

Eliciting Critical Reasoning in Retrieval-Augmented Generation via Contrastive Explanations

Leonardo Ranaldi^(†) Marco Valentino^(†) André Freitas^(†,*,‡)

[†] Idiap Research Institute, Martigny, Switzerland

^{*}Department of Computer Science, University of Manchester, UK

[‡]National Biomarker Centre (NBC), CRUK Manchester Institute, UK

[name]. [surname]@idiap.ch

Abstract

Retrieval-augmented generation (RAG) has emerged as a critical mechanism in contemporary NLP to support Large Language Models (LLMs) in systematically accessing richer factual context. However, the integration of RAG mechanisms brings its inherent challenges, as LLMs need to deal with potentially noisy contexts. Recent studies have shown that LLMs still struggle to critically analyse RAG-based in-context information, a limitation that may lead to incorrect inferences and hallucinations. In this paper, we investigate how to elicit critical reasoning in RAG via *contrastive explanations*. In particular, we propose *Contrastive-RAG* (C-RAG), a framework that (i) retrieves relevant documents given a query, (ii) selects and exemplifies relevant passages, and (iii) generates explanations that explicitly contrast the relevance of the passages to (iv) support the final answer. We show the impact of C-RAG building contrastive reasoning demonstrations from LLMs to instruct smaller models for retrieval-augmented tasks. Extensive experiments demonstrate that C-RAG improves state-of-the-art RAG models while (a) requiring significantly fewer prompts and demonstrations and (b) being robust to perturbations in the retrieved documents.

1 Introduction

Retrieval-augmented generation (RAG) aims to improve the factuality and memory access of Large Language Models (LLMs) by systematically integrating relevant knowledge from external sources via retrieval mechanisms (Lewis et al., 2020). In particular, RAG is designed to mitigate some of the well-known limitations of LLMs, including the tendency for hallucinations and the lack of specific domain knowledge in the training data (Siriwardhana et al., 2023; Zhang et al., 2023; Kandpal et al., 2023).

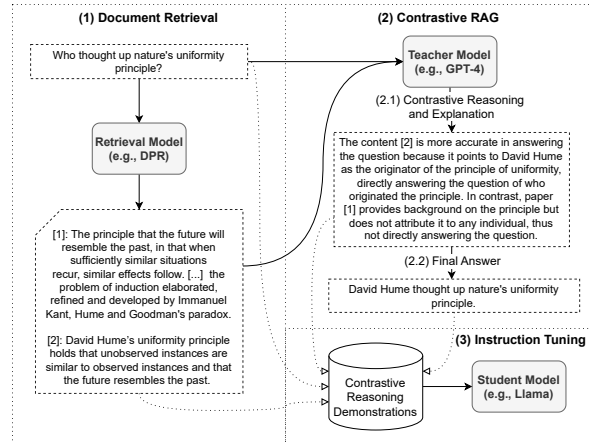


Figure 1: The overall pipeline of Contrastive-RAG (§3). (i) A set of retrieved documents is provided as a context to an LLM that (ii) generates contrastive explanatory arguments to arrive at the final answer. (iii) The explanations generated by a teacher model are used as demonstrations to improve smaller student models.

Despite the benefits of RAG, contemporary studies have identified a range of persisting limitations emerging from noisy retrieval, where irrelevant or contradictory in-context passages can introduce biases in the models, particularly on smaller LMs (Petroni et al., 2020; Shi et al., 2023). These shortcomings stem from the inability of RAG to systematically and critically assess the retrieved passages provided as context. While an emerging line of research attempted to improve the RAG pipeline by incorporating multi-step reasoning strategies to determine the relevance of in-context passages (Li et al., 2023; Yoran et al., 2024), this usually comes at the cost of significantly increasing the computational resources required for training and inference (Xia et al., 2024).

To tackle such limitations, this paper proposes *Contrastive-RAG* (C-RAG), a novel end-to-end framework designed to elicit a critical reasoning process in retrieval-augmented language models in the form of *contrastive explanations* – i.e., ex-

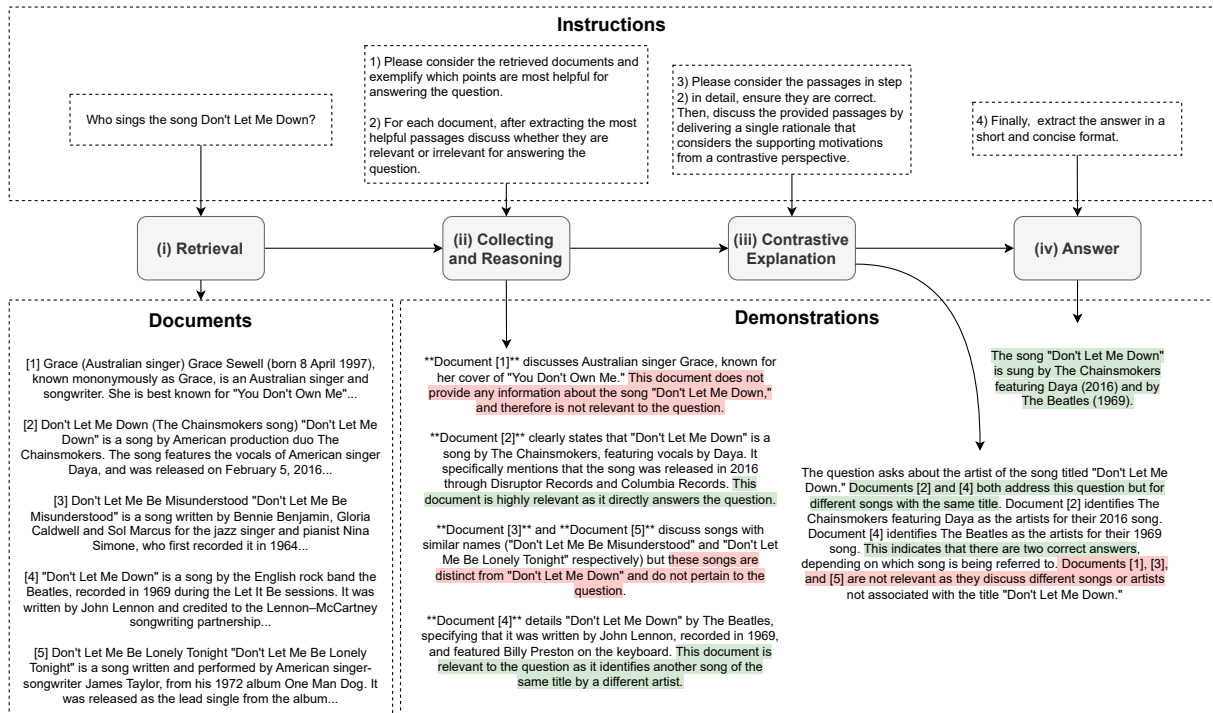


Figure 2: An illustration of the C-RAG framework for improving the traditional RAG pipeline. We demonstrate that C-RAG can significantly improve the performance of RAG models while requiring fewer prompts, training examples, and annotation steps than state-of-the-art approaches.

planations that explicitly compare and contrast the relevance of the retrieved passages with respect to the task to be addressed. Specifically, the aim of C-RAG is to improve the original RAG pipeline via a multi-step reasoning framework composed of the following steps: *i) collecting step*, where query and documents are analysed to extract passages that are relevant for the query; *ii) contrastive reasoning*, where the LLMs construct explanatory arguments about the relevance of the extracted passages, highlighting and contrasting *relevant* and *irrelevant* knowledge; *iii) explanation*, where the contrastive arguments are consolidated into a single final explanation; and *iv) answering*, where a short-form answer is generated to address the query.

We demonstrate the impact of C-RAG by building reasoning demonstrations from LLMs and instructing smaller models to address retrieval-augmented tasks via contrastive explanations. An extensive empirical evaluation on four public question-answering (QA) tasks led to the following findings and conclusions:

1. Employing C-RAG to generate contrastive reasoning demonstrations via large reasoning models can significantly improve the performance of smaller RAG models leading to an average increase in accuracy of 55.4% over

RAG and 7.2% and 1.9% over Self-RAG (Asai et al., 2023) and Self-Reasoning (Xia et al., 2024) respectively.

2. We found that C-RAG is significantly more efficient than state-of-the-art RAG frameworks requiring fewer prompts, annotation steps, and training examples (i.e., 2k vs 190k required by Self-RAG) while achieving better performance.
3. We demonstrate that C-RAG is robust to perturbations that are typically challenging for traditional RAG models, including the random shuffling of the retrieved documents and random noise applied to in-context passages.

To the best of our knowledge, this is the first work to investigate the impact of contrastive explanations on RAG and demonstrate how contrastive reasoning demonstrations can consistently boost the performance of smaller LLMs, equipping them with the ability to critically analyse external knowledge and generate contrastive explanatory reasoning for their predictions.

2 Contrastive Explanations

Contrastive explanations have been identified as fundamental modes of explanation in epistemology, artificial intelligence, and cognitive science (Lipton, 1990; Miller, 2019; Valentino and Freitas, 2024a,b). A contrastive explanation is a type of inference aimed at answering why-questions of the form “*Why P rather than Q?*”, where P is the explanandum – i.e., the *fact* to be explained – and Q is a counterfactual event – i.e., the *foil*. The function of a contrastive explanation is to describe what *makes the difference* between the occurrence of the *fact* and the *foil*.

In this paper, we aim to leverage the notion of contrastive explanation to elicit critical reasoning in RAG. Our intuition is that contrastive explanations are a particularly fitting systematic reasoning mechanism for determining the *relevance* of documents and passages provided as a context within the RAG pipeline.

Formally, given a query q and a set of documents $D = \{d_1, d_2, \dots, d_n\}$, we want to partition D into a set of relevant documents $P_q \in D$, and a set of irrelevant documents, $Q_q \in D$. In particular, we want these sets to be constructed through the generation of a natural language explanation E that describes the factors that contribute to the relevance of P_q and the factors that contribute to the irrelevance of Q_q . In the context of RAG, therefore, E aims to answer the question “*Why is P_q relevant to q rather than Q_q ?*”, where the set of relevant documents P_q corresponds to the *fact*, and the set of irrelevant documents Q_q corresponds to the *foil*.

3 Contrastive-RAG

To elicit contrastive explanations in RAG, we present a framework composed of four inference stages (Figure 2): i) *collecting step* (§3.1), where, given a query, the retrieved documents are collected and an LLM-based model extracts relevant passages; ii) *contrastive reasoning* (§3.2), where an LLM-based model generates critical and contrastive explanations about the relevance of the extracted passages, by exemplifying and contrasting *relevant* and *irrelevant* aspects; iii) *explanation* (§3.3), where the arguments constructed in *ii*) are summarised into a single explanation; iv) *answering* (§3.3), where a final short answer to the query is generated. The prompt adopted for C-RAG is illustrated in Table 6. In this paper, we are interested in employing C-RAG to generate synthetic reason-

ing demonstrations (§3.4) to enhance RAG models (§3.5) and to equip them with the capability of critically analysing the questions and the retrieved documents and generate a contrastive explanation to arrive at the answer (Figure 2).

3.1 Collecting Step

RAG models leverage retrieved knowledge to improve accuracy and reduce hallucinations in LLMs. Therefore, the first step in the proposed pipeline involves retrieving relevant documents from a given reference document base \mathcal{D} . In this paper, we use DPR (Karpukhin et al., 2020) and Contriever (Izacard et al., 2021) as the default retriever models R . Subsequently, we instruct the LLM to analyse the question and extract the *most relevant* passages from the set of retrieved documents (i.e., “**#Reference Evidence**”). We collect the retrieved documents and relevant passages for all the queries and refer to this step as α_1 .

3.2 Contrastive Reasoning

Recent work has shown that exemplifying passages from the retrieved documents and directly using them as in-context knowledge for RAG can improve LLMs’ accuracy (Asai et al., 2023). However, each passage may still contain irrelevant or contradictory knowledge that can mislead the model. Hence, after collecting the passages from the top- k documents (§3.1), we instruct the LLM to generate contrastive explanatory arguments to identify and compare relevant and irrelevant points in the passages with respect to the question (§3.1). We perform this step by setting out the instructions as reported in Table 6. Hence, we collect the outcomes by defining this phase as α_2 .

3.3 Explanation & Answer

Finally, we leverage the arguments in the previous steps to generate a final contrastive explanation to derive the answer. In particular, we instruct the LLM to explicitly consider the contrastive rationales and summarise the main points into a single explanation. We define this step as α_3 . Subsequently, we instruct the model to generate the final answer following the pattern “**#Answer:**”. This final step is defined as α_4 .

3.4 C-RAG Operability

C-RAG leverages a set of structured instructions to deliver multi-step explanations via contrastive

reasoning. Thus, the operability of C-RAG is two-fold, as it can be employed as both a prompting strategy and a synthetic annotation technique.

3.4.1 C-RAG as a Prompting Strategy

By operating through the instructions in Table 6, we adopt C-RAG to prompt GPT-4 (OpenAI, 2023). Specifically, we instruct GPT-4 to extract the most crucial passages from the retrieved documents (§3.1), explaining the relevant and irrelevant points to answer the given question by exemplifying the main passages (§3.2), provide a single exhaustive explanation that best describes the critical points (§3.3); and finally, generate the final answer in a strict format, in order to have a more straightforward and less ambiguous downstream assessment. However, although the sequence of instructions is well-structured and defined, the ability to perform sequential and complex reasoning tasks is limited to larger LLMs (such as GPT-4, as discussed in the experiments). Therefore, to transfer contrastive reasoning to smaller models, we use C-RAG to generate synthetic annotations as training demonstrations.

3.4.2 C-RAG as a Demonstration Strategy

Following recent work (Xia et al., 2024; Asai et al., 2023), we instruct smaller models via demonstrations produced by high-performing LLMs that are capable of following structured instructions. In contrast to previous methods, however, our approach is based on a single prompt composed of a series of sequential instructions (Figure 2). Although GPT-4 have demonstrated the ability to follow sequential instructions (Peng et al., 2023), we cannot formally guarantee that the generated demonstrations are correct. Therefore, we follow the method proposed by Xia et al. (2024), which computes the citation precision for the considered documents as a proxy for the quality of the demonstrations. However, since C-RAG uses a different annotation mechanism, our heuristics firstly filter out the final correct answers through a strict, exact match; then, after the filtering (cutting off about 50% of the demonstrations), it verifies that each retrieved document along the reference evidence has been considered. The starting annotations consist of approximately 10,000 training samples delivered with GPT-4, which, after the first filtering strategy, are reduced to 4,500 and finally, through quality control, are reduced to 2,000 (see Appendix C).

3.5 Training

Thanks to the operability of C-RAG (3.4), a Language Model π can be trained using the generated annotations¹, which are augmented with reasoning demonstrations α using the standard language modeling objective, maximizing likelihood:

$$\max_{\pi} \mathbb{E}_{(q, \alpha, y) \sim \mathcal{D}_A} \log p_{\pi}(y | \alpha, q) p_{\pi}(\alpha | q) \quad (1)$$

where $\alpha = \alpha_1 \oplus \alpha_2 \oplus \alpha_3 \oplus \alpha_4$ is the combination of the multiple reasoning steps performed by the model, \oplus is the concatenation operator, and α_1, α_2 , are the respective annotations generated by the above processes. q is the provided question, and y is the model output, including the intermediate steps and the final answer. \mathcal{D}_A is the training corpus augmented with contrastive reasoning demonstrations.

4 Experiments

We evaluate C-RAG on four open-domain question-answering tasks (§4.1). We perform the retrieval and evaluation phases by following standard approaches used to assess the RAG pipeline (§4.2) and perform the tuning phase by using the setup presented in §4.3.

4.1 Tasks & Datasets

We conduct an extensive experimental evaluation on the following question-answering (QA) tasks: (i) NaturalQuestion (NQ) (Kwiatkowski et al., 2019), (ii) PopQA (Mallen et al., 2023), (iii) TriviaQA (Joshi et al., 2017) and (iv) FEVER (Thorne et al., 2018). Appendix D describes the composition of each dataset.

4.2 Experimental Setup

Retriever We use DPR (Karpukhin et al., 2020) and Contriever-MS MARCO (Izacard et al., 2021) to retrieve the top top- k documents from the document base. We chose $k = 5$ to have a fair comparison to related RAG approaches using the same value for k . By default, we operate via DPR on NQ, as DPR has been fine-tuned on the dataset. On PopQA, where question and answer pairs are created based on Wikipedia, we use the Wikipedia_corpus as background knowledge as proposed in (Xia et al., 2024).

¹we select annotations as described in Section 3.4.2

Models	NQ	PopQA	TriviaQA	FEVER	Train data for LLM	
					Training Size	Annotation
Baseline (no RAG)						
Llama-2-7b	19.2	18.4	30.5	20.1	-	-
Llama-2-13b	24.0	22.6	38.5	25.2	-	-
C-RAG-7B	30.0	44.8	58.6	32.5	-	-
C-RAG-13B	31.8	46.2	61.3	34.0	-	-
GPT-4-o	35.2	52.4	64.3	36.5	-	-
RAG						
Llama-2-7b	27.8	47.8	55.6	23.2	-	-
Llama-2-13b	34.0	48.1	59.2	25.3	-	-
GPT-4-o	46.6	62.5	74.6	87.7	-	-
Llama-2-7b (C-RAG)	24.6	47.0	54.8	22.9	-	-
Llama-2-13b (C-RAG)	33.5	47.4	58.0	24.9	-	-
GPT-4-o (C-RAG)	49.4	64.8	76.4	90.3	-	-
RAG + Tuning (Llama-2-7b, -13b)						
Llama-2-7b (SFT)	36.8	54.4	61.9	67.5	2k	single-step
RECOMP	38.4	-	-	39.1	150k	external
Self-RAG-7b	37.2	54.9	66.4	70.2	190k	external
Self-RAG-13b	38.8	55.8	67.2	72.2	190k	external
Self-Reasoning-7b	38.0	54.2	-	78.6	2k	double-step
Self-Reasoning-13b	41.4	57.3	-	83.9	2k	double-step
C-RAG-7B	40.2	56.4	68.4	79.2	2k	single-step
C-RAG-13B	42.6	58.2	70.3	83.6	2k	single-step

Table 1: Overall results on QA and fact verification tasks §4.1. The models are prompted as detailed in §4.4, and the values correspond to the Exact Match (%).

Evaluation Metrics We adopt two different evaluation metrics for QA and fact-verification tasks. Specifically with regard to QA tasks (NQ, PopQA, and TriviaQA), we use a flexible exact-match accuracy following Schick et al. (2023); Mallen et al. (2023), which is based on whether or not ground-truth answers are included in the generated answers provided by the models, instead of a strict exact match. For fact verification tasks, i.e., FEVER, we report the evaluation scheme proposed in (Thorne et al., 2018) based on a three-class classification accuracy. Finally, C-RAG uses a special label ‘#Answer’ (see Table 6) through which we instruct the models to deliver a short answer.

4.3 Models Setup

To get a comprehensive evaluation of existing RAG pipelines and the impact of C-RAG, three different LLMs are used: GPT-4 (OpenAI, 2023), Llama-2-7 and Llama-2-13 (Touvron et al., 2023) along with their instruction-tuned chat version Llama-2-7-chat and Llama-2-13-chat. The models were chosen to have a common ground for comparison with state-of-the-art approaches.

Inference Settings We use greedy decoding in all experiments to ensure a more deterministic genera-

tion process. By using the prompt shown in Table 5 we set the temperature to 0.4 and maximum generation length of 2048, as we observed that these settings deliver better overall performance.

4.3.1 Training Setup

To evaluate the impact of C-RAG contrastive reasoning demonstrations on smaller models (§3), we use the annotations produced following the C-RAG strategy (§3.4.2). Additionally, for a fair comparison, we produce annotations using Llama-2-SFT, where Llama-2 is fine-tuned on training samples without C-RAG. We compare our models with several related RAG approaches trained with demonstrations generated by GPT-4 to establish strong baselines (detailed in the last column of Table 1). We fine-tune the Llama-2 models for 3 epochs with a batch size of 32 and a learning rate equal to $3e-5$ with a 0.001 weight decay. We use the cosine learning rate scheduler with a warmup ratio of 0.03. We conducted our experiments on a workstation equipped with four Nvidia RTX A6000 with 48GB of VRAM.

4.4 Prompting

We systematically prompt the models using two main settings:

Models	NQ	PQA	TQA
	(%acc)	(%acc)	(%acc)
FULL	42.6	58.2	70.3
RANDOM	38.4	54.2	61.0
w/o (2)	33.2	52.2	60.6
w/o (3)	37.0	55.7	62.3
w/o (4)	40.8	57.4	69.4

Table 2: Evaluation of impacts of each component on three QA tasks with C-RAG-13B. We eliminate (w/o) or RANDOM shuffling the four defined steps (§3). C-RAG-7Bs’ in Table 8.

Baseline (no RAG) We evaluate the baseline capabilities of selected models in a zero-shot manner and without introducing any in-context documents (without RAG) as in (Asai et al., 2023; Xia et al., 2024) using the prompt in Table 4.

RAG Models We assess the impact of retrieved knowledge by instructing the evaluated models to consider the top-5 retrieved documents. In line with (Xia et al., 2024), we use retrievers described in §4.2 as in Table 5. To complete the settings, we use C-RAG as a prompting strategy as in Table 6.

5 Results & Discussions

The results of the empirical evaluation are reported in Table 1. Overall, the experiments confirm that C-RAG can improve the capabilities of LLMs to handle retrieved documents for QA and fact verification tasks demonstrating the impact of contrastive reasoning and explanations on RAG models. We found that C-RAG is particularly effective as a demonstration strategy to improve the performance of smaller Llama-2 models, achieving state-of-the-art performance when compared to related fine-tuning approaches in the literature (Xu et al., 2023; Asai et al., 2023; Xia et al., 2024). In the following sections, we analyse the impact of C-RAG when adopted as both a prompting strategy (§5.1) and as a framework for generating annotations to instruct LLMs (§5.2). Finally, we analyse the role of the contrastive explanations (§5.3) and provide evidence of robustness on challenging perturbations and low-resource settings (§5.4).

5.1 C-RAG as a Prompting Strategy

The middle part of Table 1 reports the results of C-RAG when adopted as a prompting strategy for different models. While we can observe an overall improvement over the baseline models without

RAG (with an improvement of 68.9% for GPT-4, 56.6% for Llama-2-7b and 65.4% for Llama-2-13b), the results show that the impact of C-RAG as a prompting strategy in a RAG setting is only evident for larger models (i.e., GPT-4-o) where C-RAG achieves an overall improvement of 3.9%. For Llama-2-7b and Llama-2-13b, in fact, we observe a decrease in performance when compared to the standard RAG pipeline, indicating that such models are unable to generate the contrastive reasoning required to support their answers.

5.2 C-RAG as a Demonstration Strategy

The lower part of Table 1 reports the results of C-RAG when adopted as a demonstration strategy for different models. Here, the results show that C-RAG is highly effective in improving the performance of Llama-2 models when used to generate reasoning demonstrations via GPT-4. In particular, we found that C-RAG can outperform state-of-the-art approaches, including RECOMP (+1.8%), Self-RAG (+7.2%), and Self-Reasoning (+1.9%). Moreover, as shown in Table 1, C-RAG can achieve such results using a fraction of the reasoning demonstrations used by related approaches, also requiring fewer prompts and annotation steps. Specifically, our approach uses significantly fewer demonstrations when compared to RECOMP (Xu et al., 2023) and Self-RAG (Asai et al., 2023), (2k demonstrations vs. 150k and 190k). Similarly, in contrast to Self-Reasoning (Xia et al., 2024), C-RAG operates via a single-step prompting. This experimental setting significantly reduces the use of resources and simplifies the training process.

In addition, we observe that Llama-2 models fine-tuned via C-RAG can outperform GPT4-o by 6.1% and substantially improve the performance of all the evaluated Llama-2 models (with and without RAG). These results strongly demonstrate the impact of the training signal provided by contrastive explanations and their ability to efficiently elicit critical arguments in smaller LMs (as shown in an inference example in Appendix G).

5.3 The Role of Contrastive Explanations

Table 2 evaluates the impact of the end-to-end contrastive reasoning framework on the final performance. In particular, the table shows the effect of eliminating one of the steps or randomly shuffling them to produce the demonstrations.

The results demonstrate the importance of each stage in the multi-step reasoning process intro-

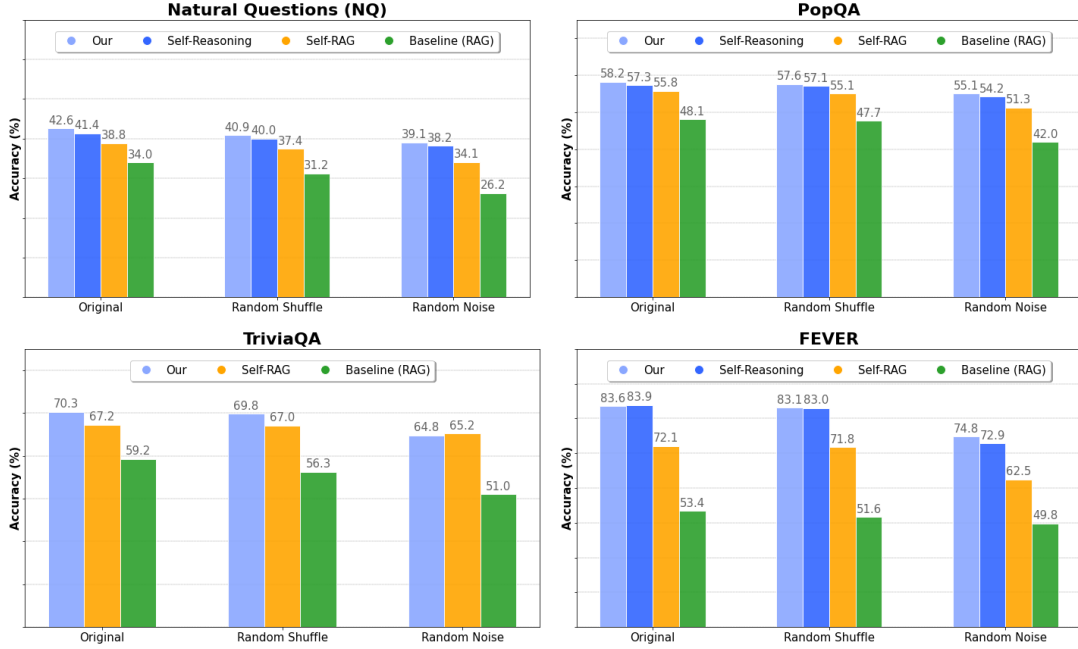


Figure 3: Robustness experiment results on four QA datasets (§4.1) using the same evaluation settings proposed for C-RAG-13B in Table 1. We provide retrieved documents by randomly shuffling them (Random Shuffle) and using 50% of irrelevant (Random Noise).

duced in §3. In particular, we observe the highest decrease in performance when removing step 2 (i.e., Contrastive Reasoning, with a decrease of -17.2%) and step 3 (i.e., Explanation, with a decrease of -10.4%), demonstrating the crucial impact of the contrastive explanatory arguments on the final performance.

5.4 Robustness & Ablation Analysis

The C-RAG framework instructs models to reason on the retrieved documents independently of their order of occurrence (i.e., invariance to documents’ permutations) and attempts to elicit critical explanations to identify irrelevant or contradictory knowledge in the extracted passages (i.e., robustness to the bias) as revealed in Figure 3. In addition, the benefits of C-RAG also emerge when reducing the number of demonstrations, showing that contrastive explanations can improve data efficiency (Figure 4). To assess such properties in more detail, we performed a robustness analysis along with an evaluation of how scaling the training demonstrations affects models’ behaviours.

Robustness to Perturbations Since noisy retrieval can negatively affect the performance of LLMs (Petroni et al., 2020), we follow the methodology introduced in (Asai et al., 2023; Xia et al., 2024) to evaluate robustness. Specifically, we shuf-

fled the order of the retrieved documents (*Random Shuffle*) and inserted two irrelevant documents (*Random Noise*). Figure 3 reports the experimental results. We found that the C-RAG framework consistently outperforms the baseline model (Llama-2-13b), as well as Self-RAG and Self-Reasoning, finding that perturbations have a lower impact on the final performance. In particular, the random shuffling of retrieved documents has a minimal impact on performance, demonstrating the permutation invariance property of C-RAG. Moreover, when noisy documents are added, all the evaluated models suffer a higher performance drop. However, the drop for C-RAG is typically lower than the RAG baseline, which shows that the proposed method is more robust even when dealing with noisier results.

Quantity of Instructions Figure 4 shows the behaviour of C-RAG when scaling-up the number of training examples. While we found that the quantity of the demonstrations used in C-RAG is important in determining the final performance, we found that C-RAG can outperform the baselines RAG models with only 50% of training demonstrations, also achieving superior training performance when compared to the fine-tuned SFT model (i.e., the model fine-tuned without contrastive reasoning demonstrations as explained in §4). This further

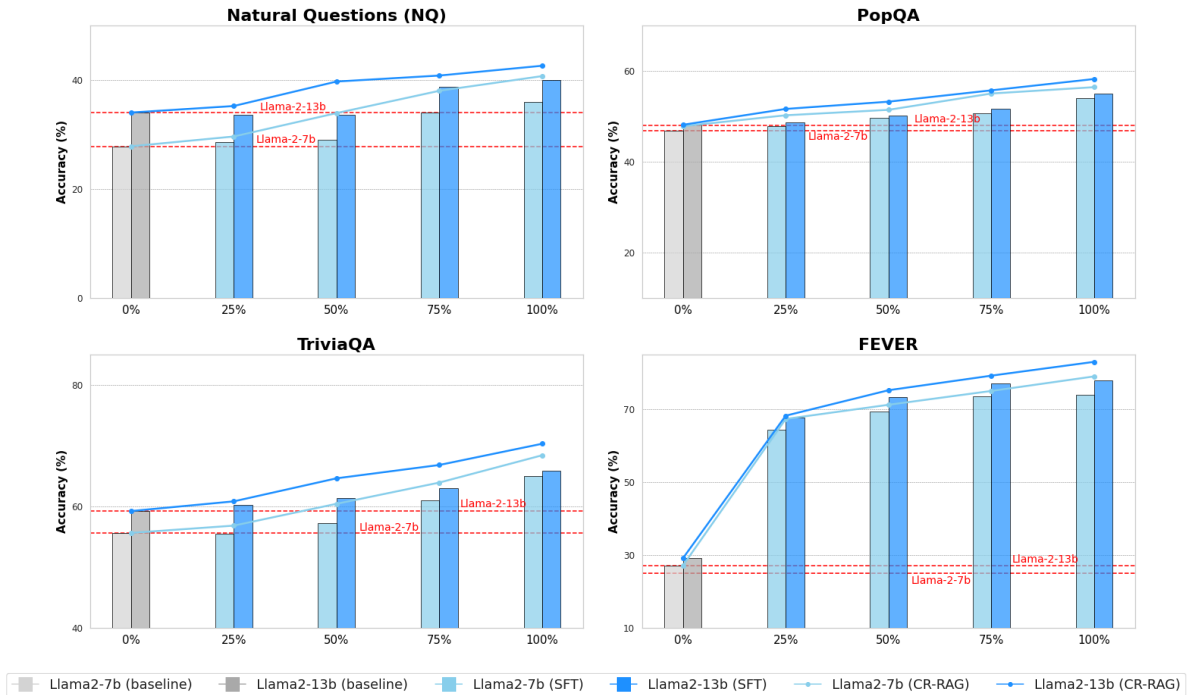


Figure 4: Performances assessment of C-RAG-7B and -13B by scaling training data. We replicated experimental settings proposed in §4.3 changing the number of tuning instructions.

Training	Model	Evaluation			
		NQ	PQA	TQA	FEV
	Baseline	19.2	18.4	30.5	20.1
	(+RAG)	27.8	47.8	55.6	23.2
NQ	SFT	36.5	48.0	56.2	41.4
	C-RAG	39.8	52.7	62.6	45.2
PopQA	SFT	31.0	50.2	54.3	40.4
	C-RAG	33.6	55.9	57.4	43.7
TQA	SFT	32.2	49.6	60.7	39.3
	CRAG	34.2	52.2	68.0	45.6
FEVER	SFT	29.4	50.3	52.9	66.6
	C-RAG	34.6	53.0	68.4	78.6

Table 3: By replicating the experimental setting (§4.3), we trained the models (Llama-2-7b) on a single task and systematically evaluated them on other tasks. As SFT, we mean models (Llama-2-7b) instructed with input-output demonstrations consisting of queries, documents, and target answers.

highlights the quality of the training signal provided by the contrastive explanations.

Quality of Instructions To complete the experiment, we assessed the transferability of C-RAG and analysed the impact of the quality of demonstrations using adversarial examples. Concerning transferability, we performed the training on a single task (we use the split reported in Table 7) and evaluated the models on another task. Table 3

shows that models trained through demonstrations can improve the performance of RAG models on tasks they have not been trained on.

Finally, we performed a sanity check to verify the impact of the quality of the demonstrations in Table 9 (in Appendix). Here, we performed adversarial experiments to study the impact of the quality of the demonstrations used to instruct the C-RAG models. We replicated the proposed experimental setting, operating via misleading instructions (defined in Appendix F) to ensure a complete understanding of the impact of the demonstrations. The results (Table 9) show that the quality of the instructions matters. Models instructed with high-quality demonstrations (filtered as detailed in Appendix C) achieve better performance than misleading demonstrations, which can degrade the accuracy below the baselines.

6 Related Work

Previous research investigated the advantages of augmenting Large Language Models (LLMs) through retrieved text passages, a technique known as Retrieval-augmented Generative (RAG) (Lewis et al., 2020; Ram et al., 2023). However, recent work showed that the benefits of RAG can be undermined by noisy retrieval, thus decreasing consistency and reliability (Liu et al., 2023; Petroni

et al., 2020; Shi et al., 2023). Hence, several works proposed techniques to improve RAG reliability by enabling the models to elicit the critical steps to arrive at the final answer (Menick et al., 2022; Gao et al., 2023b). Similarly, recent work focused on improving the retrieval phase by fine-tuning LLMs to perform dynamic retrieval via adaptive reasoning strategies Jiang et al. (2023); Yao et al. (2023); Gao et al. (2023a); Zhang et al. (2024). Nevertheless, the use of multi-step reasoning strategies usually comes at the cost of increasing the computational resources in terms of number of prompts and training examples (Fan et al., 2024; Gao et al., 2024).

For instance, Asai et al. (2023) instructed models to retrieve information using special reflection tokens. However, this solution requires the training of two external models, requiring tens of thousands of additional training samples. Xu et al. (2023) attempted to lower the computational cost for multi-step RAG pipelines, but their approach still requires additional models to summarise the retrieved documents. Finally, Xia et al. (2024) eliminated the dependence on special tokens and external components by introducing reasoning trajectories that are employed to boost the performance of LLMs directly. Although the solution improves efficiency, the framework operates through a multi-step mechanism that requires multiple annotation phases.

Similarly to recent work investigating the impact of natural language explanations on LLMs (He et al., 2024; Dalal et al., 2024; Ye and Durrett, 2022; Quan et al., 2024b,a; Ranaldi and Freitas, 2024b), we propose a method to integrate via multi-step explanations into RAG. To the best of our knowledge, however, this is the first work to investigate the impact of contrastive explanations on RAG and demonstrate how contrastive reasoning demonstrations can boost the performance of smaller LMs.

7 Conclusion & Future Work

RAG has shown great potential in boosting the performance of LLMs on knowledge-intensive tasks. Despite the success of RAG, noisy retrieval represents a major limitation. To tackle such challenges, we introduce *Contrastive-RAG* (C-RAG), a framework that (i) retrieves relevant documents given a query, (ii) selects and exemplifies relevant passages, and (iii) generates explanations that explicitly contrast the relevance of the passages to (iv) support the final answer. We demonstrate that C-RAG can

outperform state-of-the-art models while requiring fewer prompts and demonstrations and being robust to perturbations in the retrieved documents, laying the foundation for integrating natural language explanations within RAG-based architectures and improving efficiency, consistency and reliability.

In future developments, we plan to extend our contribution to non-English languages to broaden the beneficial impacts and operability of C-RAG for multi-lingual alignment (Ranaldi and Freitas, 2024a) and augmentation approaches (Ranaldi et al., 2024). Another possible extension for C-RAG is to investigate the integration of contrastive explanations with different types of explanatory arguments studied in Philosophy of Science (Valentino and Freitas, 2024a; Valentino et al., 2022, 2021). A potential direction, in fact, would be to explore how different types of explanations can be leveraged to further improve reasoning in LLMs and RAG and develop more efficient instruction tuning techniques to improve smaller models using a lower amount of data.

8 Limitations

In this work, we proposed an approach to instruct Large Language Models (LLMs) to improve performance in RAG settings. Specifically, we introduced a framework eliciting LLMs to deliver contrastive reasoning for open-domain question answering and fact verification. Although the framework can handle the distribution of realistic questions, as evaluated standard datasets, it remains to be assessed the impact of this framework on more complex reasoning scenarios, such as multi-hop reasoning, code generation, and arithmetic reasoning. In future work, we will tackle more challenging reasoning tasks. While C-RAG can serve as a mechanism to mitigate factual hallucinations in LLMs and improve the robustness of RAG, there is still a risk that our approach might induce hallucinations.

Acknowledgements

This work was funded by the Swiss National Science Foundation (SNSF) project “NeuMath” (200021_204617), by the Innosuisse project “SINFONIA” (n. 104.170 IP-ICT), by the CRUK National Biomarker Centre, and supported by the Manchester Experimental Cancer Medicine Centre and the NIHR Manchester Biomedical Research Centre.

References

- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#).
- Dhairya Dalal, Marco Valentino, Andre Freitas, and Paul Buitelaar. 2024. [Inference to the best explanation in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 217–235, Bangkok, Thailand. Association for Computational Linguistics.
- Wenqi Fan, Yujian Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. [A survey on rag meeting llms: Towards retrieval-augmented large language models](#). In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '24*, page 6491–6501, New York, NY, USA. Association for Computing Machinery.
- Luyu Gao, Zhuyun Dai, Panupong Pasupat, Anthony Chen, Arun Tejasvi Chaganty, Yicheng Fan, Vincent Zhao, Ni Lao, Hongrae Lee, Da-Cheng Juan, and Kelvin Guu. 2023a. [RARR: Researching and revisiting what language models say, using language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16477–16508, Toronto, Canada. Association for Computational Linguistics.
- Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023b. [Enabling large language models to generate text with citations](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488, Singapore. Association for Computational Linguistics.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-augmented generation for large language models: A survey](#).
- Xuanli He, Yuxiang Wu, Oana-Maria Camburu, Pasquale Minervini, and Pontus Stenetorp. 2024. [Using natural language explanations to improve robustness of in-context learning](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13477–13499, Bangkok, Thailand. Association for Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2021. Unsupervised dense information retrieval with contrastive learning. *arXiv preprint arXiv:2112.09118*.
- Zhengbao Jiang, Frank Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. [Active retrieval augmented generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7969–7992, Singapore. Association for Computational Linguistics.
- 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.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. [Large language models struggle to learn long-tail knowledge](#).
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. [Natural questions: A benchmark for question answering research](#). *Transactions of the Association for Computational Linguistics*, 7:452–466.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Daliang Li, Ankit Singh Rawat, Manzil Zaheer, Xin Wang, Michal Lukasik, Andreas Veit, Felix Yu, and Sanjiv Kumar. 2023. [Large language models with controllable working memory](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1774–1793, Toronto, Canada. Association for Computational Linguistics.
- Peter Lipton. 1990. Contrastive explanation. *Royal Institute of Philosophy Supplements*, 27:247–266.
- Nelson Liu, Tianyi Zhang, and Percy Liang. 2023. [Evaluating verifiability in generative search engines](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7001–7025, Singapore. Association for Computational Linguistics.
- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khatabi, and Hannaneh Hajishirzi. 2023. [When not to trust language models: Investigating effectiveness of parametric and non-parametric memories](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9802–9822, Toronto, Canada. Association for Computational Linguistics.
- Jacob Menick, Maja Trebacz, Vladimir Mikulik, John Aslanides, Francis Song, Martin Chadwick,

- Mia Glaese, Susannah Young, Lucy Campbell-Gillingham, Geoffrey Irving, and Nat McAleese. 2022. [Teaching language models to support answers with verified quotes](#).
- Tim Miller. 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial intelligence*, 267:1–38.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023. [Instruction tuning with gpt-4](#).
- Fabio Petroni, Patrick Lewis, Aleksandra Piktus, Tim Rocktäschel, Yuxiang Wu, Alexander H. Miller, and Sebastian Riedel. 2020. [How context affects language models’ factual predictions](#).
- Xin Quan, Marco Valentino, Louise Dennis, and Andre Freitas. 2024a. [Enhancing ethical explanations of large language models through iterative symbolic refinement](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1–22, St. Julian’s, Malta. Association for Computational Linguistics.
- Xin Quan, Marco Valentino, Louise A. Dennis, and Andre Freitas. 2024b. [Verification and refinement of natural language explanations through LLM-symbolic theorem proving](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2933–2958, Miami, Florida, USA. Association for Computational Linguistics.
- Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. [In-context retrieval-augmented language models](#). *Transactions of the Association for Computational Linguistics*, 11:1316–1331.
- Leonardo Ranaldi and Andre Freitas. 2024a. [Aligning large and small language models via chain-of-thought reasoning](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1812–1827, St. Julian’s, Malta. Association for Computational Linguistics.
- Leonardo Ranaldi and Andre Freitas. 2024b. [Self-refine instruction-tuning for aligning reasoning in language models](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2325–2347, Miami, Florida, USA. Association for Computational Linguistics.
- Leonardo Ranaldi, Giulia Pucci, and Andre Freitas. 2024. [Empowering cross-lingual abilities of instruction-tuned large language models by translation-following demonstrations](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 7961–7973, Bangkok, Thailand. Association for Computational Linguistics.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). *arXiv preprint arXiv:2302.04761*.
- Freda Shi, Xinyun Chen, Kanishka Misra, Nathan Scales, David Dohan, Ed Chi, Nathanael Schärli, and Denny Zhou. 2023. [Large language models can be easily distracted by irrelevant context](#).
- Shamane Siriwardhana, Rivindu Weerasekera, Elliott Wen, Tharindu Kaluarachchi, Rajib Rana, and Suranga Nanayakkara. 2023. [Improving the domain adaptation of retrieval augmented generation \(RAG\) models for open domain question answering](#). *Transactions of the Association for Computational Linguistics*, 11:1–17.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [FEVER: a large-scale dataset for fact extraction and VERification](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 809–819, New Orleans, Louisiana. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Marco Valentino and André Freitas. 2024a. On the nature of explanation: An epistemological-linguistic perspective for explanation-based natural language inference. *Philosophy & Technology*, 37(3):88.
- Marco Valentino and André Freitas. 2024b. [Reasoning with natural language explanations](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 25–31, Miami, Florida, USA. Association for Computational Linguistics.
- Marco Valentino, Mokanarangan Thayaparan, Deborah Ferreira, and André Freitas. 2022. [Hybrid autoregressive inference for scalable multi-hop explanation regeneration](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(10):11403–11411.
- Marco Valentino, Mokanarangan Thayaparan, and André Freitas. 2021. [Unification-based reconstruction of multi-hop explanations for science questions](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 200–211, Online. Association for Computational Linguistics.
- Yuan Xia, Jingbo Zhou, Zhenhui Shi, Jun Chen, and Haifeng Huang. 2024. [Improving retrieval augmented language model with self-reasoning](#).

- Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023. [Re-comp: Improving retrieval-augmented lms with compression and selective augmentation](#).
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. 2023. [React: Synergizing reasoning and acting in language models](#).
- Xi Ye and Greg Durrett. 2022. The unreliability of explanations in few-shot prompting for textual reasoning. *Advances in neural information processing systems*, 35:30378–30392.
- Ori Yoran, Tomer Wolfson, Ori Ram, and Jonathan Berant. 2024. [Making retrieval-augmented language models robust to irrelevant context](#).
- Tianjun Zhang, Shishir G. Patil, Naman Jain, Sheng Shen, Matei Zaharia, Ion Stoica, and Joseph E. Gonzalez. 2024. [Raft: Adapting language model to domain specific rag](#).
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemaoy Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. 2023. [Siren’s song in the ai ocean: A survey on hallucination in large language models](#).

A Prompting Approaches

#Role You are an experienced expert skilled in answering various questions.
#Task Please answer the question following the detailed requirements.
#Requirements Please answer the question based on your knowledge using the format “ #Answer: ”
#Question {question}

Table 4: Baseline prompting example.

#Role You are an experienced expert skilled in answering various questions.
#Task Please answer the question based on the documents provided and following the detailed requirements using the format “ #Answer: ”
#Reference Documents [1] {Document ₁ } [2] {Document ₂ } [3] {Document ₃ } [4] {Document ₄ } [5] {Document ₅ }
#Requirements Please consider the retrieved documents provided “ #Reference Documents ” and answer the question.
#Question {question}

Table 5: Retrieval-augmented Generation prompting example.

B C-RAG prompting Requirements

<p>#Role You are an experienced expert skilled in answering various questions.</p>
<p>#Task Please answer the question based on the documents provided and following the detailed requirements.</p>
<p>#Reference Documents [1] {Document₁} [2] {Document₂} [3] {Document₃} [4] {Document₄} [5] {Document₅}</p>
<p>#Requirements</p> <ol style="list-style-type: none">1) Please consider the retrieved documents provided “#Reference Documents” and understand the main points. Follow the directions in detail and use only the information in the Documents, exemplifying which points are most helpful for answering the question #Question. <i>Do not forget any documents, and be as precise as possible.</i>2) For each document, after extracting the most helpful passages discuss whether they are actually relevant or irrelevant for answering the #Question. <i>For clarity in your answer, provide the exact passages of each document referring to the document number, organizing the explanation as follows: passage claims [1] that [1], in contrast [4] claims....</i>3) Please consider the passages in step 2) in detail, ensure they are correct. Then, discuss the provided passages by delivering a single rationale that considers the supporting motivations from a <i>contrastive</i> perspective as concern relevant and irrelevant passages. <i>For clarity, provide a detailed explanation by marking it as “#Explanation:”</i>4) Finally, after explaining the rationale supporting the final answer to facilitate the final evaluation, extract the answer in a short and concise format by marking it as “#Answer:”
<p>#Question {question}</p>

Table 6: The Contrastive RAG (C-RAG) framework instructs the model to deliver multi-step reasoning paths that lead the models to solve the task by providing an explanation that contrasts the perspectives that have emerged.

C Instructions Data

As introduced in §3, we use our Contrastive RAG (C-RAG) to instruct Llama-2-7b and Llama-2-13 to address knowledge-intensive tasks using retrieved documents using *contrastive* perspective (§3.5). Since C-RAG alone do not benefit from the abilities of the baseline models (particularly Llama-2-7b and Llama-2-13b without further tuning as reported in Table 1), we use GPT-4 (version gpt-4o-2024-05-13) as the annotation model. We systematically prompt GPT-4 using the instructions reported in Table 6.

We operate through GPT-4, which produced synthetic demonstrations, to instruct models to deliver C-RAG multi-step reasoned-solving strategies. Although this model addresses the tasks by following the instructions provided exhaustively (Peng et al., 2023), these may still be incorrect and contain misleading information. Therefore, we checked the quality by filtering out high-quality demonstrations to refine the instruction set. Hence, by examining the answers, we eliminated all incorrect ones (i.e., all generations that do not contain the final target string metric better known as *Exact Match*). We then checked that all the necessary steps were encoded in the remaining answers.

D Data Composition

Task	Total	Correct	C-RAG	Used
NQ	2,5k	1,9k	1.10k	515
PQA	2,5k	1,1k	0,75K	500
TQA	2,5k	1,5k	0,51K	500
FEVER	2,5k	0,9k	0,48K	485
Total	10k	6,0k	2,8k	2,0k

Table 7: Data used to construct C-RAG instructions. We applied the annotation as explained in §3.4.2. We obtained the following correct answers, filtered according to the heuristics in Appendix C, and balanced for the four tasks. *(1k is equal to 1000).

E Eliminating Components

We reproduced the experiment discussed in §5 Table 2. However, we used the -7b version unlike the previous one.

Models	NQ	PopQA	TriviaQA
	(%acc)	(%acc)	(%acc)
COMPLETE	40.2	56.4	68.4
RANDOM	32.0	52.6	57.0
w/o (2)	31.6	50.8	59.8
w/o (3)	36.8	54.2	61.3
w/o (4)	38.7	55.6	66.2

Table 8: Ablation study on three QA task with C-RAG-7B. We analyze the impact of each component on tuning by eliminating (w/o) or RANDOM shuffling the four defined steps (§3).

F Misleading C-RAG

Since the annotations produced through C-RAG are not always of good quality and correct (see Table 7), we define these demonstrations as misleading (obtained through prompting C-RAG but with an incorrect final target). To observe their impact on tuning, we produced the experimental setting of §4 by varying the demonstrations as shown in Table 9.

Models	NQ	PopQA	TriviaQA	FEVER
	(%acc)	(%acc)	(%acc)	(%acc)
<i>baseline</i>	27.8	47.8	55.6	23.2
CL-RAG-7B	40.2	56.4	68.4	79.2
<i>mixed</i>	28.0	49.8	56.7	67.3
<i>misleading</i>	27.2	48.6	54.2	30.8

Table 9: The instructions used for our CL-RAG-7B and -13B are selected and filtered as detailed in §3.4.2. We conducted tuning on incorrect (i.e., *misleading*) and *mixed* (50% *misleading* and 50% filtered correct demonstrations) and reported the *baseline* (Llama-2-7b).

G Example Generations

Question: *When was catch me if you can made?*

Llama-2-7b (baseline)

The movie "Catch Me If You Can" is a biographical crime drama directed by Spielberg released in **2000**.

Llama-2-7b (RAG)

Based on the information retrieved from Documents [1-5], we can determine that the Catch Me if You Can development began in 1980 and was released on **1989**.

Self-Reasoning-7b (extracted from (Xia et al., 2024))

Relevant: True

Relevant Reason: The provided documents are relevant with question.

Cite content: [1] the original start date was January 2002, but was pushed to February 7 in Los Angeles

Reason to cite: This piece provides information on the commencement and location of filming for 'Catch Me If You Can', indicating that it started in April 2002

Cite content: [3] When production finished, the studio that had promised to distribute the movie had gone out of business, and it was eventually released by MCA Inc. in July 1989.

Reason to cite: This piece provides information on the release date of 'Catch Me If You Can', indicating that it was eventually released by MCA Inc. in July.

The film 'Catch Me If You Can' was made in 2002. It started filming in April 2002 in Park Avenue, just outside the Waldorf-Astoria Hotel, and moved to Orange, New Jersey, before returning to Brooklyn for bank and courthouse scenes [1]

C-RAG-7B

Reference Documents:

- [1]: The film due to her busy schedule [...]

- [2]: Development for the film started in 1980 but did not book [...]

- [3]: The film was Sommers's directorial debut and was written by him. When production finished [...]

[other # Reference Documents]

Analysis:

[1] *Passage claims:* "The original start date was January 2002, but was pushed to February 7 in Los Angeles, California." *Relevance:* This is directly relevant as it provides the specific production start date for *Catch Me If You Can* in early 2002.

[2] *Passage claims:* "Development for the film started in 1980, but did not progress until 1997, when Spielberg's DreamWorks bought the film rights to Abagnale's book." *Relevance:* Relevant for providing context on the film's development timeline, but no specific filming dates.

[3] *Passage claims:* "It was funded independently and had a budget of 800,000...released by MCA Inc. in July 1989." *Relevance: Irrelevant as it pertains to a different film.*

Explanation:

Document #1 provides the direct answer with filming starting in February 2002. *Document #2* gives development context but no dates for filming. *Document #3* is unrelated and irrelevant. (explanations for the remaining Docs)

Answer:

Catch Me If You Can was made in early 2002, with filming beginning in February **2002**.

Table 10: An example of C-RAG inference compared to other approaches. C-RAG generates reasoned explanations, delivering the final correct answer.