

Removal of Hallucination on Hallucination: Debate-Augmented RAG

Wentao Hu[♣] Wengyu Zhang[♣] Yiyang Jiang[♣]
Chen Jason Zhang[♣] Xiaoyong Wei^{♡,♣,*} Qing Li[♣]
[♣]The Hong Kong Polytechnic University [♡]Sichuan University
wayne-wt.hu@connect.polyu.hk
{jason-c.zhang, cs007.wei}@polyu.edu.hk

Abstract

Retrieval-Augmented Generation (RAG) enhances factual accuracy by integrating external knowledge, yet it introduces a critical issue: erroneous or biased retrieval can mislead generation, compounding hallucinations, a phenomenon we term *Hallucination on Hallucination*. To address this, we propose **Debate-Augmented RAG (DRAG)**, a training-free framework that integrates Multi-Agent Debate (MAD) mechanisms into both retrieval and generation stages. In retrieval, DRAG employs structured debates among proponents, opponents, and judges to refine retrieval quality and ensure factual reliability. In generation, DRAG introduces asymmetric information roles and adversarial debates, enhancing reasoning robustness and mitigating factual inconsistencies. Evaluations across multiple tasks demonstrate that DRAG improves retrieval reliability, reduces RAG-induced hallucinations, and significantly enhances overall factual accuracy. Our code is available at <https://github.com/Huenao/Debate-Augmented-RAG>.

1 Introduction

Large Language Models (LLMs) have demonstrated remarkable natural language understanding and reasoning capabilities (Achiam et al., 2023; Touvron et al., 2023). However, their reliance on parametric knowledge introduces a critical challenge, hallucination, where the generated content deviates from factual correctness (Ji et al., 2023; Huang et al., 2024). This issue severely limits their reliability, particularly in knowledge-intensive tasks where factual accuracy is paramount.

To mitigate hallucinations, Retrieval-Augmented Generation (RAG) has been proposed as a framework that integrates external knowledge retrieval to enhance LLM outputs (Gao et al., 2023). By conditioning responses on retrieved documents, RAG

* Corresponding author

(a) Hallucinations Originate in the Retrieval Stage

Q: Who is the girl playing with guns and roses?
A: **Melissa Reese**.

Retrieval

Doc 1: "Guns and Roses (2012 film)" Xiao Dongbei tries to beg for money. By luck, he runs into **Gu Xixi** (Cheng Yuanyuan), the daughter of the bank's president. He saves her from an air raid while being injured himself.
Doc 2: "Guns and Roses (Numbers)" his former girlfriend from Albuquerque, New Mexico, Bureau of Alcohol,
Doc 3: "Rose Hip Rose" Rose Hip Rose The series primarily deals with Kasumi Asakura, a.k.a. Rose Hip

Generation
Gu Xixi

(b) Hallucinations Originate in the Generation Stage

Q: Who does Fez marry in that 70's show?
A: **Laurie Forman**.

Retrieval

Doc 1: "Fez (That '70s Show)". Fez worked at DMV from season 5 to season 7, where he met **Nina**, his short stint girlfriend.....
Doc 2: "Green card marriage ": In the season 5 finale of "That '70s Show", **Fez marries Laurie Forman** to avoid deportation. In the two-part season two finale of the sitcom "Melissa & Joey", Joe proposes to and marries Elena, his Russian girlfriend,
Doc 3: "Fez (That '70s Show)" Fez is a fictional character and one of the four male leads on the Fox Network's "That '70s Show", ...

Generation
Nina

Figure 1: Demonstration of Hallucination in Retrieval-Augmented Generation. In the first example, the error stems from retrieving information about a movie with the same name instead of the correct entity (a band). In the second, despite retrieving accurate information, retrieval noise still leads to an incorrect response.

reduces reliance on the model's parametric knowledge, thereby aiming to improve factual correctness. However, as shown in Figure 1, RAG introduces a new challenge: *Biased or erroneous retrieval results can mislead the generation, compounding the hallucination problem rather than solving it*. In other words, the LLM's inherent hallucinations are further amplified when retrieval provides incomplete, biased, or misleading information, leading to what we term *Hallucination on Hallucination*. Addressing this challenge requires optimizing the entire RAG pipeline, encompassing both retrieval and generation stages.

In the retrieval stage, hallucinations frequently arise from insufficient retrieval, where incomplete or biased results mislead the generation phase. Even when retrieval strategies are optimized, the generation stage remains vulnerable to retrieval noise and misinformation (Ngo et al., 2007), which can propagate hallucinations rather than mitigate them. Existing approaches to reducing retrieval bias include iterative retrieval (Trivedi et al., 2023; Shao et al., 2023; Feng et al., 2024) and autonomous retrieval (Jiang et al., 2023; Su et al., 2024; Yu et al., 2024; Wang et al., 2025), both of which seek to counteract hallucination stemming from knowledge deficiency and retrieval noise. Additionally, reflection during generation (Asai et al., 2023) and retrieval summarization (Kim et al., 2024) have been proposed to enhance the robustness and factual consistency of generation. Existing methods do not systematically optimize the entire pipeline. These methods also rely on LLMs for self-decision-making and self-optimization, making them susceptible to inherent biases, which limits their ability to generate diverse queries and robustness (Wei and Ngo, 2008).

Recent advancements in Multi-Agent Debate (MAD) offer a promising paradigm for enhancing LLM robustness, factual accuracy, and reasoning diversity (Becker, 2024). MAD introduces multiple independent LLM agents that engage in structured debates, iteratively refining their responses through critical evaluation and multi-agent verification. This process improves factual consistency and enhances robustness in reasoning (Du et al., 2024). Furthermore, MAD fosters perspective diversity by assigning distinct roles to agents, thereby mitigating biases that arise in single-agent approaches (Yin et al., 2023; Liang et al., 2024).

Motivated by this, we leverage MAD as a mechanism to address the challenge in RAG and propose **Debate Augmented RAG (DRAG)**. DRAG is a novel training-free framework integrating multi-agent debate mechanisms across retrieval and generation phases. In the retrieval phase, multiple agents engage in structured debates to assess retrieval adequacy. DRAG maintains a dynamic query pool by incorporating proponents advocating retrieval strategies, opponents questioning query sufficiency, and judges evaluating completeness, ensuring broader knowledge coverage and reduced factual biases. In the generation phase, we establish asymmetric information roles among agents, reducing the LLM’s over-reliance on retrieved content

while promoting structured adversarial debate to enhance reasoning robustness and mitigate factual inconsistencies. We evaluate DRAG on multiple tasks and demonstrate that this structured debate process leads to more reliable retrieval, enhanced reasoning robustness, and substantially reduced RAG-induced hallucinations. Overall, our main contributions can be summarized as follows:

- We propose a novel perspective on the challenges in RAG: *Hallucination on Hallucination*.
- We introduce DRAG, a novel training-free framework that integrates Multi-Agent Debate (MAD) mechanisms into both retrieval and generation phases of RAG.
- We evaluate DRAG on multiple tasks, demonstrating its effectiveness in improving retrieval reliability, reasoning robustness, and reducing RAG-induced hallucinations.

2 Related Work

2.1 Retrieval-Augmented Generation

LLMs have demonstrated remarkable success across a wide range of tasks. However, they are prone to hallucinations, particularly when handling queries beyond their training data, such as in domain-specific or knowledge-intensive tasks. To mitigate this issue, Retrieval-Augmented Generation (RAG) enhances LLMs by incorporating external knowledge retrieval into the response generation process (Gao et al., 2023). The most straightforward approach involves using the initial input as a query to retrieve information from an external corpus, which is then integrated into the model’s input as supplementary knowledge (Guu et al., 2020; Lewis et al., 2020; Izacard and Grave, 2021; Izacard et al., 2022; Shi et al., 2024). However, single-round retrieval based solely on the initial input may fail to retrieve essential external knowledge for complex tasks and introduce harmful noise.

To meet complex retrieval needs, IRCot (Trivedi et al., 2023) and ITER-RETGEN (Shao et al., 2023) introduce iterative retrieval-generation loops, allowing LLMs to refine their retrieval strategy dynamically to address insufficient retrieval. Auto-RAG (Yu et al., 2024) further extends this by enabling autonomous query decision-making, enhancing retrieval flexibility. FLARE (Jiang et al., 2023) and DRAGIN (Su et al., 2024) adopt a token-triggered

retrieval mechanism, activating retrieval only when necessary and improving efficiency. However, focusing on local tokens may neglect global reasoning needs. In contrast, Auto-RAG leverages LLM reasoning to optimize queries autonomously. Self-RAG (Asai et al., 2023) autonomously retrieves via reflection and follows a reflection-retrieval-regeneration process to improve factual consistency between generated and retrieved content. RECOMP (Xu et al., 2023) and SuRe (Kim et al., 2024) optimize LLM generation through summarized retrieval, enhancing the utilization of external information, reducing information overload, and improving factual consistency. Existing methods lack a comprehensive analysis and optimization of RAG. Moreover, relying on a single-agent decision-making and reasoning process may introduce biases that impact performance.

2.2 Multi-Agent Debate

Recent studies have demonstrated that multi-agent interactions significantly enhance LLM capabilities across various dimensions (Becker, 2024). EoT (Yin et al., 2023) mitigates biases and blind spots inherent in single-agent approaches by integrating external insights from other agents. Moreover, Multi-Agent Debate (MAD) enables agents to iteratively challenge and refine each other’s claims, improving factual consistency (Du et al., 2024). Beyond factual alignment, MAD fosters divergent thinking, reducing reasoning stagnation and enhancing reasoning depth (Liang et al., 2024). BDoG (Zheng et al., 2024) extends this paradigm to multimodal reasoning while introducing a blueprint to mitigate the trivialization of opinions and focus diversion. The capabilities of MAD align well with the second-order hallucination problem in RAG, and this paper represents an initial endeavor to explore extending these capabilities to address this specific challenge.

3 Debate-Augmented RAG

In this section, we introduce Debate-Augmented RAG (DRAG), a general framework designed to enhance Retrieval-Augmented Generation (RAG) through a structured adversarial debate mechanism.

We first give an overview of DRAG in Section 3.1, followed by a comprehensive discussion. DRAG consists of a two-stage debate framework that comprehensively addresses the hallucinations introduced by RAG in both the retrieval and gen-

eration stages, as illustrated in Figure 2. Stage one is the Retrieval Debate, which is introduced in Section 3.2, while stage two is the Response Debate, discussed in Section 3.3.

3.1 Overview

Formally, given a question x , a retriever \mathcal{R} retrieves top- k supporting paragraphs C_x^k from a knowledge corpus C :

$$C_x^k = \mathcal{R}(x, C, k) \quad (1)$$

Once retrieval concludes, the retrieved paragraphs C_x^k and the question x are used to prompt a large language model (LLM) \mathcal{M} , which then generates an answer y :

$$y = \mathcal{M}(x, C_x^k) \quad (2)$$

To mitigate hallucinations introduced at both retrieval and generation stages in RAG, DRAG employs a two-stage adversarial debate mechanism: **Retrieval Debate** and **Response Debate**. This debate-driven framework refines the retrieved content and ensures a robust, factually grounded response. The following subsections introduce the debate mechanisms for each stage separately.

3.2 Retrieval Debate

The retrieval process in standard RAG models often suffers from biases due to single-pass retrieval and suboptimal query formulation. DRAG introduces the Retrieval Debate $f_{\text{RetDebate}}$, a dynamic multi-agent debate mechanism to iteratively refine retrieval and retrieved contents:

$$\hat{Q}, C_{\hat{Q}} = f_{\text{RetDebate}}(x, \mathcal{R}) \quad (3)$$

where, $\hat{Q} = \{q_i\}_{i=1}^n$ represents the optimal query pool with n queries for the question x , while $C_{\hat{Q}} = \{C_{q_i}^k\}_{i=1}^n$ denotes the corresponding retrieved paragraph set.

Debate Structure. The debate structure is a multi-round debate on optimizing the retrieval pool Q between m distinct agents $A = \{a_i\}_{i=1}^m$. The debate at the j^{th} round can be formulated as follows:

$$\mathcal{T}^j = \text{Debate}(x, Q^j, C_{Q^j}, A) \quad (4)$$

Each debate round aims to determine the appropriate refinement operation \mathcal{T}^j applied to the retrieval pool Q^j , which is subsequently updated to Q^{j+1} to initiate the next round.

Agents and Roles. Just like a real debate, this debate is also structured around agents having distinct

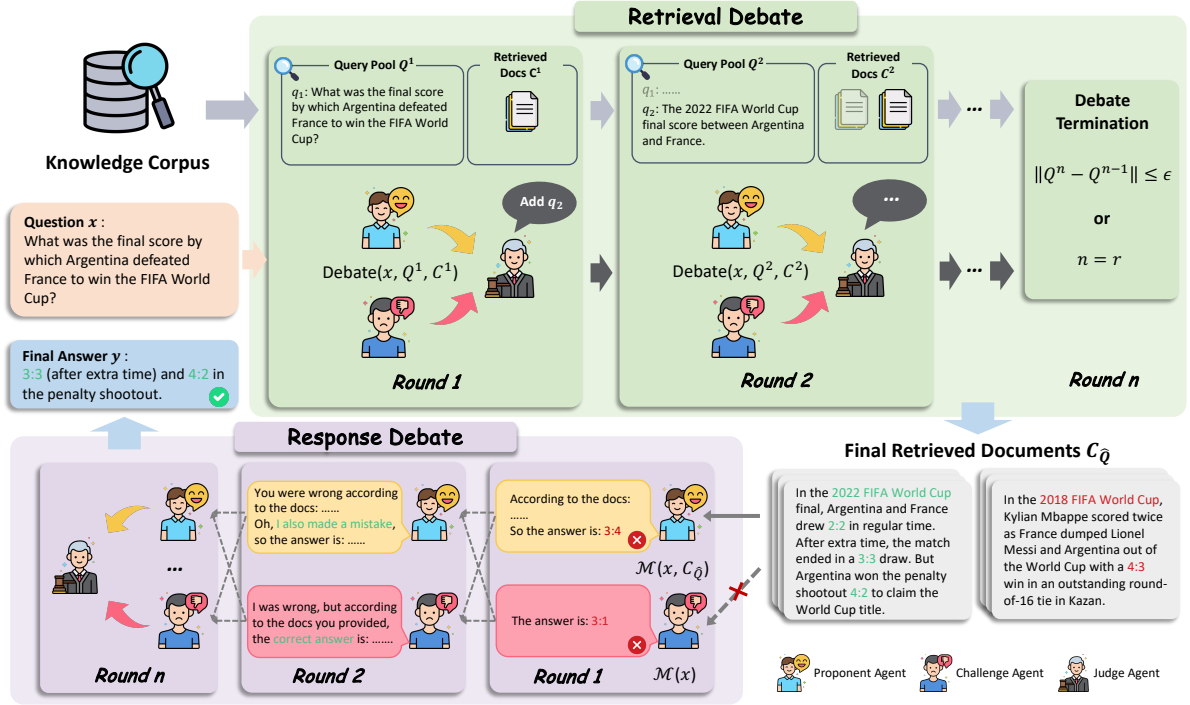


Figure 2: An overview of our Debate-Augmented RAG (DRAG) framework. It iteratively refines the retrieval strategy and enhances factual consistency.

roles assigned. These different roles have distinct responsibilities, aiming to ensure a comprehensive and in-depth exploration of retrieval strategies by promoting perspective diversity and critical thinking, ultimately determining the optimal refinement operations for the retrieval pool. In this stage, we define three agent roles.

Proponent Agent is required to provide justifications to support that the current query Q^j is reasonable and that the retrieved results C_{Q^j} are sufficient. Therefore, the operation supported by the proponent agent is:

$$\mathcal{T}_p : Q^{j+1} = Q^j \quad (5)$$

that is, to keep Q^j unchanged.

Challenger Agent critiques Q^j , identifying deficiencies and proposing modifications. Specifically, the challenger agent is responsible for proposing modifications to unreasonable queries or initiating new retrieval queries to address further knowledge needs. So challenger agent insists on the refinement operation for the query pool Q^j :

$$\mathcal{T}_c : Q^{j+1} = (Q^j \cap Q_{\text{useful}}) \cup Q_{\text{new}} \quad (6)$$

where $Q_{\text{useful}} = \{q_i \in Q^j\}$ represents the subset of queries in Q^j that the challenger agent deems necessary to retain. Q_{new} consists of additional queries

introduced by the challenger agent to supplement missing information or enhance retrieval quality.

Judge Agent evaluates the arguments presented by the proponent agent and the challenger agent and determines the ultimate operation \mathcal{T}^j for the current round of debate:

$$\mathcal{T}^j = \text{Select}(\mathcal{T}_p, \mathcal{T}_c) \quad (7)$$

where Select represents the decision-making function used by the judge agent.

Update Query Pool Q . The debate starts with an initial query set $Q^0 = \{x\}$. After each round of debate, the query pool Q^j is iteratively refined following the judge agent's decision and gets $Q^{j+1} = \mathcal{T}^j(Q^j)$, the specific operations as Eq. 5 and Eq. 6.

The debate process continues iteratively until a convergence criterion is met:

- The query set converges when $\|Q^{j+1} - Q^j\| \leq \epsilon$, where ϵ is a predefined threshold ensuring minimal query modification.
- The maximum number of debate rounds r is reached.

Once the debate terminates, the final query set \hat{Q} is used for retrieval, forming the evidence set $C_{\hat{Q}}$ for the next stage, the Response Debate.

3.3 Response Debate

Even after optimizing retrieval through the Retrieval Debate, the model may still be influenced by noisy or incorrect retrieval results. To address this, DRAG introduces the Response Debate $f_{\text{ResDebate}}$, which iteratively enhances factual consistency and reasoning robustness:

$$\hat{y} = f_{\text{ResDebate}}(x, C_{\hat{Q}}) \quad (8)$$

where, \hat{y} represents the final factually verified and robust response.

Structure and Roles. As shown in Figure 2, unlike the independent debates centered on the retrieval pool in the Retrieval Debate stage, the Response Debate fosters direct exchanges between agents. To prevent over-reliance on retrieved information, the Response Debate employs an asymmetric information setting, ensuring agents engage in substantive argumentation rather than merely restating retrieved facts. If all agents have equal access to retrieval results, they may uncritically accept the retrieved content, reinforcing biases instead of evaluating reliability. Our experiments further confirm the effectiveness of this information asymmetry.

The information asymmetry is enforced through role assignments, defining three distinct agent roles:

Proponent Agent initiates the response based on retrieved information $C_{\hat{Q}}$ in first round:

$$y_{A_p}^1 = \mathcal{M}_{A_p}(x, C_{\hat{Q}}) \quad (9)$$

and in the i^{th} round, refines its response by incorporating the challenger agent’s answer:

$$y_{A_p}^i = \mathcal{M}_{A_p}(x, \{y_{A_c}^{<i}\}, y_{A_c}^{i-1}) \quad (10)$$

where $y_{A_c}^{i-1}$ represents the response from the challenger agent, which we defined later, in the $(i-1)^{\text{th}}$ round.

Challenger Agent generates an initial response based on internal knowledge:

$$y_{A_c}^1 = \mathcal{M}_{A_c}(x) \quad (11)$$

and iteratively updates its response based on the proponent agent’s answer:

$$y_{A_c}^i = \mathcal{M}_{A_c}(x, \{y_{A_p}^{<i}\}, y_{A_p}^{i-1}) \quad (12)$$

Judge Agent: Evaluates the final responses from both agents and selects the optimal answer:

$$\hat{y} = \mathcal{M}_{A_j}(x, y_{A_p}^r, y_{A_c}^r) \quad (13)$$

where r denotes the maximum number of debate rounds.

4 Experimental Setup

4.1 Datasets

We experimented on six benchmark datasets of three tasks. We follow previous work (Jiang et al., 2023) and randomly select 500 examples from each dataset for testing.

Open-domain QA. Open-domain QA aims to answer a question in the form of natural language based on large-scale unstructured documents. Here, we provide the model with only the questions without reference documents. We choose three datasets for this task: NQ (Kwiatkowski et al., 2019), TriviaQA (Joshi et al., 2017), PopQA (Mallen et al., 2023). For the evaluation metrics, we extract the final answer from the output using regular expressions and calculate the exact match (EM) and token-level F1 score.

Multi-hop QA. Multi-hop QA focuses on answering questions that require the ability to gather information from multiple sources and conduct multi-step reasoning to arrive at a comprehensive answer. We choose the 2WikiMultihopQA (Ho et al., 2020) and HotpotQA (Yang et al., 2018). We also report EM and token-level F1 score for these datasets.

Commonsense Reasoning. Commonsense reasoning requires world and commonsense knowledge to generate answers. We utilize the StrategyQA (Geva et al., 2021) to evaluate DRAG and other baselines. We extract the yes/no response and compare it with the standard correct answer using EM.

4.2 Baselines

We select **Naive Gen**, which is generated directly by an LLM, and **MAD** (Du et al., 2024) as baselines that do not incorporate retrieval. For retrieval-based baselines, we use **Naive RAG**, which represents the standard RAG method without any optimization. Additionally, we compare DRAG with RAG optimization baselines, including **IRCoT** (Trivedi et al., 2023), **ITer-RetGen** (Shao et al., 2023), and **FLARE** (Jiang et al., 2023), which focus on improving the retrieval stage, as well as **Self-RAG** (Asai et al., 2023) and **SuRe** (Kim et al., 2024), which enhance the generation stage.

4.3 Implementation Details

We selected Llama-3.1-8B-Instruct (Dubey et al., 2024) as our base LLM model. Our implementation of DRAG and the baselines are built upon FlashRAG (Jin et al., 2024), a Python toolkit designed for the reproduction and development of

Method	NQ		TriviaQA		PopQA		2Wiki		HotpotQA		StrategyQA
	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1	EM
<i>Without Retrieval</i>											
Naive Gen	17.40	26.27	56.60	65.50	19.20	23.33	8.60	16.50	16.80	23.88	63.20
MAD	21.80	33.11	56.40	66.39	21.40	28.28	18.20	25.13	23.00	32.79	66.20
<i>With Retrieval</i>											
Naive RAG	38.20	50.08	60.80	69.55	37.60	45.69	14.80	24.27	25.80	35.80	62.60
IRCoT	28.60	37.36	47.20	54.56	27.00	33.02	22.80	31.19	25.20	34.40	53.60
Iter-RetGen	40.80	52.31	63.00	72.23	39.60	46.41	15.00	24.75	27.80	38.93	62.00
FLARE	19.40	27.68	53.60	63.05	21.60	24.35	9.20	20.13	16.60	23.74	42.80
SuRe	34.80	51.35	47.60	64.11	41.80	48.99	10.20	18.60	19.00	32.39	0.00
Self-RAG	44.00	52.20	46.40	58.37	22.00	34.38	13.00	26.63	14.80	28.81	0.40
DRAG	36.80	50.38	60.80	69.93	38.60	46.50	28.80	36.97	30.80	41.74	69.20

Table 1: The overall evaluation results of DRAG and other baselines on six benchmarks. Blue marks the best-performing method, light blue represents the second-best-performing method. All methods are evaluated under the same settings.

RAG-based systems. Following FlashRAG, we employed E5-base-v2 (Wang et al., 2022) as the retriever and used the widely adopted Wikipedia dump from December 2018 as the retrieval corpus (Karpukhin et al., 2020). We retrieve top-3 paragraphs for each query for our DRAG and other iterative baselines. The maximum debate interactions r is set to 3. We set the number of agents per stage to 3, with one agent assigned to each role. The threshold ϵ is set to 0.

5 Experimental Results

5.1 Comparison with Baselines

We conduct comprehensive experiments to evaluate the performance of DRAG against various baselines across six benchmark datasets. As summarized in Table 1, DRAG achieves strong and consistent improvements on multi-hop reasoning tasks, underscoring its robustness in settings that demand complex retrieval and multi-step inference. In comparison, on single-hop QA tasks, DRAG exhibits competitive performance, ranking second on several benchmarks with only marginal differences from the top-performing methods.

On multi-hop reasoning tasks such as 2WikiMultiHopQA and HotpotQA, DRAG achieves substantial improvements in Exact Match (EM) scores, outperforming the best baseline methods by 6 and 3, respectively. This superior performance can be attributed to DRAG’s structured multi-agent debate mechanism, which enhances retrieval adequacy and improves generation robustness. Compared to generation-focused models such as SuRe and

Self-RAG, DRAG demonstrates that merely optimizing the generation stage is insufficient for complex reasoning tasks; instead, retrieval quality plays a crucial role. This aligns with our hypothesis that low-quality and insufficient queries negatively impact RAG model performance. Additionally, when compared to retrieval-optimized models such as IRCoT and Iter-RetGen, DRAG exhibits a stronger ability to dynamically refine retrieval strategies, ensuring broader and more relevant knowledge retrieval.

While DRAG excels in multi-hop reasoning tasks, its performance on single-hop tasks is slightly lower than some retrieval-optimized baselines. We attribute this to the phenomenon of problem drift induced by the debate mechanism (Becker, 2024). Specifically, excessive debates among agents may introduce unnecessary complexity, generating additional queries or conflicting perspectives that deviate from the optimal answer. However, this issue can be effectively mitigated by adjusting the number of debate rounds, as further analyzed in Section 5.2.

5.2 Analysis of Debate Rounds

To systematically assess the impact of debate rounds on DRAG’s performance, we analyze the effects of Retrieval Debate and Response Debate rounds separately. In each experiment, we varied the number of rounds in one stage (either retrieval or response) from 0 to 4, while keeping the number of rounds in the other stage fixed at 3.

Retrieval debates help refine the retrieval pool.

Table 2 shows that increasing the Retrieval Debate

Settings	NQ	TriviaQA	PopQA	2Wiki	HotpotQA
	EM	EM	EM	EM	EM
DRAG _{RetIter=0}	34.40	60.40	37.60	26.00	30.20
DRAG _{RetIter=1}	36.80	60.80	37.40	29.20	31.40
DRAG _{RetIter=2}	36.80	60.76	37.50	29.60	30.80
DRAG _{RetIter=3}	36.80	60.80	37.60	28.80	30.80
DRAG _{RetIter=4}	37.00	60.80	37.60	29.00	31.00
DRAG _{ResIter=0}	38.20	57.20	39.60	16.00	28.20
DRAG _{ResIter=1}	37.60	62.40	39.20	27.80	33.40
DRAG _{ResIter=2}	38.20	60.20	38.80	28.20	31.20
DRAG _{ResIter=3}	36.80	60.80	38.60	28.80	30.80
DRAG _{ResIter=4}	37.00	60.60	38.80	27.20	31.00

Table 2: Model performance with respect to the iteration round of debate.

	Benchmark	Debate Rounds	Query Counts
<i>Single-hop</i>	NQ	1.27	1.06
	TriviaQA	1.33	1.08
	PopQA	1.31	1.07
	StrategyQA	1.18	1.07
<i>Multi-hop</i>	2Wiki	1.57	1.17
	HotpotQA	1.46	1.17

Table 3: The average number of debate rounds and the counts of queries in the retrieval debate stage of DRAG on single-hop QA tasks and multi-hop QA tasks.

rounds generally enhances performance by progressively refining the retrieval pool. For instance, on 2WikiMultiHopQA, performance improves from 26.00 (RetIter=0) to 29.60 (RetIter=2) before stabilizing. This trend indicates that iterative refinement helps mitigate retrieval noise and improves knowledge coverage. However, beyond a certain threshold (e.g., RetIter=3), additional debate rounds contribute marginal gains. We attribute this to the convergence of the retrieval pool in the Retrieval Debate stage.

Response debates enhance reasoning but risk problem drift. In the Response Debate stage, performance initially improves as adversarial discussions enhance factual consistency, but excessive debate rounds lead to performance degradation. This is particularly evident in TriviaQA, where the EM score drops from 62.4 (ResIter=1) to 60.20 (ResIter=2), suggesting that early problem drift can emerge when excessive rounds introduce unnecessary complexity. Similar effects are observed in HotpotQA and 2WikiMultiHopQA, reinforcing the idea that over-iterating debates may amplify inconsistencies rather than resolve them.

	NQ	TriviaQA	2Wiki	HotpotQA
Question Counts	139	176	338	331
w/o $f_{\text{ResDebate}}$				
Avg EM	3.60	12.50	7.10	9.32
Avg F1	12.93	23.42	18.01	19.04
with $f_{\text{ResDebate}}$				
Avg EM	5.04	28.98	19.23	17.36
Avg F1	18.50	40.21	28.52	29.64

Table 4: Comparison of performance with and without $f_{\text{ResDebate}}$ across different datasets.

5.3 Analysis of the Retrieval Debate

DRAG dynamically adapts retrieval strategies to task complexity. To examine DRAG’s retrieval debates, we analyze debate rounds and query counts across single-hop and multi-hop QA tasks. Multi-hop QA requires retrieving interrelated evidence, demanding more iterative refinement than single-hop QA, which typically retrieves a single fact. Table 3 shows a clear adaptation pattern in DRAG’s retrieval behavior. Table 3 shows that for single-hop QA, DRAG converges quickly with minimal refinement (1.18 to 1.33 debate rounds, 1.06 to 1.08 queries per instance). For multi-hop QA, it engages in more debates (1.57 rounds for 2Wiki, 1.46 for HotpotQA) and higher query counts (1.17 rounds), ensuring adequate contextual retrieval. These results highlight DRAG’s adaptive retrieval mechanism, dynamically adjusting debate intensity to optimize efficiency and retrieval sufficiency based on task complexity.

5.4 Analysis of the Response Debate

Response debate enhances robustness against retrieval deficiency. To further examine DRAG’s ability to handle incomplete retrieval, we calculate the average EM with and without the Response Debate on a subset of questions where the gold answer is absent from the retrieved documents. As shown in Table 4, this result further reinforces DRAG’s advantage over conventional RAG models by mitigating the impact of incomplete retrieval and ensuring a more robust and verifiable response generation process.

Enhances the final response’s alignment with the retrieved evidence. To assess the impact of response debates on factual consistency, we analyze the Response Debate stage in DRAG. As shown in Figure 3, this case study illustrates how multi-agent debates iteratively refine responses, correcting inaccuracies and reinforcing evidence-based reasoning.

Q: Who does Fez marry in that 70's show?

Retrieval

Query 1: who does Fez marry in that 70's show

Retrieved Content:

Doc 1(Title: "Fez (That '70s Show)"): Fez worked at DMV (Department of Motor Vehicle) from season 5 to season 7, where he met Nina, his short stint girlfriend. With Nina, Fez managed to lose his virginity. Fez continued to work there after finishing school and being granted a green card ...
Doc 2(Title: "Green card marriage"): forced into a green card marriage so Catalina could come back to America. In the season 5 finale of "That '70s Show", Fez marries Laurie Forman to avoid deportation. In the two-part season two finale of the sitcom "Melissa & Joey", Joe proposes to and marries Elena, his Russian girlfriend, who will likely be arrested upon her return to Russia after her company is shut down due to a Ponzi scheme that she was either unaware of or not a participant in ...
Doc 3(Title: "Fez (That '70s Show)"): Fez (That '70s Show) Fez is a fictional character and one of the four male leads on the Fox Network's "That '70s Show", portrayed by Wilmer Valderrama. He was a foreign exchange student in a group of six local teenagers. Fez was born on August 4, 1960 ...

Round 1

Proponent Agent: To find the answer, I will look at the given documents. In Doc 1, it is mentioned that Fez worked at DMV and met Nina, his short-term girlfriend... In Doc 2, it is mentioned that Fez marries Laurie Forman to avoid deportation in the season 5 finale of "That '70s Show". In Doc 3, it does not mention Fez's marriage. The answer is: **Nina**.

Opponent Agent: The character Fez, played by Wilmer Valderrama, is a main character in the popular American sitcom "That '70s Show." Throughout the series, Fez's romantic relationships are a recurring plot point. However, Fez's most notable and long-term relationship is with Jackie Burkhart, played by Mila Kunis. The answer is **Jackie Burkhart**.

Round 2

Proponent Agent: The opponent agent's answer is incorrect because, in the given documents, it is mentioned in Doc 2 that Fez marries Laurie Forman to avoid deportation in the season 5 finale of "That '70s Show." This is a specific instance of Fez getting married in the show. Opponent Agent O's answer does not mention Fez's marriage, but rather a long-term relationship. The answer is: **Laurie Forman**

Opponent Agent: The Proponent Agent seems to have considered a few documents and didn't look at the entire series. While it's true that Fez does have a short-term girlfriend named Nina and that Doc 3 doesn't mention Fez's marriage, Doc 2's claim that Fez marries Laurie Forman to avoid deportation in the season 5 finale is actually a plot point in the show. So the answer is: **Laurie Forman**

Figure 3: Case study of the response debate.

Method	NQ		TriviaQA		PopQA		2wiki		HotpotQA		StrategyQA
	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1	EM
Naive RAG	38.20	50.08	60.80	69.55	37.60	45.69	14.80	24.27	25.80	35.80	62.60
+ $f_{RetDebate}$	38.20	50.79	57.20	65.31	39.60	47.04	16.00	25.91	28.20	37.48	64.60
+ $f_{ResDebate}$	34.40	47.61	60.40	69.60	37.60	46.43	26.00	34.30	30.20	41.00	68.80
DRAG	36.80	50.38	60.80	69.93	38.60	46.50	28.80	36.97	30.80	41.74	69.20
- <i>Info Asymmetry</i>	35.20	50.09	57.40	66.21	37.80	45.95	22.00	29.36	30.00	38.62	68.20

Table 5: Ablation study for key components of DRAG on six benchmarks. Blue marks the best-performing method, light blue represents the second-best-performing method.

5.5 Ablation Study

As shown in Table 5, we conduct an ablation study to evaluate the impact of key components in DRAG across six benchmarks.

First, we independently assess the effectiveness of the two debate stages: retrieval debate ($f_{RetDebate}$) and response debate ($f_{ResDebate}$). To this end, we introduce two DRAG variants: (1) + $f_{RetDebate}$: Incorporates retrieval debate in the retrieval stage but excludes the Response Debate. (2) + $f_{ResDebate}$: Incorporates response debate while omitting retrieval debate. Compared to Naive RAG, both variants improve performance across most tasks. This confirms that the Retrieval Debate enhances retrieval strategies by reducing information bias and improving factual coverage, leading to better answer generation. Meanwhile, response debate ensures proper utilization of retrieved information, mitigating hallucinations and enhancing factual accuracy.

To examine the effect of information asymmetry in response debate, we introduce the - *Info Asym-*




Figure 4: Average LLM and Retriever calls for DRAG and other baselines methods on the StrategyQA.

metry variant, where all agents have equal access to retrieval evidence. Results show consistent performance deterioration across all datasets, indicating that information asymmetry is crucial for promoting factual consistency. By enforcing role-based knowledge distribution, it encourages adversarial interactions, prevents agents from over-relying on retrieved content, and improves robustness in factual reasoning.

5.6 Efficiency Analysis

We provide detailed comparisons of average LLM and retrieval calls on StrategyQA, as shown in Figure 4. While DRAG requires more LLM calls than simpler baselines, making an average of 2.08 Retriever Calls and 10.36 LLM Calls (with 3.36 during the retrieval phase and 7 during the response phase), its overall cost remains comparable to other multi-turn frameworks such as IRCOT. This cost is further mitigated by DRAG’s adaptive termination mechanism in the retrieval debate stage, which helps eliminate unnecessary debate rounds. As a result, the number of LLM calls in the retrieval stage is significantly lower than in the response stage (7 rounds on average). Moreover, as discussed in Section 5.2, for simpler single-hop questions, only 1–2 debate rounds are often sufficient. In such cases, the number of LLM calls in the response stage can be reduced from 7 to 5 or fewer, achieving a more favorable trade-off between accuracy and efficiency. A more effective adaptive stopping mechanism in the response phase may further enhance DRAG’s efficiency, and we will elaborate in the Limitations section.

6 Conclusion

In this paper, we introduce Debate-Augmented RAG (DRAG), a novel framework designed to mitigate hallucinations in Retrieval-Augmented Generation (RAG) by leveraging multi-agent debate mechanisms. DRAG employs a two-stage structured debate: (1) Retrieval Debate, where multiple agents iteratively refine retrieval queries to improve knowledge coverage; (2) Response Debate, where agents adopt asymmetric roles and engage in structured debates to strengthen reasoning and mitigate reliance on flawed retrievals. Extensive experiments across six datasets spanning three knowledge-intensive tasks validate the efficacy of DRAG.

Limitations

A key limitation of our approach is the increased computational overhead introduced by multi-agent debates in both the retrieval and response stages. While the retrieval phase employs a Judge Agent to dynamically terminate early when sufficient evidence is gathered, the response phase currently uses a fixed number of debate rounds, which may lead to unnecessary LLM calls, especially in simple single-hop tasks.

This also introduces another limitation: problem drift, where excessive reasoning adds unnecessary complexity and reduces effectiveness in straightforward scenarios. Although DRAG adapts retrieval depth well, its response stage lacks such flexibility.

Future work should explore adaptive stopping criteria, agent pruning, or confidence-based early termination in the response stage to better balance reasoning depth and efficiency. We believe these enhancements will improve DRAG’s flexibility and performance across both single-hop and multi-hop settings. Investigating how to make the response debate more task-sensitive, by learning to infer when enough reasoning has been done, is a promising direction for future research.

Acknowledgement

This work was supported by the National Natural Science Foundation of China (Grant No.: 62372314), Hong Kong Research Grants Council under General Research Fund (project no. 15200023), as well as Research Impact Fund (project no R1015-23). This work was also supported by the following funding sources: industry donations from Accel Group Holding Limited and Minshang Creative Technology Holdings Limited (P0045948, P0046453); PolyU internal research funding (P0046701, P0046703); Research Matching Grant Scheme funded by the University Grants Committee (P0048183, P0048191); the Innovation and Technology Fund - ITSP, ITS/028/22FP (P0048887); Postdoc Matching Fund Scheme (P0048984); the RGC Early Career Scheme, 25600624 (P0051906); and Two Square Capital Limited donation (P0054482).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. [Self-rag: Learning to retrieve, generate, and critique through self-reflection](#). *arXiv preprint arXiv:2310.11511*.
- Jonas Becker. 2024. [Multi-agent large language models for conversational task-solving](#). *arXiv preprint arXiv:2410.22932*.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. 2024. [Improving](#)

- factuality and reasoning in language models through multiagent debate. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*. JMLR.org.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. *The llama 3 herd of models*. *arXiv preprint arXiv:2407.21783*.
- Zhangyin Feng, Xiaocheng Feng, Dezhi Zhao, Maojin Yang, and Bing Qin. 2024. *Retrieval-generation synergy augmented large language models*. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 11661–11665. IEEE.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. *Retrieval-augmented generation for large language models: A survey*. *arXiv preprint arXiv:2312.10997*.
- Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. *Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies*. *Transactions of the Association for Computational Linguistics*, 9:346–361.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. *Retrieval augmented language model pre-training*. In *International conference on machine learning*, pages 3929–3938. PMLR.
- Xanh Ho, Anh-Khoa Duong Nguyen, Saku Sugawara, and Akiko Aizawa. 2020. *Constructing a multi-hop QA dataset for comprehensive evaluation of reasoning steps*. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6609–6625, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2024. *A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions*. *ACM Transactions on Information Systems*.
- Gautier Izacard and Edouard Grave. 2021. *Leveraging passage retrieval with generative models for open domain question answering*. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 874–880, Online. Association for Computational Linguistics.
- Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and Edouard Grave. 2022. *Few-shot learning with retrieval augmented language models*. *arXiv preprint arXiv:2208.03299*.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. *Survey of hallucination in natural language generation*. *ACM Computing Surveys*, 55(12):1–38.
- 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.
- Jiajie Jin, Yutao Zhu, Xinyu Yang, Chenghao Zhang, and Zhicheng Dou. 2024. *Flashrag: A modular toolkit for efficient retrieval-augmented generation research*. *CoRR*, abs/2405.13576.
- 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.
- 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.
- Jaehyung Kim, Jaehyun Nam, Sangwoo Mo, Jongjin Park, Sang-Woo Lee, Minjoon Seo, Jung-Woo Ha, and Jinwoo Shin. 2024. *Sure: Summarizing retrievals using answer candidates for open-domain qa of llms*. *arXiv preprint arXiv:2404.13081*.
- 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.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. *Encouraging divergent thinking in large language models through multi-agent debate*. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 17889–17904, Miami, Florida, USA. Association for Computational Linguistics.

- Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, 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.
- Chong-Wah Ngo, Yu-Gang Jiang, Xiaoyong Wei, Feng Wang, Wanlei Zhao, Hung-Khoon Tan, and Xiao Wu. 2007. [Experimenting vireo-374: Bag-of-visual-words and visual-based ontology for semantic video indexing and search](#).
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. 2023. [Toolformer: Language models can teach themselves to use tools](#). *Advances in Neural Information Processing Systems*, 36:68539–68551.
- Zhihong Shao, Yeyun Gong, Yelong Shen, Minlie Huang, Nan Duan, and Weizhu Chen. 2023. [Enhancing retrieval-augmented large language models with iterative retrieval-generation synergy](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9248–9274, Singapore. Association for Computational Linguistics.
- Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2024. [REPLUG: Retrieval-augmented black-box language models](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 8371–8384, Mexico City, Mexico. Association for Computational Linguistics.
- Weihang Su, Yichen Tang, Qingyao Ai, Zhijing Wu, and Yiqun Liu. 2024. [DRAGIN: Dynamic retrieval augmented generation based on the real-time information needs of large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12991–13013, Bangkok, Thailand. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, et al. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2023. [Interleaving retrieval with chain-of-thought reasoning for knowledge-intensive multi-step questions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10014–10037, Toronto, Canada. Association for Computational Linguistics.
- Keheng Wang, Feiyu Duan, Peiguang Li, Sirui Wang, and Xunliang Cai. 2025. [LLMs know what they need: Leveraging a missing information guided framework to empower retrieval-augmented generation](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 2379–2400, Abu Dhabi, UAE. Association for Computational Linguistics.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. [Text embeddings by weakly-supervised contrastive pre-training](#). *arXiv preprint arXiv:2212.03533*.
- Xiao-Yong Wei and Chong-Wah Ngo. 2008. [Fusing semantics, observability, reliability and diversity of concept detectors for video search](#). In *Proceedings of the 16th ACM international conference on Multimedia*, pages 81–90.
- Fangyuan Xu, Weijia Shi, and Eunsol Choi. 2023. [Re-comp: Improving retrieval-augmented lms with compression and selective augmentation](#). *arXiv preprint arXiv:2310.04408*.
- Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.
- Zhangyue Yin, Qiushi Sun, Cheng Chang, Qipeng Guo, Junqi Dai, Xuanjing Huang, and Xipeng Qiu. 2023. [Exchange-of-thought: Enhancing large language model capabilities through cross-model communication](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15135–15153, Singapore. Association for Computational Linguistics.
- Tian Yu, Shaolei Zhang, and Yang Feng. 2024. [Auto-rag: Autonomous retrieval-augmented generation for large language models](#). *arXiv preprint arXiv:2411.19443*.
- Changmeng Zheng, Dayong Liang, Wengyu Zhang, Xiao-Yong Wei, Tat-Seng Chua, and Qing Li. 2024. [A picture is worth a graph: A blueprint debate paradigm for multimodal reasoning](#). In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 419–428.

A Experimental Details

A.1 More Implementation Details

We run our models on eight NVIDIA RTX A6000 GPU devices. To get a better performance, for IR-CoT, we set the iteration number to 5. For all baselines, we reproduced the key settings from their original papers as faithfully as possible under partially unified conditions, including the few-shot setting. In contrast, our DRAG performs inference under a zero-shot setting. As a training-based approach, Self-RAG utilizes the Self-RAG-Llama-2-7B¹ model provided by its authors for experimentation. Additionally, since SuRe and Self-RAG are not well-suited for answering questions with definitive yes/no answers, their performance is relatively poor.

A.2 Prompts for Retrieval Debate

An example of a proponent agent prompt in the Retrieval Debate stage is

Prompt for Proponent Agent

You are a debater. Argue that the current retrieved content is sufficient to answer the question and no further retrieval is needed. Deliver a brief, strong argument with clear reasoning. Do not suggest further retrieval.

Question:
{question}

Queries:
{queries in Q }

Retrieved Documents:
{retrieved results in C_Q }

In the Retrieval Debate stage, the challenger agent is responsible for proposing modifications to unreasonable queries or initiating new retrieval queries to address further knowledge needs. Inspired by Toolformer (Schick et al., 2023) and FLARE (Jiang et al., 2023), we can directly prompt the agent to choose two types of operations: Query Optimization and Query Expansion, and an example prompt is

¹https://huggingface.co/selfrag/selfrag_llama2_7b

Prompt for Challenger Agent

You are a critical thinker and debater, and your task is to challenge the sufficiency of the current retrieved content. Argue that the current information is insufficient to generate a reliable answer and propose either query optimization or query expansion.

The action you can choose:

1. Query Optimization: If the retrieved content is somewhat relevant but has expression or scope issues.

Optimize the query using this format:
Query Optimization: [Original Query] → [New Query].

2. Query Expansion: If critical information is missing.

Propose a new query using this format:
Query Expansion: [New Query].

Deliver a brief, strong argument with clear reasoning, and then you must choose only one action. The output must be in the exact format after your reasoning, without additional explanation, and keep the new query short and precise.

Question:
{question}

Queries:
{queries in Q }

Retrieved Documents:
{retrieved results in C_Q }

Instead of calculating the convergence conditions, we directly use the judge agent to evaluate the arguments presented by the proponent agent and the challenger agent and determine the winning agent for the current round of debate. An example prompt is

Prompt for Judge Agent

You are the judge in a debate. Your task is to evaluate the arguments from agents.

There are two types of agents:

1. Proponent Agent: Argue that the current retrieved content is sufficient.

2. Challenger Agent: Argue that the current

retrieved content is insufficient and propose query refinement.

Question:
{question}

Queries:
{queries in Q }

Retrieved Documents:
{retrieved results in C_Q }

Agents Arguments:
{agents arguments}

Output only the agent's name.

A.3 Prompts for Response Debate

In the Response Debate phase, the discussion is initialized based on the retrieval pool and its associated documents produced in the Retrieval Debate phase. An example of a proponent agent prompt in the Response Debate stage is

Prompt for Initializing Proponent Agent

Answer the question based on the given document.

The following are given documents:

Query 1: $q_1 \in Q$

Retrieved Documents: $C_{q_1} \in C_Q$

Query 2: $q_2 \in Q$

Retrieved Documents: $C_{q_2} \in C_Q$

...

Query n : $q_n \in Q$

Retrieved Documents: $C_{q_n} \in C_Q$

Question:
{question}

The challenger agent is initialized based on internal knowledge, and an example of a challenger agent prompt is

Prompt for Initializing Challenger Agent

Answer the question based on your own knowledge.

Question:
{question}

Factual consistency in the Response Debate phase is achieved through iterative exchanges and debates among agents. An example of the debate prompt is

Prompt for Debate

I will give the answers and arguments to this question from other agents. Use their solution as additional advice; note that they may be wrong.

Explain your answer.

{Agent's Name}:
{Other agent's response}

Question:
{question}

In the Response Debate phase, the Judge Agent aggregates the responses from all agents and selects the most probable answer as the final output. An example of a judge agent prompt is

Prompt for Judge Agent

You are a moderator in a debate competition. Your task is to determine the correct final answer based on the arguments presented by the debaters. Output only the final answer with no explanations or additional text.

{Agent's Name}:
{Agent's Response}

Question:
{question}

B More Detailed Results

Table 6 shows the detailed performance metrics of DRAG across six datasets under different debate round settings in the Retrieval Debate and Response Debate stages.

C Case Study

We show the case study of the Retrieval Debate in Figure 5, and the error case in the Response Debate stage in Figure 6.

Settings	NQ		TriviaQA		PopQA		2Wiki		HotpotQA		StrategyQA
	EM	F1	EM	F1	EM	F1	EM	F1	EM	F1	EM
DRAG _{RetIter=0}	34.40	47.61	60.40	69.60	37.60	45.43	26.00	34.30	30.20	41.00	68.80
DRAG _{RetIter=1}	36.80	50.64	60.80	69.90	37.40	45.60	29.20	37.49	31.40	42.41	69.60
DRAG _{RetIter=2}	36.80	50.32	60.76	69.93	37.50	45.83	29.60	37.84	30.80	41.85	69.40
DRAG _{RetIter=3}	36.80	50.38	60.80	69.93	37.60	46.10	28.80	36.97	30.80	41.74	69.20
DRAG _{RetIter=4}	37.00	50.58	60.80	69.91	37.60	46.10	29.00	37.33	31.00	41.86	69.40
DRAG _{ResIter=0}	38.20	50.79	57.20	65.31	39.60	47.04	16.00	25.91	28.20	37.48	64.60
DRAG _{ResIter=1}	37.60	51.77	62.40	71.47	39.20	47.52	27.80	36.91	33.40	43.19	71.00
DRAG _{ResIter=2}	38.20	51.61	60.20	68.59	38.80	46.52	28.20	36.47	31.20	41.53	68.00
DRAG _{ResIter=3}	36.80	50.38	60.80	69.93	38.60	46.50	28.80	36.97	30.80	41.74	69.20
DRAG _{ResIter=4}	37.00	50.09	60.60	70.03	38.80	46.98	27.20	34.67	31.00	41.30	67.00

Table 6: The detailed performance metrics of DRAG across six datasets under different debate round settings in the Retrieval Debate and Response Debate stages.

Q: Who is the girl playing with guns and roses?

Retrieval Debate Round 1

Opponent Agent: The current retrieved content appears to be insufficient to generate a reliable answer to the question, "Who is the girl playing with guns and roses." The content includes three unrelated documents, each with a different title and context. Doc 1 is about a 2012 film, Doc 2 is about a TV episode, and Doc 3 is about a manga series. None of the documents provide information about a girl playing with Guns N' Roses, the famous rock band. *Query Optimization: who is the girl playing with guns and roses -> who is the girl playing with Guns N' Roses*

Proponent Agent: Based on the retrieved content, I argue that the current information is sufficient to answer the question, "Who is the girl playing with guns and roses."

Judge Agent: Opponent Agent

Retrieval Debate Round 2

Opponent Agent: The current retrieved content is insufficient to generate a reliable answer to the question "Who is the girl playing with Guns N' Roses?" because it only mentions the current lineup of the band Therefore, I propose a query expansion to gather more specific information about Melissa Reese's role in Guns N' Roses and any context in which she plays with the band. *Query Expansion: What is Melissa Reese's role in Guns N' Roses, and what is her involvement with the band?*

Proponent Agent: Based on the retrieved content, it is clear that the question "Who is the girl playing with Guns N' Roses?" can be answered with confidence. Doc 1 explicitly states that *Melissa Reese* is a member of the hard rock band Guns N' Roses, and Doc 2 further confirms this, listing her as a current band member. Therefore, the answer to the question is Melissa Reese, and no further retrieval is needed.

Judge Agent: Proponent Agent

Retrieval

Query 1: Who is the girl playing with Guns N' Roses?

Retrieved Content:
 Doc 1(Title: "Melissa Reese"): *Melissa Reese* is an American musician and model best known for her collaborations with Bryan "Brain" Mantia and her current position in the hard rock band Guns N' Roses. Reese was born in Seattle, Washington. She is the youngest of three girls born into a household of Chinese, English, Spanish, Japanese, Filipino, Irish, Scottish, and Danish descent

Doc 2(Title: "Guns N' Roses"): Guns N' Roses Guns N' Roses, often abbreviated as GNR, is an American hard rock band from Los Angeles, California, formed in 1985. When they signed to Geffen Records in 1986, the band comprised vocalist Axl Rose, lead guitarist Slash, rhythm guitarist Izzy Stradlin, bassist Duff McKagan, and drummer Steven Adler. The current lineup consists of Rose, Slash, McKagan, keyboardist Dizzy Reed, guitarist Richard Fortus, drummer Frank Ferrer, and keyboardist *Melissa Reese*. Guns N' Roses' debut album, "Appetite for Destruction" (1987), reached number one on the "Billboard" 200 a year after its release, on the strength of "Sweet Child o'

Doc 3(Title: "Sweet Child o' Mine"): to the demo in a loop, Rose started saying to himself, "Where do we go? Where do we go now?" and Proffer suggested that he sing that. The "Sweet Child of Mine" video depicts the band rehearsing in the Huntington Ballroom at Huntington Beach, surrounded by crew members. All of the band members' girlfriends at the time were shown in the clip: Rose's girlfriend, Erin Everly.....

Response Debate Round 1

Proponent Agent: The song "Girl, You'll Be a Woman Soon" is not by Guns N' Roses, but the song "Paradise City" and "Welcome to the Jungle" are. However, I think you might be referring to the song "Paradise City" or "Welcome to the Jungle"

Opponent Agent: The answer is: *Melissa Reese*

Response Debate Round 2

Proponent Agent: The answer is: *Melissa Reese*

Opponent Agent: The answer is: *Melissa Reese*

Response Debate Round 3

Proponent Agent: *Melissa Reese*

Opponent Agent: *Melissa Reese*

Judge Agent: *Melissa Reese*

Figure 5: Case study of the Retrieval Debate.

Q: Where was the father of Elizabeth Of Rhuddlan born?
A: Westminster

Retrieval

Query 1: Who was the father of Elizabeth Of Rhuddlan and what was his birthplace?

Retrieved Content:
 Doc 1 (Title: "Elizabeth of Rhuddlan"): Elizabeth of Rhuddlan Elizabeth of Rhuddlan (7 August 1282 – 5 May 1316) was the eighth and youngest daughter of King Edward I and Queen Eleanor of Castile. Of all of her siblings, she was closest to her younger brother King Edward II, as they were only two years apart in age. In April 1285 there were negotiations with Floris V for Elizabeth's betrothal to his son John I,

Doc 2 (Title: "Elizabeth of Rhuddlan"): of Hereford, 3rd of Essex, also Constable of England, at [Westminster Abbey](#). The children of Elizabeth and Humphrey de Bohun, 4th Earl of Hereford were: During Christmas 1315, Elizabeth, who was pregnant with her eleventh child, was visited by her sister-in-law, Queen Isabella of France. On 5 May 1316 she went into labour, giving birth to her daughter Isabella.....

Doc 3 (Title: "Rhuddlan Castle"): Rhuddlan Castle Rhuddlan Castle () is a castle located in Rhuddlan, Denbighshire, Wales. It was erected by Edward I in 1277, following the First Welsh War. Much of the work was overseen by master mason James of Saint George. Rhuddlan, which was not completed until 1282, was built concurrently with Flint Castle,

Response Debate Round 1

Proponent Agent: To answer this question, I need to identify the father of Elizabeth of Rhuddlan. Elizabeth of Rhuddlan was the daughter of John, 1st Baron Strange of Knokyn,

The answer is: [Oxford Castle, Oxford, England](#). ❌

Opponent Agent: To answer this question, we need to identify the father of Elizabeth of Rhuddlan. According to the given documents, Elizabeth of Rhuddlan was the daughter of King Edward I and Queen Eleanor of Castile. King Edward I was born in Westminster, England. The answer is: [Westminster](#). ✅

Response Debate Round 2

Proponent Agent: Proponent Agent 0's solution suggests that Elizabeth of Rhuddlan's father is King Edward I, which is incorrect. Elizabeth of Rhuddlan was actually the daughter of John of England, 1st Baron of Rhuddlan, who was the son of King John of England. However, considering the information provided by Proponent Agent 0, we can infer that the father of Elizabeth of Rhuddlan is indeed King Edward I, but this is incorrect. To confirm, we can look at the birthplace of King Edward I, which is indeed Westminster, England. However, this does not necessarily confirm the birthplace of the father of Elizabeth of Rhuddlan. The answer is: [Westminster, England](#) (Note: This answer is incorrect due to the incorrect identification of the father of Elizabeth of Rhuddlan). ✅

Opponent Agent: The other agent's response is incorrect. According to the given documents, Elizabeth of Rhuddlan was the daughter of King Edward I and Queen Eleanor of Castile, not John of England, 1st Baron of Rhuddlan. The documents do not mention John of England, 1st Baron of Rhuddlan, as the father of Elizabeth of Rhuddlan. As I previously mentioned, King Edward I, the father of Elizabeth of Rhuddlan, was born in Westminster, England. The answer is: [Westminster](#). ✅

Response Debate Round 3

Proponent Agent: Proponent Agent 0's solution is incorrect. Elizabeth of Rhuddlan was actually the daughter of John of England, 1st Baron of Rhuddlan, who was the son of King John of England. This information is not mentioned in Proponent Agent 0's response. Proponent Agent 0's response is based on incorrect information, and it does not provide any evidence to support the claim that King Edward I is the father of Elizabeth of Rhuddlan. To answer the question correctly, we need to identify the father of Elizabeth of Rhuddlan, which is John of England, 1st Baron of Rhuddlan. As previously mentioned, King John of England, the father of John of England, 1st Baron of Rhuddlan, was born in Oxford Castle, England. The answer is: [Oxford Castle, England](#). ❌

Opponent Agent: The other agent's response is incorrect in identifying the father of Elizabeth of Rhuddlan as John of England, 1st Baron of Rhuddlan, and King John of England

However, the other agent's response does correctly identify the birthplace of King Edward I as Westminster, England. Since King Edward I is indeed the father of Elizabeth of Rhuddlan, we can confirm that the birthplace of the father of Elizabeth of Rhuddlan is indeed Westminster, England. The answer is: [Westminster](#). ✅

Judge Agent: [Oxford Castle, England](#). ❌

Figure 6: Error Case of the Response Debate.