites for each query based on a wide array of downstream tasks and retrieval data, serving as our preliminary draft dataset. We then engage a white-box LLM for each rewritten text to determine its supportiveness. Based on the supportiveness score, we refined the draft dataset, and the refined datasets were subsequently used for supervised fine-tuning (SFT).

The second mechanism is Supportiveness-guided Alignment, a training strategy that aligns the generated rewrite to optimal supportiveness using Direct Preference Optimization (DPO) (Rafailov et al., 2023). Specifically, we pair two different rewrites for the same query and calculate the supportiveness difference within each pair. Pairs with more significant discrepancies are selected for DPO training. In each selected pair, the rewrite with higher supportiveness serves as the positive sample, while the one with lower sup-

portiveness serves as the negative sample. This process encourages SKR to generate rewrites that better support LLMs in performing downstream tasks. Our contributions are as follows:

- We conceptualized "supportiveness", offering a novel perspective to assess how effectively an augmented knowledge piece contributes to a specific query.
- Leveraging supportiveness, we devise an effective knowledge rewriter by incorporating (1) a supportiveness-data rewrite data generation method that improves data efficacy and (2) a direct preference training mechanism to align rewritten text to optimal supportiveness better. Experiments conducted on six datasets have thoroughly validated the effectiveness and generalization of SKR. With merely 7B parameters, SKR has a knowledge rewriting capability surpassing GPT-4.
- Further experimental analysis validated that SKR effectively removes noise and misleading information in retrieval data, achieving a compression rate exceeding 7x.

We will make our source code public to encourage reproducibility.

2 Method

2.1 Architecture

Retriever We utilize DRAGON+ (Lin et al., 2023a) as the retriever, which is a state-of-the-art dual-encoder model. Given a query q and a corpus \mathcal{C} , the query encoder \mathbf{E}_q maps q to an embedding $\mathbf{E}_q(q)$ and the document encoder \mathbf{E}_d maps each text chunk $c \in \mathcal{C}$ to an embeddings $\mathbf{E}_d(c)$. After that, we compute the similarity score $s(q, c)$ for each $c \in \mathcal{C}$ by dot product: $s(q, c) = \mathbf{E}_d(c) \cdot \mathbf{E}_q(q)$. The top-k relevant text chunks for q are selected based on their similarity scores.

In-Context Retrieval-Augmentation For a given query q , retrieved knowledge chunks or the rewrites generated by the rewriter, we map them to a prompt S according to the template to construct the input of language model to generate the response y to the query q :

$$P_{LM}(y_1, y_2, \dots, y_n) = \prod_{i=1}^n P_{LM}(y_i | S).$$

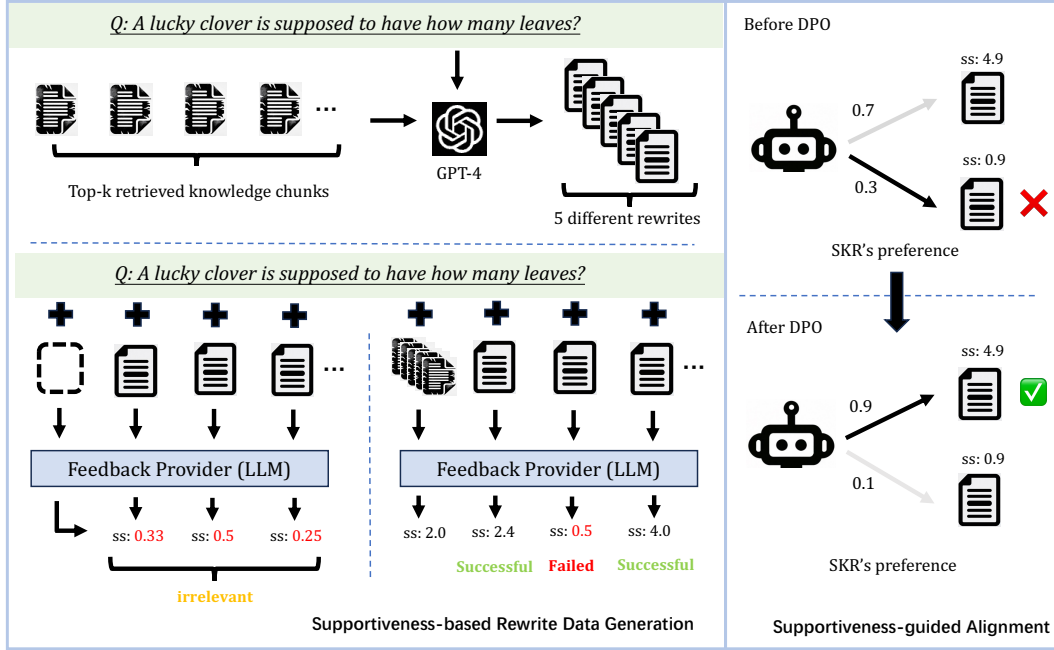


Figure 2: A visual illustration of the proposed Supportiveness-based Rewrite Data Generation and Supportiveness-guided Alignment. “ss” denotes the supportiveness score.

Rewriter To eliminate unhelpful or misleading information from retrieved knowledge, we proposed SKR as the knowledge rewriter. The rewriter is designed to refine retrieved knowledge through rewriting, aiming to optimize the utilization efficiency of external knowledge by LLMs. Given a query q and k candidate retrieval chunks $c_q \in \mathcal{C}_q$, the input text T to the rewriter is a concatenation of the query q and k retrieved knowledge chunks, separated by newline characters. The rewriter generates rewrites in an autoregressive manner. The generated rewrite will serve as part of the prompt to assist the LLM in handling downstream tasks.

2.2 Supportiveness

Supportiveness indicates the degree to which a piece of knowledge assists in answering a specific query. To estimate supportiveness, we first need to employ an LLM (e.g. Mistral) to attempt to answer the current query twice, once with the integration of retrieved knowledge and once without it. Then, based on the correct label of the query, we calculate the perplexity of both outputs. The formal description of this process is as follows:

$$\begin{aligned} \mathcal{P}_{\text{raw}} &= \exp\left(-\frac{1}{N} \sum_{i \in \text{target}} \log(P_{LM}(w_i | S_q \circ w_{<i}))\right) \\ \mathcal{P}_{\text{retrieval}} &= \exp\left(-\frac{1}{N} \sum_{i \in \text{target}} \log(P_{LM}(w_i | S_q^r \circ w_{<i}))\right), \end{aligned} \quad (1)$$

where $\mathcal{P}_{\text{retrieval}}$ and \mathcal{P}_{raw} denote the perplexity of the output with retrieval and the output without retrieval, respectively. $\{w_1, \dots, w_n\}$ denotes the target sequence of the query and \circ denote the text concatenation. For cases with multiple target sequences, we calculate the average of the perplexities computed from the numerous target sequences. S_q^r and S_q represent the prompt with and without retrieval, which are shown in Appendix A. Consequently, we calculate the ratio of \mathcal{P}_{raw} to $\mathcal{P}_{\text{retrieval}}$, denoted as $r(q, c)$, where c represents a specific knowledge piece. Intuitively, the ratio $r(q, c)$ can represent the degree to which a specific knowledge piece supports the query. However, we note that when the LLM can easily answer the question without any retrieval knowledge (e.g., \mathcal{P}_{raw} is relatively low), $r(q, c)$ cannot accurately measure the supportiveness of the knowledge piece c . In such cases, even if the knowledge piece c contains key information to answer the query, it will not be reflected in the metric because LLM already possesses this knowledge.

To address this issue, we propose supportiveness score, denoted as $ss(q, c)$, which is calculated as:

$$\begin{aligned} r(q, c) &= \mathcal{P}_{\text{raw}} / \mathcal{P}_{\text{retrieval}}, \\ ss(q, c) &= r(q, c) / \sigma(\mathcal{P}_{\text{raw}}), \end{aligned} \quad (2)$$

where σ denotes the sigmoid function. Compared to directly using the ratio of \mathcal{P}_{raw} to $\mathcal{P}_{\text{retrieval}}$, the

supportiveness score more accurately represents the supportiveness of a knowledge piece for a query, mitigating the bias introduced by LLM’s internal knowledge.

2.3 Supportiveness-based Rewrite Data Generation

In this section, we provide a detailed description of Supportiveness-guided Rewrite Data Generation (SRDG). The left part of Figure 2 demonstrates a visual depiction of these approaches.

2.3.1 Automated Rewrite Data Generation

We first need to collect a substantial amount of rewrite data to train the rewriter. Considering that manual annotation is resource-intensive and it is well-known that powerful general-purpose LLMs can serve as data annotators (Yu et al., 2023), we employed the current state-of-the-art LLM, GPT-4, as a data annotator to generate our draft dataset. To ensure the diversity of the data, we used the training sets from a total of 10 datasets across three tasks (Open-Domain QA, Reading Comprehension, and Reasoning) as the sources for generating rewrite data. The prompt for GPT-4 is in Appendix A.

To enhance the SKR’s ability to generalize across varying lengths of context, we use the top-5 retrieved chunks for 80% of the samples. For the remaining 20% of the samples, we evenly distribute the use of the top-1, top-2, top-3, and top-4 retrieved knowledge chunks as prompts. To further enrich the rewrite data and also to enable subsequent adjustments to the generation preference of SKR, we sampled five different rewrites for each query.

2.3.2 Supportiveness-based Data Refinement

Then, we introduce how to refine the draft dataset based on supportiveness. Through the supportiveness score, we are able to assess the degree to which a specific knowledge piece aids the query. Next, we will refine the draft dataset based on the supportiveness score. We categorize rewrites into three groups: “irrelevant,” “failed,” and “successful”.

Given a query q , its retrieval data c , and the rewrites $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$ generated by GPT-4, we will first label the sample as “irrelevant” if the following conditions are met:

$$\begin{aligned} ss_{r_i} < 1 \quad \text{for } r_i \in \mathcal{R}, \\ ss_c < 1, \end{aligned} \quad (3)$$

where ss_{r_i} and ss_c denote the supportiveness score of r_i and raw retrieval data, respectively. In this

situation, we eliminate all of the rewritten content generated by GPT-4 and change the label rewrite of c to the word “irrelevant”. This scenario suggests that regardless of whether it is rewritten or not, the retrieved knowledge cannot assist the LLM in answering the query. Therefore, this refinement enhances the rewriter’s ability to recognize unhelpful retrieved knowledge and further compresses the length of the rewrites to improve the overall efficiency of RALM. Then, we will label the rewrite as “failed” if the following conditions are met: $ss_{r_i}/ss_c < \delta$. This situation implies that the specific rewrite r_i is failed, as this rewrite, compared to the raw retrieval data, does not offer any additional assistance to the LLM. Therefore, we simply discard these rewrites. The remaining samples are labeled as “successful”. These, along with the ‘irrelevant’ samples, are then used for the training of SKR, employing a standard supervised fine-tuning approach:

$$\pi_{\text{SFT}} = \max_{\pi} \mathbb{E} \log \pi(r \mid q \circ c), \quad (4)$$

where $\{r_i \mid i = 1 \dots k\}$ denotes k different rewrites for a given pair of (q, c) and \circ denotes text concatenation. Through supportiveness-based data refinement, approximately 9 % of the samples were labeled as “irrelevant” and modified, while 16% were labeled as “failed” and discarded.

2.4 Supportiveness-guided Alignment

After Supervised fine-tuning on the refined dataset, SKR is already capable of generating decent rewrites. To further strengthen SKR and make its outputs more supportive of LLMs in completing downstream tasks, we will alter the generative preferences of SKR to make it more inclined towards generating more supportive rewrites.

First, we construct preference data based on the draft dataset and the supportiveness score calculated for each rewrite. Given a query q , the raw retrieval data c and the rewrites set $\mathcal{R} = \{r_1, r_2, \dots, r_n\}$ from draft dataset (we additionally include “irrelevant” in \mathcal{R} as an extra version of the rewrite), we construct the preference data \mathcal{D}_p as follows:

$$\mathcal{D}_p = \{(x, r_w, r_l) \mid r_w, r_l \in \mathcal{R}, ss_w - ss_l > 1\},$$

where ss_w and ss_l respectively represent the supportiveness scores of the corresponding rewrites. $x = q \circ c$ denotes the input of SKR. This construction implies that the supportiveness score of r_l in

the preference data is significantly lower than that of r_w . Then, we employed DPO (Rafailov et al., 2023) to align the generative preferences of SKR with the supportiveness of the rewrites.

3 Experiment Setup

3.1 Implementation Details

Automated Rewrite Data Generation To construct the draft dataset, we utilized GPT-4 to automatically annotate rewrite data. We set the temperature to 1, "n" (the number of chat completion choices) to 10, and randomly sampled responses for rewrite data, keeping all other settings at their defaults.

SKR SKR was initialized with Mistral-7B, which also served as the feedback provider during supportiveness calculation. In both SFT and DPO, the hyperparameters used are consistent with those in zephyr (Tunstall et al., 2023), except for the learning rate, which we set to 1e-05 during the SFT process. We used 10 datasets to train SKR, including four Open-domain QA datasets, four Reading comprehension tasks and two reasoning datasets (Detailed dataset statistics can be found in Appendix B). Note that we did not use any evaluation datasets for training.

Others Most of our evaluations are conducted using vLLM (Kwon et al., 2023). We conducted all training and evaluations on 4x NVIDIA A100 80GB GPUs. The process of SFT for training SKR took approximately 30 hours, and the DPO process took about 40 hours. The retrieval corpus used in this paper is Wikipedia dump 2018.

3.2 Evaluation

Datasets We have conducted evaluations on knowledge-intensive tasks and ensured that the dataset used for training the SKR does not appear in them. Specifically, we use Massive Multitask Language Understanding (MMLU) (Hendrycks et al., 2021), Natural Questions (NQ) (Kwiatkowski et al., 2019), TriviaQA (Joshi et al., 2017), HotpotQA (HoPo) (Yang et al., 2018), WebQuestions (WQ) (Berant et al., 2013), and zero-shot Relation Extraction (Levy et al., 2017) (zsRE) from KILT tasks (Petroni et al., 2020). We use dev split for Natural Questions, TriviaQA, and zsRE, and test split for MMLU, WebQuestions, and HotpotQA. All experimental results are presented using the EM (Exact Match) metrics. For the evaluations of these datasets, we use the same prompts as mentioned in

Appendix A.

Baselines To ensure the fairness of the evaluation, we selected GPT-3.5 and GPT-4 (OpenAI, 2023b), Filco (Wang et al., 2023b), and RECOMP (Xu et al., 2023a) as baselines. We used the same retrieval corpus and the same prompt (for generating the final answer), while other methods that could not guarantee fairness were not included as baselines.

4 Experiment Results

In this section, we report the experimental results (statistically significant with $p < 0.05$). Due to the page limit, we have placed the analysis results regarding the choice of retriever in Appendix D. Another analysis experiment related to the feedback provider is included in Appendix E.

4.1 Main Results

We report the main results in Table 1. The experimental results using Llama-2-13B and Mistral-7B as the base model can be found in Appendix C. From these results, we make the following four observations:

Negative Impact of Noise Our first observation is that adding the Top-5 knowledge chunks, as opposed to just the Top-1 knowledge chunk, does not significantly enhance model performance as expected, even though the Top-5 knowledge chunks clearly contain more useful information (compared to the Top-1 setting, the Top-5 setting resulted in changes of -4.67, -2.48, +0.45, and -5.49 in the EM metric for Llama-2-7B, -13B, -70B, and Mistral-7B, respectively, on average across six different tasks). *The additional noisy data introduced by more retrieved knowledge chunks has a significant negative impact on the LLMs, affirming the necessity of rewriting the retrieved knowledge.*

Effectiveness of Supportiveness-based Rewrite Our second observation is that *the proposed SKR significantly enhances the LLMs' ability to utilize retrieved knowledge, markedly reducing the negative impact of noise on various LLMs* (compared to the Top-5 setting without a rewriter, SKR consistently yielded improvements of +10.01 and +3.30 in the EM metric for Llama-2-7B and Llama-2-70B, respectively, on average across six different tasks). This result validates the capability of SKR to rewrite retrieved knowledge, confirming the effectiveness of this approach.

Effectiveness of Supportiveness-based Training

Model	Rewriter	Retrieval	NQ	TQA	HoPo	MMLU	zsRE	WQ	Avg.	Δ_{EM}
Llama-2-7B	-	-	19.61	52.04	18.04	39.96	17.07	19.01	27.62	-
	-	Top-1	23.55	58.82	28.04	43.17	56.39	19.59	38.26	-
	-	Top-5	17.5	56.8	25.43	44.53	42.4	14.91	33.6	-
	GPT-3.5	Top-5	27.78	61.71	30.16	45.99	59.22	21.3	41.03	7.43
	GPT-4	Top-5	29.82	63.76	30.52	47.77	60.04	22.83	42.46	8.86
	RECOMP	Top-5	27.02	61.04	28.29	-	-	-	-	-
	FILCO†	Top-5	28.27	60.78	28.44	-	59.27	20.18	-	-
	LongLLMLingua	Top-5	27.44	60.28	27.92	44.23	59.44	21.01	40.05	-
	SKR-7B	Top-5	31.42	63.62	30.98	48.45	61.97	25.24	43.61	10.01
Llama-2-70B	-	-	30.61	68.82	26.59	62.84	28.65	22.21	39.95	-
	-	Top-1	33.57	69.14	33.6	65.58	63.89	25.01	48.47	-
	-	Top-5	35.32	69.54	32.16	66.15	65.17	25.17	48.92	-
	GPT-3.5	Top-5	35.68	69.47	36.24	66.8	65.91	28.22	50.39	1.47
	GPT-4	Top-5	36.51	69.83	38.22	67.09	67.15	30.15	51.49	2.57
	RECOMP	Top-5	35.47	69.44	34.47	-	-	-	-	-
	FILCO†	Top-5	35.5	69.81	35.68	-	66.12	26.44	-	-
	LongLLMLingua	Top-5	33.42	69.11	34.17	65.78	62.88	28.78	49.02	-
	SKR-7B	Top-5	37.97	70.68	38.24	67.87	66.91	31.67	52.22	3.3

Table 1: Main results. All experimental results are presented using the EM (Exact Match) metrics, “ Δ_{EM} ” represents the average performance difference of the corresponding rewriting method compared to using “Top-5” retrieved data without any rewriting. “†” denotes our implementations.

Models	Rewriter	NQ			WQ		
		None	Bad	Random	None	Bad	Random
Llama-2-7B	None	19.61	10.08	10.31	19.01	9.30	12.84
	SKR-7B	-	20.11 (+10.03)	19.55 (+9.24)	-	22.95 (+13.65)	18.68 (+5.84)
Llama-2-13B	None	26.5	15.37	23.13	24.05	16.14	21.34
	SKR-7B	-	24.91 (+9.54)	26.23(+3.10)	-	27.77 (+11.63)	24.26 (+2.92)
Mistral-7B	None	25.73	18.25	19.78	21.8	15.80	19.13
	SKR-7B	-	22.32 (+4.07)	24.19(+4.41)	-	24.37 (+8.57)	21.22 (+2.09)

Table 2: Impact of noise and misleading information. “Bad” and “Random” denote using the Top-5 knowledge chunks respectively disturbed by the two methods described in section 4.2, while “None” indicates not using any retrieved knowledge.

Strategy Our third observation is that *guided by the concept of supportiveness, SKR, with only 7B parameters, has exhibited a knowledge rewriting capability that surpasses representative powerful off-the-shelf LLMs, including the current state-of-the-art general-purpose LLM, GPT-4.* Compared to GPT-4, SKR yielded consistent improvements in the EM metric for Llama-2-7B and Llama-2-70B of +1.15 and +0.73, respectively, on average across six different tasks. Compared to GPT-3.5, the improvements brought by SKR are even more pronounced. These results demonstrate the effectiveness of our proposed supportiveness-based rewrite data generation and supportiveness-based alignment methods. A detailed analysis of the individual contributions of these two methods will be presented in Ablation Studies.

Performance Superiority Compared to Other

Rewriting Methods The final observation is that *SKR’s performance significantly surpasses other rewriting methods.* SKR shows substantial performance gains across all six tasks compared to FILCO, RECOMP and LongLLMLingua. It’s important to note that FILCO and RECOMP require training the rewriter on corresponding training sets, whereas SKR is evaluated in a zero-shot setting, further validating SKR’s performance advantage.

4.2 Impact of Noise and Misleading Information

To further validate SKR’s ability to eliminate noisy or misleading information, we replaced retrieved knowledge with two types of interference data: 1) Bad retrieval, where the order of the top-30 retrieved knowledge pieces is reversed to simulate scenarios where the retrieved knowledge is mislead-

ing. This type of data may still contain helpful information, but it also includes a significant amount of misleading information; 2) Random, where all retrieved knowledge is replaced with random samples from the entire corpus, to simulate a scenario where the retrieved information is irrelevant. The experimental results are presented in Table 2. We observed that both types of interference data negatively impacted the model. Taking Llama-2-7B as an example, the "Bad" type caused a decline of -9.53 and -9.71 in EM on NQ and WQ respectively, compared to not using retrieval data, while the "Random" type led to a decline of -9.3 and -7.17 in EM on NQ and WQ respectively. We then observed that SKR significantly mitigated the negative impact caused by severe interference in the retrieved data. Taking Llama-2-7B as an example, SKR under the "Bad" setting brought about EM improvements of +10.03 and +13.65 on NQ and WQ respectively. Under the "Random" setting, it resulted in EM improvements of +9.24 and +5.84 on NQ and WQ respectively. These experimental results fully demonstrate that *SKR can effectively eliminate noise and misleading information.*

4.3 Effects on Compression Rate

To verify SKR’s capability to compress retrieved knowledge through rewriting, and to analyze the impact of supportiveness-based rewrite data generation and supportiveness-based alignment from the perspective of compression rate, we conducted the experiments presented in Table 3.

In this experiment, we used four different settings: 1) SKR w/o SRDG & SA, which refers to training directly using the draft dataset generated by GPT-4; 2) SKR w/o SA, which means training using only the supportiveness-based rewrite data generation; 3) SKR w/o SRDG, which implies training using only the supportiveness-based alignment; 4) SKR, which represents our complete method. We first observed that the full SKR achieved remarkable compression rates of 7.57x and 7.60x on the NQ and TQA datasets, respectively. Furthermore, we noted that supportiveness-based rewrite data generation significantly positively impacts the compression rate (compression rate of 6.27x under the SKR w/o SRDG & SA setting, and a rate of 7.49x under the SKR w/o SA setting). This can be attributed to the capability of recognizing irrelevant knowledge that SRDG brings to SKR.

Rewriter	NQ		WQ	
	length	comp.	length	comp.
-	815.4	1x	842.7	1x
GPT-4	140.9	5.79x	161.1	5.23x
w/o SRDG & SA	130.1	6.27x	130.9	6.44x
w/o SA	108.8	7.49x	106.8	7.89x
w/o SRDG	121.4	6.72x	124.4	6.77x
SKR	107.7	7.57x	110.9	7.60x

Table 3: Effects on Compression Rate. We used Llama-2-7B as the evaluation model. “length” denotes the average length of the token sequences, and “comp.” represents the compression rate.

Model	Rewriter	NQ	
		EM	Irr. rate
Llama-2-7B	GPT-4	28.2	0%
	GPT-4†	29.82	38.9%
	SKR	31.22	7.1%

Table 4: Analysis of Recognizing Irrelevant Knowledge. “Irr. rate” denotes the rate at which the rewriter output “irrelevant” and “GPT-4†” denotes using GPT-4 as the rewriter with modified prompt described in section 4.4.

4.4 Analysis of Recognizing Irrelevant Knowledge

By employing the supportiveness-based rewrite data generation described in section 2.3, SKR has gained the capability to recognize irrelevant knowledge. Existing general-purpose LLMs like GPT-4 can also discern irrelevant knowledge through instruction prompts. To analyze and compare these two modes, we have constructed the evaluation shown in Table 4. Specifically, we modified GPT-4’s prompt to enable it to independently assess whether the retrieved knowledge contains information that aids in answering the query and to output "irrelevant" if it determines the information is not helpful. Subsequently, we recorded the rate at which different rewriter output "irrelevant", denoted as "Irr. rate," as shown in Table 4. The modified prompt for GPT-4 is shown in Appendix A.

It can be observed that allowing GPT-4 to independently assess whether the retrieved knowledge contains information that aids in answering the query does not help improve the final performance. Additionally, it is noted that GPT-4 is quite stringent in identifying whether retrieved knowledge contains irrelevant information, with its Irr. rate significantly higher than SKR’s (38.9% vs. 7.1%). We believe that GPT-4 is overly strict in its judg-

ments, to the extent that it overlooks some knowledge that, while not overtly obvious, could provide supplementary lateral information to LLMs. Therefore, we utilize the proposed supportiveness-based rewrite data generation method to enable SKR to recognize irrelevant information, rather than relying on GPT-4 as the annotator for this task.

4.5 Ablation Studies

To thoroughly analyze the contributions of supportiveness-based rewrite data generation and supportiveness-based alignment to the performance of SKR, we conducted ablation experiments on the aforementioned two modules. The results of these experiments are presented in Table 5. All settings uniformly utilize the top-5 knowledge chunks as retrieval data. ‘‘SRDG’’ denotes supportiveness-based rewrite data generation, and ‘‘SA’’ denotes supportiveness-based alignment.

We observed that even without utilizing the two proposed supportiveness-based methods, SKR’s rewriting performance is commendable. However, compared to GPT-4, there is a decrease in performance (SKR results in an average decrease of -0.79 in the EM metric for Llama-2-7B, across two tasks). Another observation is that SRDG and SA individually contributed to performance improvements in SKR, ultimately surpassing the performance of GPT-4. SKR achieved an average increase of +2.00 in the EM metric for Llama-2-7B, across two tasks.

Rewriter	NQ	WQ	Avg.
-	17.5	14.91	16.21
GPT-4	29.82	22.83	26.33
SKR w/o SRDG & SA	29.33	21.75	25.54
SKR w/o SA	30.01	22.61	26.31
SKR w/o SRDG	30.48	24.7	27.59
SKR	31.42	25.24	28.33

Table 5: Ablation with Llama-2-7B.

5 Related Works

Existing RALM research can be broadly classified into two categories. The first one primarily focuses on enhancing LLMs via pre-training or fine-tuning, aiming to improve their ability to effectively harness the retrieved knowledge (Borgeaud et al., 2021; Guu et al., 2020a; Lin et al., 2023b; Yu et al., 2023). Another line of RALM efforts explores directions beyond training LLMs, encompassing but not limited to 1) merging multiple generations based on diverse knowledge-enhanced queries (Shi

et al., 2023a; Lin et al., 2023b), 2) re-ranking candidate knowledge chunk (Ram et al., 2023a; Jiang et al., 2023), or 3) improving retrievers by training them with specific language tasks (Shi et al., 2023a; Izacard et al., 2022c). RALM can address the issues faced by large language models, such as the difficulty in timely updating information and the unreliable memory of long-tail knowledge (Wang et al., 2023a; Huang et al., 2022; Carlini et al., 2020; Khandelwal et al., 2020; Guu et al., 2020b; Ram et al., 2023b; Izacard et al., 2022b; Asai et al., 2023; Shi et al., 2023b; Khattab et al., 2022).

However, the presence of noise or misleading information in the retrieved content can impair the performance of RALM and the extensive context information also introduces additional computational overhead (Wang et al., 2023b; Chen et al., 2023; Xu et al., 2023b; Zhu et al., 2024; Yoran et al., 2024). This has inspired research into compressing or rewriting the retrieved content. RECOMP (Xu et al., 2023a) trains an extractive compressor based on the similarity between queries and sentences and uses GPT-3.5 as a teacher model to train an abstractive rewriter. CoN (Yu et al., 2023) employs GPT-3.5 as a teacher model, targeting answer prediction and context summarization simultaneously. Most recently, Filco (Wang et al., 2023b) identifies the importance of sentences within the context by integrating multiple strategies. Jiang et al. (2024) and Jin et al. (2024) have sought to evaluate the quality of retrieved knowledge by examining the changes in perplexity (or probability) in the outputs of LLMs caused by the retrieved knowledge, thereby providing supervision signals for model training.

6 Conclusion

In this paper, we introduced the innovative concept of ‘‘supportiveness’’, providing a simple yet powerful approach to evaluate how effectively a knowledge piece facilitates language tasks. Based on the supportiveness, we have developed the supportiveness-based rewrite data generation and supportiveness-guided alignment methods, harnessing them to train SKR, an effective knowledge rewriter. Our comprehensive experiments across six datasets and four LLMs validate SKR’s effectiveness and generalizability. With only 7B parameters, SKR demonstrates superior knowledge rewriting capabilities compared to GPT-4.

Limitations

Like other abstractive context rewrite methods, SKR generates rewrites using an autoregressive approach. This type of method carries the risk of potentially altering the retrieved context and producing hallucinations. Currently, in this paper, we propose a supportiveness score to evaluate the quality of the rewrite, but we have yet to assess whether the generated rewrite remains faithful to the raw retrieval context. We will consider this in future work.

Ethics Statement

This work was conducted in compliance with the ACL Ethics Policy. All datasets and models used for evaluation are publicly available. We do not anticipate any negative ethical impacts from our work.

References

- Shourya Aggarwal, Divyanshu Mandowara, Vishwajeet Agrawal, Dinesh Khandelwal, Parag Singla, and Dinesh Garg. 2021. [Explanations for commonsenseqa: New dataset and models](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Rohan Anil, Andrew M. Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, Eric Chu, Jonathan H. Clark, Laurent El Shafey, Yanping Huang, Kathy Meier-Hellstern, Gaurav Mishra, Erica Moreira, Mark Omernick, Kevin Robinson, Sebastian Ruder, Yi Tay, Kefan Xiao, Yuanzhong Xu, Yujing Zhang, Gustavo Hernandez Abrego, Junwhan Ahn, Jacob Austin, Paul Barham, Jan Botha, James Bradbury, Siddhartha Brahma, Kevin Brooks, Michele Catasta, Yong Cheng, Colin Cherry, Christopher A. Choquette-Choo, Aakanksha Chowdhery, Clément Crepy, Shachi Dave, Mostafa Dehghani, Sunipa Dev, Jacob Devlin, Mark Díaz, Nan Du, Ethan Dyer, Vlad Feinberg, Fangxiaoyu Feng, Vlad Fienber, Markus Freitag, Xavier Garcia, Sebastian Gehrmann, Lucas Gonzalez, Guy Gur-Ari, Steven Hand, Hadi Hashemi, Le Hou, Joshua Howland, Andrea Hu, Jeffrey Hui, Jeremy Hurwitz, Michael Isard, Abe Ittycheriah, Matthew Jagielski, Wenhao Jia, Kathleen Kenealy, Maxim Krikun, Sneha Kudugunta, Chang Lan, Katherine Lee, Benjamin Lee, Eric Li, Music Li, Wei Li, YaGuang Li, Jian Li, Hyeontaek Lim, Hanzhao Lin, Zhongtao Liu, Frederick Liu, Marcello Maggioni, Aroma Mahendru, Joshua Maynez, Vedant Misra, Maysam Moussalem, Zachary Nado, John Nham, Eric Ni, Andrew Nystrom, Alicia Parrish, Marie Pellat, Martin Polacek, Alex Polozov, Reiner Pope, Siyuan Qiao, Emily Reif, Bryan Richter, Parker Riley, Alex Castro Ros, Auroko Roy, Brennan Saeta, Rajkumar Samuel, Renee Shelby, Ambrose Slone, Daniel Smilkov, David R. So, Daniel Sohn, Simon Tokumine, Dasha Valter, Vijay Vasudevan, Kiran Vodrahalli, Xuezhi Wang, Pidong Wang, Zirui Wang, Tao Wang, John Wieting, Yuhuai Wu, Kelvin Xu, Yunhan Xu, Linting Xue, Pengcheng Yin, Jiahui Yu, Qiao Zhang, Steven Zheng, Ce Zheng, Weikang Zhou, Denny Zhou, Slav Petrov, and Yonghui Wu. 2023. [Palm 2 technical report](#). *Preprint*, arXiv:2305.10403.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). *Preprint*, arXiv:2310.11511.
- Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. [Semantic parsing on Freebase from question-answer pairs](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1533–1544, Seattle, Washington, USA. Association for Computational Linguistics.
- Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, T. W. Hennigan, Saffron Huang, Lorenzo Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and L. Sifre. 2021. [Improving language models by retrieving from trillions of tokens](#). In *International Conference on Machine Learning*.
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom B. Brown, Dawn Xiaodong Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2020. [Extracting training data from large language models](#). In *USENIX Security Symposium*.
- Howard Chen, Ramakanth Pasunuru, Jason Weston, and Asli Celikyilmaz. 2023. [Walking down the memory maze: Beyond context limit through interactive reading](#). *ArXiv*, abs/2310.05029.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *ArXiv*, abs/2110.14168.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.

- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020a. [Realm: Retrieval-augmented language model pre-training](#). *ArXiv*, abs/2002.08909.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020b. REALM: Retrieval-augmented language model pre-training. *arXiv preprint arXiv:2002.08909*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Jie Huang, Hanyin Shao, and Kevin Chen-Chuan Chang. 2022. [Are large pre-trained language models leaking your personal information?](#) *ArXiv*, abs/2205.12628.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022a. [Unsupervised dense information retrieval with contrastive learning](#). *Transactions on Machine Learning Research*.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022b. [Atlas: Few-shot learning with retrieval augmented language models](#). *Preprint*, arXiv:2208.03299.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane A. Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022c. [Few-shot learning with retrieval augmented language models](#). *ArXiv*, abs/2208.03299.
- Huiqiang Jiang, Qianhui Wu, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2023. [LLMLingua: Compressing prompts for accelerated inference of large language models](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13358–13376, Singapore. Association for Computational Linguistics.
- Huiqiang Jiang, Qianhui Wu, Xufang Luo, Dongsheng Li, Chin-Yew Lin, Yuqing Yang, and Lili Qiu. 2024. [Longllmlingua: Accelerating and enhancing llms in long context scenarios via prompt compression](#). *Preprint*, arXiv:2310.06839.
- Kelvin Jiang, Dekun Wu, and Hui Jiang. 2019. [FreebaseQA: A new factoid QA data set matching trivia-style question-answer pairs with Freebase](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 318–323, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jiajie Jin, Yutao Zhu, Yujia Zhou, and Zhicheng Dou. 2024. [Bider: Bridging knowledge inconsistency for efficient retrieval-augmented llms via key supporting evidence](#). *Preprint*, arXiv:2402.12174.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1601–1611, Vancouver, Canada. Association for Computational Linguistics.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. [Generalization through memorization: Nearest neighbor language models](#). In *International Conference on Learning Representations*.
- Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts, and Matei Zaharia. 2022. [Demonstrate-search-predict: Composing retrieval and language models for knowledge-intensive nlp](#). *arXiv preprint arXiv:2212.14024*.
- 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.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Haoteng Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). *Proceedings of the 29th Symposium on Operating Systems Principles*.
- Omer Levy, Minjoon Seo, Eunsol Choi, and Luke Zettlemoyer. 2017. [Zero-shot relation extraction via reading comprehension](#). *ArXiv*, abs/1706.04115.
- Sheng-Chieh Lin, Akari Asai, Minghan Li, Barlas Oguz, Jimmy Lin, Yashar Mehdad, Wen-tau Yih, and Xilun Chen. 2023a. [How to train your DRAGON: Diverse augmentation towards generalizable dense retrieval](#). *arXiv preprint arXiv:2302.07452*.
- Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Rich James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvasy, Mike Lewis, Luke Zettlemoyer, and Scott Yih. 2023b. [Ra-dit: Retrieval-augmented dual instruction tuning](#). *ArXiv*, abs/2310.01352.
- OpenAI. 2023a. [Gpt-4 technical report](#). *ArXiv*, abs/2303.08774.
- OpenAI. 2023b. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vassilis Plachouras, Tim Rocktaschel, and Sebastian Riedel. 2020. [Kilt: a benchmark for](#)

- knowledge intensive language tasks. In *North American Chapter of the Association for Computational Linguistics*.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. [Direct preference optimization: Your language model is secretly a reward model](#). *ArXiv*, abs/2305.18290.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don't know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023a. [In-context retrieval-augmented language models](#). *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023b. [In-context retrieval-augmented language models](#). *Preprint*, arXiv:2302.00083.
- Siva Reddy, Danqi Chen, and Christopher D. Manning. 2018. [Coqa: A conversational question answering challenge](#). *Transactions of the Association for Computational Linguistics*, 7:249–266.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023a. [Replug: Retrieval-augmented black-box language models](#). *ArXiv*, abs/2301.12652.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023b. [Replug: Retrieval-augmented black-box language models](#). *Preprint*, arXiv:2301.12652.
- Oyvind Tafjord, Peter Clark, Matt Gardner, Wen tau Yih, and Ashish Sabharwal. 2018. [Quarel: A dataset and models for answering questions about qualitative relationships](#). In *AAAI Conference on Artificial Intelligence*.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. [Commonsenseqa: A question answering challenge targeting commonsense knowledge](#). *ArXiv*, abs/1811.00937.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. [Llama: Open and efficient foundation language models](#). *ArXiv*, abs/2302.13971.
- Lewis Tunstall, Edward Beeching, Nathan Lambert, Nazneen Rajani, Kashif Rasul, Younes Belkada, Shengyi Huang, Leandro von Werra, Clémentine Fourrier, Nathan Habib, Nathan Sarrazin, Omar Sanseviero, Alexander M. Rush, and Thomas Wolf. 2023. [Zephyr: Direct distillation of lm alignment](#). *Preprint*, arXiv:2310.16944.
- Boxin Wang, Wei Ping, Peng Xu, Lawrence C. McAfee, Zihan Liu, Mohammad Shoeybi, Yi Dong, Oleksii Kuchaiev, Bo Li, Chaowei Xiao, Anima Anandkumar, and Bryan Catanzaro. 2023a. [Shall we pretrain autoregressive language models with retrieval? a comprehensive study](#). *ArXiv*, abs/2304.06762.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2022. [Self-instruct: Aligning language models with self-generated instructions](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Zhiruo Wang, Jun Araki, Zhengbao Jiang, Md. Rizwan Parvez, and Graham Neubig. 2023b. [Learning to filter context for retrieval-augmented generation](#). *ArXiv*, abs/2311.08377.
- Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023a. [Recomp: Improving retrieval-augmented lms with compression and selective augmentation](#). *ArXiv*, abs/2310.04408.
- Peng Xu, Wei Ping, Xianchao Wu, Lawrence C. McAfee, Chen Zhu, Zihan Liu, Sandeep Subramanian, Evelina Bakhturina, Mohammad Shoeybi, and Bryan Catanzaro. 2023b. [Retrieval meets long context large language models](#). *ArXiv*, abs/2310.03025.
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. [WikiQA: A challenge dataset for open-domain question answering](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018, Lisbon, Portugal. Association for Computational Linguistics.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [Hotpotqa: A dataset for diverse, explainable multi-hop question answering](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2024. [Making retrieval-augmented language models robust to irrelevant context](#). *Preprint*, arXiv:2310.01558.
- Wenhao Yu, Hongming Zhang, Xiaoman Pan, Kaixin Ma, Hongwei Wang, and Dong Yu. 2023. [Chain-of-note: Enhancing robustness in retrieval-augmented language models](#). *Preprint*, arXiv:2311.09210.
- Kun Zhu, Xiaocheng Feng, Xiyuan Du, Yuxuan Gu, Weijiang Yu, Haotian Wang, Qianglong Chen, Zheng Chu, Jingchang Chen, and Bing Qin. 2024. [An information bottleneck perspective for effective noise](#)

A Prompts

This section introduces the prompts mentioned in this paper. Figures 3, 4, and 5 illustrate the prompt for GPT-4, prompts for calculating supportiveness, and the modified prompt for GPT-4, which are discussed in the section Analysis of Recognizing Irrelevant Knowledge.

```
Q:
<question>

Background:
<background>

There is a question Q and a text about
question Q (the Background section).
Please summarize the Background section,
removing any irrelevant information and
keeping content that assists in answering Q.
Please only output the summary.
```

Figure 3: Prompt for GPT-4.

```
Prompt without retrieval:
Answer the following question:
Q:
<question>
A:

Prompt with retrieval:
Answer the following question:
Background:
<background>

Q:
<question>
A:
```

Figure 4: Prompts for calculating supportiveness.

```
Q:
<question>

Background:
<background>

There is a question Q and a text about
question Q (the Background section).
Please summarize the Background section,
removing any irrelevant information and
keeping content that assists in answering Q.
If you believe the content helps answer Q,
please provide the summary, only the
summary. If you think these contents do not
contain information helpful in answering Q,
please output 'irrelevant'.
```

Figure 5: Modified prompt for GPT-4.

B Datasets

We utilized three tasks encompassing a total of 10 datasets to train SKR. “#Size” denotes the number of samples in the original training set of the dataset, while “#Train” represents the number of training set samples used for training SKR after random sampling. The detailed statistics of these datasets can be found in Table 1.

C Supplement to Main Results

Due to space constraints, we have placed the experimental results using Llama-2-13B and Mistral-7B as the base models in Appendix C, Table 2. All settings are consistent with those described in the main results section of the paper.

D Impact of Retriever Selection

Furthermore, to investigate the generalizability of SKR, we also employed another popular method, Contriever (Izacard et al., 2022a), as the retriever in our evaluation. The usage of Contriever slightly differs from DRAGON+. Specifically, Contriever uses a shared encoder \mathbf{E} to encode both query and knowledge chunks, the similarity score is calculated as:

$$s(q, c) = \mathbf{E}(c) \cdot \mathbf{E}(q). \quad (5)$$

The experimental results are shown in Table 8. We observed that due to Contriever’s weaker performance compared to DRAGON+, the average performance of Llama-2-7B with top-5 retrieved knowledge dropped by -1.56 EM on the two datasets. Both SKR and GPT-4 can successfully rewrite knowledge under different retriever settings, and SKR still maintains its performance lead over GPT-4 (+1.93 EM). This confirms the robustness of SKR to different retrievers.

E Effects on Feedback Providers

Table 9 resents the performance of SKR across four tasks when utilizing different feedback providers. In this experiment, we consistently employ the Llama-2-7B model for answer prediction, with all other settings aligned with those described in the section Experiment Results. Our observations indicate that the variations in performance due to different feedback providers are negligible, affirming the robustness of the proposed supportiveness evaluation methodology.

Task	Dataset name	#Size	#Train
Open-Domain QA	CommonsenseQA (Talmor et al., 2019)	9,741	9,741
	Yahoo! Answers QA	87,362	10,000
	Wiki Question Answering (Yang et al., 2015)	20,360	10,000
	FreebaseQA (Jiang et al., 2019)	20,358	10,000
Reading Comprehension	Conversational Question Answering (Reddy et al., 2018)	108,647	10,000
	Discrete Reasoning Over Paragraphs (Dua et al., 2019)	77,400	10,000
	QuaRel (Tafjord et al., 2018)	1,941	1,941
	SQuAD v2 (Rajpurkar et al., 2018)	130,319	10,000
Reasoning	Grade School Math 8K (Cobbe et al., 2021)	7,473	7,473
	Explanations for CommonsenseQ (Aggarwal et al., 2021)	7,598	7,598

Table 6: Datasets used for training SKR.

Model	Rewriter	Retrieval	NQ	TQA	HoPo	MMLU	zsRE	WQ	Avg.	Δ_{EM}
Mistral-7B	-	-	25.73	59.16	21.28	58.92	19.66	21.8	34.43	-
	-	Top-1	28.25	62.54	30.18	59.41	54.43	21.16	42.66	-
	-	Top-5	26.62	60.42	27.17	60.91	30.61	17.32	37.18	-
	GPT-3.5	Top-5	28.99	61.06	30.67	61.23	53.27	20.32	42.59	5.41
	GPT-4	Top-5	32.04	62.8	33.39	61.88	56.46	20.08	44.44	7.27
	RECOMP	Top-5	28.02	60.99	30.63	-	-	-	-	-
	FILCO†	Top-5	31.01	62.67	-	-	59.24	19.42	-	-
	LongLLMLingua	Top-5	27.74	61.44	30.1	60.47	59.11	20.19	43.18	-
	SKR-7B	Top-5	33.25	64.92	32.43	61.99	60.78	24.89	46.38	9.20
	Llama-2-13B	-	-	26.5	61.02	21.63	49.15	23.2	24.05	34.26
-		Top-1	31.83	64.32	31.35	50.92	63.21	25.14	44.46	-
-		Top-5	26.12	62.32	29.51	52.32	61.98	19.63	41.98	-
GPT-3.5		Top-5	35.23	66.12	29.44	52.91	64.28	27.22	45.87	3.89
GPT-4		Top-5	36.66	67.19	34.3	52.94	65.64	28.73	47.58	5.6
RECOMP		Top-5	34.17	66.08	29.25	-	-	-	-	-
FILCO†		Top-5	35.42	65.43	32.03	-	64.17	27.3	-	-
LongLLMLingua		Top-5	34.19	64.42	28.85	50.76	64.12	27.19	44.92	-
SKR-7B		Top-5	38.11	67.21	35.07	54.29	65.98	29.71	48.40	6.41

Table 7: Supplement to Main Results. “†” denotes our implementations.

Retriever	Rewriter	NQ	WQ	Avg.
DRAGON+	-	17.5	14.91	16.21
	GPT-4	29.82	22.83	26.33
	SKR	31.42	25.24	28.33
Contriever	-	14.8	14.13	14.47
	GPT-4	27.12	22.56	24.84
	SKR	28.74	24.79	26.77

Table 8: Impact of Retriever Selection. “Retriever” denotes the use of different models as the retriever. “Avg.” denotes the average performance across the two datasets. In this experiment, we use the top-5 knowledge pieces as retrieval data and evaluate EM using Llama-2-7B.

Feedback	NQ	TQA	HoPo	WQ	Avg.
Mistral-7B	31.42	63.62	30.98	25.24	37.82
Llama-2-7B	31.44	63.58	30.86	25.11	37.75
Llama-2-13B	31.37	63.44	31.04	25.17	37.76

Table 9: Effects on Feedback Providers.