

# RISE: Reasoning Enhancement via Iterative Self-Exploration in Multi-hop Question Answering

Anonymous ACL submission

## Abstract

Large Language Models (LLMs) excel in many areas but continue to face challenges with complex reasoning tasks, such as Multi-Hop Question Answering (MHQA). MHQA requires integrating evidence from diverse sources while managing intricate logical dependencies, often leads to errors in reasoning. Retrieval-Augmented Generation (RAG), widely employed in MHQA tasks, faces challenges in effectively filtering noisy data and retrieving all necessary evidence, thereby limiting its effectiveness in addressing MHQA challenges. To address these challenges, we propose **RISE: Reasoning Enhancement via Iterative Self-Exploration**, a novel framework designed to enhance models' reasoning capability through iterative self-exploration. Specifically, RISE involves three key steps in addressing MHQA tasks: question decomposition, retrieve-then-read, and self-critique. By leveraging continuous self-exploration, RISE identifies accurate reasoning paths, iteratively self-improving the model's capability to integrate evidence, maintain logical consistency, and enhance performance in MHQA tasks. Extensive experiments on multiple MHQA benchmarks demonstrate that RISE significantly improves reasoning accuracy and task performance.

## 1 Introduction

Large language models (LLMs) demonstrate outstanding capabilities in natural language understanding and generation (Brown et al., 2020; Zhang et al., 2022; Zeng et al., 2022; Chowdhery et al., 2023; Touvron et al., 2023). However, LLMs still face challenges with complex Multi-Hop Question Answering (MHQA) tasks. MHQA requires models to integrate evidence from multiple sources and manage intricate logical relationships. This involves both retrieving and combining various pieces of evidence and constructing coherent reasoning chains. Prompt-based methods, such as

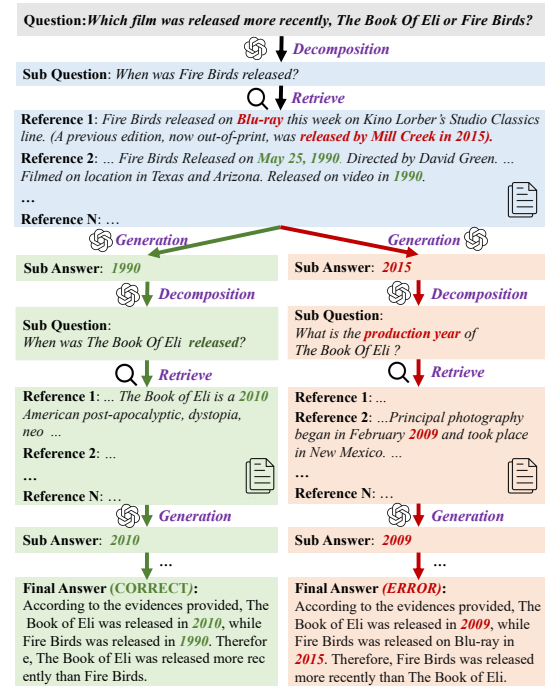


Figure 1: The upper part of the figure (blue) illustrates an Evidence Aggregation Error, where the Blu-ray release year of *Fire Birds* (2015) is mistaken for its theatrical release year. The lower part (green and red) shows a Reasoning Decomposition Error. The incorrect path formulates the sub-question as the production year of *The Book of Eli* (2009) instead of its release year (2010).

Chain-of-Thought (CoT) (Wei et al., 2022b; Wang et al., 2023a; Yu et al., 2023), are employed to address MHQA by split complex problems into smaller, thereby harnessing the reasoning potential of LLMs. However, these methods often lack external knowledge, resulting in key evidence being overlooked and generate hallucinations (Rawte et al., 2023; Ji et al., 2023; Ye et al., 2023).

Retrieval-Augmented Generation (RAG) methods (Guu et al., 2020; Lewis et al., 2020; Izacard et al., 2022; Nakano et al., 2021; Asai et al., 2023; Ma et al., 2023; Yu et al., 2024; Shi et al., 2024a) have been proposed to address the afore-

mentioned challenges. By incorporating external knowledge, RAG effectively mitigates hallucination phenomena and achieves significant results in MHQA tasks through multiple retrievals. However, RAG is constrained by the performance of the retrievers, inevitably introducing noise. Additionally, the multi-round retrieval process may lead to error propagation, resulting in two main types of errors: **Evidence Aggregation Errors** and **Reasoning Decomposition Errors**. As illustrated in Figure 1, Evidence Aggregation Errors occur when the model fails to accurately integrate evidence from multiple evidences, leading to hallucinations. Reasoning Decomposition Errors arise when problem decomposition phase generates sub-questions that do not align with original question’s intent. These issues are particularly pronounced in smaller models with weaker reasoning capabilities.

Distillation and fine-tuning (Uesato et al., 2022; Luo et al., 2023; Shridhar et al., 2023) effectively enhance the reasoning capabilities of LLMs by leveraging large-scale models or high-quality, manually annotated data to improve performance. However, biases brought by human subjective annotations may undermine the performance of fine-tuning (Casper et al., 2023; Lightman et al., 2023), and these methods are costly, requiring substantial human or computational resources. Meanwhile, self-iteration methods (Yuan et al., 2024; Wang et al., 2024; Madaan et al., 2024) demonstrate tremendous potential in complex reasoning tasks. Unlike approaches that depend on large-scale models and manual annotations, self-iteration methods enable models to generate and learn from their own data, achieving outstanding results in complex tasks such as code generation and intelligent agents (Jiang et al., 2023; Ni et al., 2024; Qiao et al., 2024). Nevertheless, research on combination self-iteration methods with RAG remains limited. The integration of these two approaches has the potential to improve performance in complex reasoning tasks and leads to cost reduction.

In this paper, we introduce an innovative framework, **RISE** (**R**easoning **E**nhancement via **I**terative **S**elf-**E**xploration), which combines the paradigms of RAG and self-iteration to address key challenges in MHQA tasks. Specifically, RISE defines three core actions: question decomposition, retrieve-then-read, and self-critique. By repeatedly executing these actions, the model autonomously explores accurate reasoning paths for problems. During this process, RISE accumulates experience

datasets for the three actions and updates the model based on this experience. Through multiple iterations, RISE significantly enhances the model’s reasoning capabilities in MHQA tasks. Experimental results demonstrate that RISE outperforms baseline methods on several MHQA benchmark datasets, strongly validating its effectiveness in solving MHQA tasks while offering lower usage costs. Our main contributions are as follows:

- We propose RISE, which combines RAG and self-iteration to address two key challenges in MHQA tasks: Evidence Aggregation Errors and Reasoning Decomposition Errors.
- We design self-exploration mechanism, converts MHQA in RAG into multi-objective optimization problem, thus improving model’s reasoning capability and reducing costs.
- We integrate self-iteration paradigm with RAG, bridging gap in applying self-iteration strategies within MHQA RAG framework.

## 2 Methods

### 2.1 Overview

In this section, we provide a concise description of **RISE**, focusing on its algorithmic process. As shown in algorithm 1, RISE begins with a seed question set  $\mathcal{Q}_0$  and an initial model  $M_0$ . The model iteratively performs self-exploration for each question  $q \in \mathcal{Q}_i$ , with details presented in Section 2.2. The exploration results are stored as historical experience  $\mathcal{H}$ . After completing the exploration for all questions, the accumulated experiences optimize the model through multi-objective training, yielding an enhanced model  $M_{i+1}$ . Subsequently,  $M_{i+1}$  expands the question set based on the previous seed questions  $\mathcal{Q}_i$ , generating  $\mathcal{Q}_{i+1}$  to initiate the next round of exploration. This self-iterative process enables RISE to continuously improve capabilities without external supervision, leveraging the model’s intrinsic potential.

### 2.2 Self-Exploration Mechanism

The self-exploration mechanism enables the model to address complex problems through iterative reasoning, comprising three core actions: question decomposition, retrieve-then-read, and self-critique. These actions collectively form a structured exploration pathway, with the resulting information collected as historical data to support the model’s self-

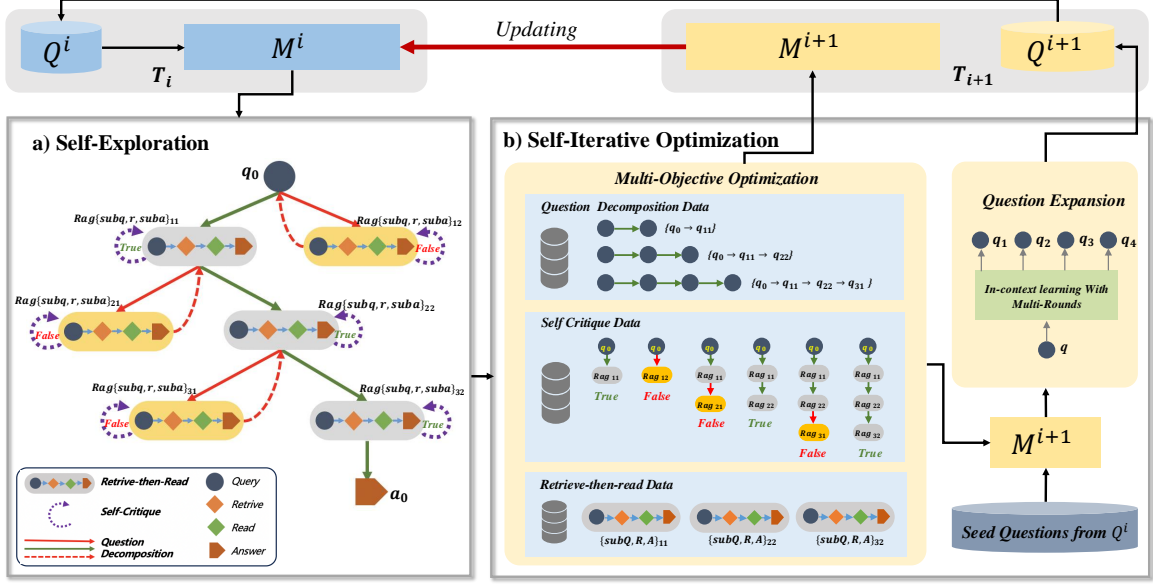


Figure 2: A complete iteration cycle in RISE. *a) Self-Exploration:* Model  $M^i$  decomposes complex questions  $q_0$  into simpler sub-questions, generates sub-answers via retrieve-then-read, and evaluates their validity, leading to a final answer  $a_0$ . Interactions are stored as historical data  $\mathcal{D}$ . *b) Iterative Optimization:* RISE optimizes model  $M^i$  using historical data  $\mathcal{D}$  to create an enhanced model  $M^{i+1}$ , which generates new questions  $Q^{i+1}$  for the next cycle.

improvement in complex problem-solving. The related prompts are provided in Appendix A.1.1.

**Question Decomposition.** In this task, the model incrementally decomposes the initial complex question into fine-grained sub-questions. At the  $t$ -th exploration node, the model uses previously explored sub-questions and answers as historical information, denoted as  $\mathcal{H} = (subq_1, suba_1), \dots, (subq_{t-1}, suba_{t-1})$ . The original question  $q_0$  is combined with  $\mathcal{H}$  and input into model  $M$  to generate the next sub-question. The model ends exploration by generating the final answer if the historical information suffices for  $q_0$ . Formally, this process is represented as Formula 1:

$$subq_t = \mathcal{F}_d(M, \mathcal{H}, q_0) \quad (1)$$

$$a_0 = M(q_0, \mathcal{H}), \quad \text{if } \mathcal{H} \text{ is sufficient.} \quad (2)$$

Additionally, all decomposition steps, including the original question and generated sub-questions, are recorded to form the dataset  $\mathcal{D}_d = \{\{q_0, \mathcal{H}, subq\}_{i=1}^{n_p}\}^{N_q}$ . By leveraging this fine-grained and structured dataset, the model learns the logical dependencies and relationships between questions and sub-questions, thereby improving its ability to decompose complex problems.

**Retrieve-then-Read.** This task follows the standard RAG paradigm to provide evidence-based an-

swers for sub-questions. At the  $t$ -th exploration node, a retriever obtains relevant fragments  $r_t$  based on the sub-question, and model  $M$  generates answer using the retrieved evidence:

$$suba_t = \mathcal{F}_g(M, subq_t^i, r_t) \quad (3)$$

Each sub-question and its answer form an exploration node  $(subq_i, suba_i)$ , added to the historical information  $\mathcal{H}_{t+1} = \mathcal{H}_t \cup \{(subq_i, suba_i)\}$ . All nodes are recorded to construct the dataset  $\mathcal{D}_r = \{(subq, r, suba)_{i=1}^{n_p}\}^{N_q}$ . Training on this dataset helps the model integrate evidence into reasoning, improving answer accuracy and reliability.

**Self-Critique.** In this task, the model’s critique capability is incorporated into the exploration process. Specifically, after completing the question decomposition and retrieve-then-read tasks at the  $t$ -th exploration node, the model  $M$  critiques the relevance and utility of the node for solving the original question and outputs a binary decision. If critiqued as True, it is retained, and exploration proceeds to the next step. If critiqued as False, the node is temporarily stored, and the process reverts to the preceding valid node to generate a new node. This process is formalized in Formula 4:

**Algorithm 1** RISE

**Input:** Seed question set  $\mathcal{Q}_0$ , Initial model  $M_0$ , Retriever  $R$ , Maximum nodes  $N_{max} = 20$

```

1: Initialize: History  $\mathcal{H} \leftarrow \emptyset$ , Model index  $i \leftarrow 0$ 
2: while True do
3:   for each question  $q \in \mathcal{Q}_i$  do
4:      $n \leftarrow 0$  ▷ Start self-exploration.
5:     while  $M_i(q, \mathcal{H}) = \text{More information}$ 
       needed and  $n < N_{max}$  do
6:        $subq \leftarrow M_i(\mathcal{H})$ 
7:        $r \leftarrow R(subq)$ 
8:        $suba \leftarrow M_i(subq, r)$ 
9:        $\sigma \leftarrow M_i(subq, suba)$ 
10:      if  $\sigma = 1$  then
11:        Add  $(subq, suba)$  to  $\mathcal{H}$ 
12:      end if
13:       $n \leftarrow n + 1$ 
14:    end while
15:     $a \leftarrow M_i(\mathcal{H})$ 
16:  end for ▷ End self-exploration.
17:   $M_{i+1} \leftarrow \text{Multi-Objective Train}(M_i, \mathcal{H})$ 
18:   $\mathcal{Q}_{i+1} \leftarrow \text{Question Expansion}(M_{i+1}, \mathcal{Q}_i)$ 
19:   $i \leftarrow i + 1$ 
20: end while

```

**Output:** Final model  $M_i$

	weights ratios				
$\alpha$	1	1	4	1	1
$\beta$	1	4	1	2	3
$\gamma$	4	1	1	3	2
Accuracy	41.55	40.32	41.94	38.71	37.63
$\alpha$	2	2	3	3	2
$\beta$	1	3	1	2	2
$\gamma$	3	1	2	1	2
Accuracy	39.78	40.32	43.01	41.13	44.27

Table 1: Performance on the 2WikiMultiHopQA dataset under varying weight ratios of  $\alpha$ ,  $\beta$ , and  $\gamma$ .

automatically expand and collect  $\mathcal{D}_d$ ,  $\mathcal{D}_r$ , and  $\mathcal{D}_c$  datasets for subsequent model training.

**Multi-Objective Optimization.** These three datasets,  $\mathcal{D}_d$ ,  $\mathcal{D}_r$ , and  $\mathcal{D}_c$ , are interconnected, with sample sizes ranging from 2k to 8k (detailed statistics are provided in Appendix 5). We believe that joint training facilitates complementary learning and enhances model capabilities. Therefore, we adopt a multi-objective optimization approach to integrate the objectives of different tasks into a unified optimization goal. The effectiveness of this approach is validated through ablation study at Section 4.3. The overall loss function is defined as follows Formula 5:

$$\mathcal{L} = \alpha\mathcal{L}_d + \beta\mathcal{L}_r + \gamma\mathcal{L}_c \quad (5)$$

Here,  $\alpha$ ,  $\beta$ , and  $\gamma$  represent the task weights, which are primarily determined based on the proportion of each task within the dataset. Meanwhile, Experimental results (see Table 1) indicate that the weights assigned to different tasks have an impact on model performance, and appropriate weight adjustments facilitate fine-grained performance optimization. Notably, to avoid potential overfitting caused by manual weight tuning which may affect the final evaluation we do not perform any fine-tuning of the task weights in our experiments. Instead, we adopt a uniform weighting strategy, assigning equal weights to all three tasks.

**Question Expansion** After completing multi-objective optimization, we use the questions generated in the previous iteration as seed data for  $M^{i+1}$  to perform question expansion, thereby acquiring training data for the next iteration. This method is inspired by (Wang et al., 2023c), leveraging multi-round in-context learning to ensure the diversity and richness of the newly generated questions. Detailed information about the question expansion prompts is provided in the appendix Figure 10.

$$\sigma_t = \mathcal{F}_c(M, subqt, subat), \quad \sigma_t \in \{0, 1\} \quad (4)$$

Similarly, we record critique historical information and then construct the dataset  $\mathcal{D}_c = \{\{\langle subq, suba \rangle, \sigma\}_{i=1}^{n_p}\}^{N_q}$ , to enhance the self-critique capabilities of the model, ensuring logical consistency and relevance within the exploration.

### 2.3 Self-Iterative Optimization

RISE is a self-iterative fine-tuning framework that optimizes the model in each training round based on data generated by the model itself, gradually enhancing its generalization ability and reasoning performance. Through a closed-loop iteration of data generation and model training, RISE effectively uncovers the model’s reasoning potential in complex tasks, driving continuous self-improvement.

**Initialization.** We initially use randomly sampled question set  $\mathcal{Q}_0$  from the training sets of three tasks—2Wiki, HotpotQA, and Musique, with 800 examples from each. Subsequently, we employ self-exploration mechanism to automatically expand and collect the three types of dataset. Subsequently, we employ self-exploration mechanism to



### 3 Experiments Setup

**Datasets:** For the main experiments, we use three QA datasets: 2WikiMultiHopQA (2WIKI) (Ho et al., 2020), HotpotQA (Hotpot) (Yang et al., 2018), and MuSiQue (MSQ) (Trivedi et al., 2022), which provide diverse reasoning challenges to evaluate the robustness of our framework. Additionally, for the analysis experiments, we include Natural Questions (NQ) (Kwiatkowski et al., 2019), Web Questions (WebQ) (Berant et al., 2013) and TriviaQA (Joshi et al., 2017) to assess the model’s performance on open-domain Question Answering tasks, further extending the evaluation scope.

**Models and Methods:** In our experiments, we use LLaMA-3.1-8B (Dubey et al., 2024) as the base model for our method in main experiments. Similarly, most of the reproduced methods are also implemented using LLaMA-3.1-8B. Additionally, based on the characteristics of MHQA tasks, we select and reproduce a variety of methods, categorized into non-retrieval-based methods and retrieval-based methods. Non-retrieval-based methods include Naive LLM (LLaMA-3.1-8B, GPT-3.5-turbo), CoT (Wei et al., 2022b), CoT-SC (Wang et al., 2023a) and GenRead (Yu et al., 2023), while the retrieval-based methods consist of Naive RAG, Self-Ask (Press et al., 2023), WebGLM (Liu et al., 2023), Self-RAG (Asai et al., 2023), RRR (Ma et al., 2023), and GenGround (Shi et al., 2024a). In the analysis experiments, we employ GPT-4o<sup>1</sup> as the evaluation model, combining subjective analysis with specific metrics to comprehensively assess model performance.

**Retrieval:** We adopt a two-stage retrieval framework (Liu et al., 2023), consisting of coarse-grained web search (via Chrome) followed by fine-grained LLM-enhanced retrieval. We consistently use the same retrieval method to reproduce results for other approaches that incorporate retrievers.

**Evaluation Metrics:** We assess performance using Accuracy (Acc), F1 score (F1), and Exact Match (EM) to evaluate QA quality. Additionally, we evaluate the quality of the reasoning chains from the perspectives of chain length and four subjective dimensions: conciseness, rationality, sequencing, and goal orientation.

We provide comprehensive experimental details in Appendix A.2, including implementation details, datasets, and other relevant information.

<sup>1</sup>We use GPT models accessed via the OpenAI API: <https://openai.com/api/>.

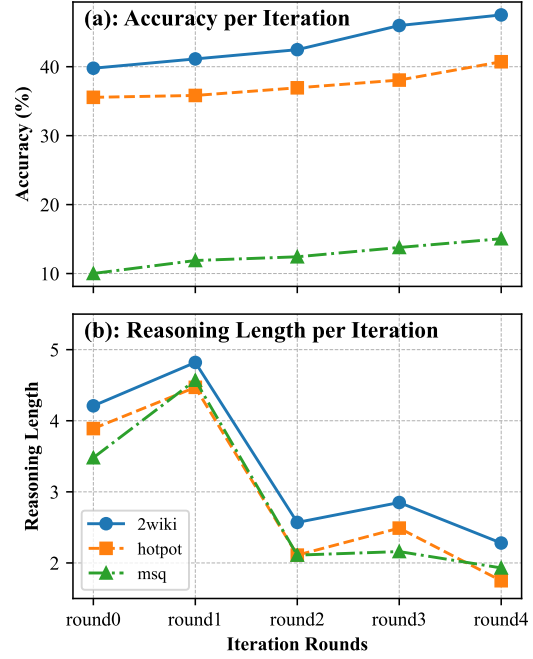


Figure 3: Changes in model accuracy (a) and reasoning length (b) across datasets. Accuracy consistently improves across datasets, while reasoning length, despite some fluctuations, shows an overall decreasing trend.

### 4 Results and Analysis

In this section, we evaluate RISE from three aspects. First, we validate effectiveness of multiround self-iterative and compare RISE with mainstream MHQA methods. Second, we conduct an in-depth analysis of the performance of question decomposition, retrieve-then-read, and self-critique using objective metrics and AI-based evaluations. Finally, we conduct ablation studies to verify the importance of different tasks in enhancing performance.

#### 4.1 Overall Performance

**RISE Outperforms Other Methods:** Table 2 presents the experimental results across three MHQA datasets and three SHQA datasets. We observe that retrieval-based enhancement is crucial for MHQA tasks. While CoT achieves relatively good performance, other non-retrieval methods generally perform worse than most RAG approaches with the same model. For the relatively simpler SHQA tasks, retrieval-based enhancement does not seem to offer significant advantages. Notably, RISE achieves outstanding results in both task types over all datasets, even surpassing GPT-3.5. Furthermore, our method excels in F1 and EM metrics, demonstrating its efficiency (additional metrics are provided in Appendix A.4).

Method	Model	MHQA			SHQA		
		2WIKI	Hotpot	MSQ	NQ	WebQ	Trival
w/o retrieval							
Naive LLM	LLaMA-3.1-8B	35.90	27.30	11.30	57.50	61.25	71.50
	GPT-3.5-turbo	47.10	41.50	19.10	57.25	58.30	80.25
CoT (Wei et al., 2022b)	LLaMA-3.1-8B	43.00	34.60	16.20	56.75	62.00	71.75
CoT-SC* (Wang et al., 2023a)	GPT-3.5-turbo	21.00	30.60	8.90	39.75	38.50	79.25
GenRead (Yu et al., 2023)	LLaMA-3.1-8B	20.00	28.40	10.30	46.25	43.00	67.00
w retrieval							
Naive RAG	LLaMA-3.1-8B	43.50	37.50	15.00	51.75	57.75	73.75
Self-Ask (Press et al., 2023)	LLaMA-3.1-8B	22.90	29.70	12.70	44.50	46.50	66.00
WebGLM (Liu et al., 2023)	LLaMA-3.1-8B	37.60	36.50	13.30	55.75	62.25	73.75
Self-RAG* (Asai et al., 2023)	LLaMA2-7B	32.00	30.20	8.00	54.50	60.25	74.25
	LLaMA2-13B	30.80	29.40	8.80	55.25	58.75	75.00
RRR (Ma et al., 2023)	LLaMA-3.1-8B	23.70	11.80	5.40	22.50	30.50	36.75
	GPT-3.5-turbo	28.20	29.70	8.70	57.75	56.50	80.50
GenGround (Shi et al., 2024a)	LLaMA-3.1-8B	37.90	36.10	17.80	48.50	44.50	75.25
RISE(Ours)	LLaMA-3.1-8B	49.40	40.50	21.70	59.50	62.50	80.25

Table 2: Comparison of RISE’s accuracy with other methods on 2WikiMultiHopQA, HotpotQA, MuSiQue, Natural Questions, Web Questions, and TriviaQA. Methods marked with an asterisk (\*) involve specific considerations: CoT-SC uses GPT-3.5 due to LLaMA-3.1’s instruction-following limitations, and Self-RAG employs public model weights as its dataset is unavailable. Other methods use LLaMA-3.1-8B.   denote *Prompting-based Methods*, while   denote *Finetuning-based Methods*. Due to space constraints, **F1** and **EM** metrics are in Appendix A.4.

**Steady Performance Improvement:** Meanwhile, as shown in Figure 3 (a) *Accuracy per Iteration*, we illustrate how the model’s accuracy evolves over four iterations on multiple datasets. The results demonstrate a consistent upward trend in accuracy with each iteration, further validating the effectiveness of our proposed self-training method in improving the model’s overall performance.

## 4.2 Analysis Experiments

**Question Decomposition Capability:** To evaluate improvement in the model’s decomposition capability for MHQA tasks, we first analyze the changes in reasoning length. As shown in Figure 3 (b) *Reasoning Length per Iteration*, accuracy steadily improves, while reasoning length initially increases and then decreases, ultimately showing downward trend. This trend reflects model’s decomposition ability progressively improves over iterations.

To further analyze changes in decomposition ability, we using GPT-4o as a judge to evaluate the model’s query decomposition across four dimensions (including conciseness, rationality, sequencing and goal orientation, see Appendix A.1.2 for more details.). As illustrated in Figure 4, we compare the performance of the model across iterations and observe newer model consistently outperforms the previous iteration. These findings demonstrate that self-training not only improves reasoning paths but also enhances the rationality of decomposition.

**Retrieve-then-Read Capability:** In MHQA tasks, models often struggle to integrate logical information from extensive evidence, especially in filtering irrelevant content. To evaluate the changes in the model’s summarization capability over iterations, we disable the decomposition functionality and instead allow model to perform single-round retrieval and direct question-answering. To ensure robustness in the experiments, we introduce relatively simpler datasets such as NQ, WebQ, and TriviaQA (Figure 5 (a) *Simple Questions*) while retaining the complex datasets from main experiments (Figure 5 (b) *Complex Questions*). The experimental results show that, as iterations progress, RISE consistently improves its performance across six datasets. This demonstrates the advantage of RISE in MHQA tasks and its effectiveness in conventional QA tasks, further validating its generalizability.

**Self-Critique Capability:** To evaluate the changes in the model’s self-critique capability, we designed a third set of experiments. In this experiment, both our model and GPT-4o assess the same set of decomposition results, with GPT-4o serving as a reference. By analyzing the consistency between our model and GPT-4o evaluations, we measure the improvement in the model’s self-critique capability. As shown in Table 3, the consistency between our model and GPT-4o steadily increases with each iteration. This indicates that the iterative




Figure 4: Evaluating the win rates between the current and previous iterations using GPT-4o to assess model’s question decomposition capability. Results indicate that each new iteration consistently outperforms the previous one in subjective effectiveness, demonstrating RISE’s continuously enhance the model’s decomposition capability.

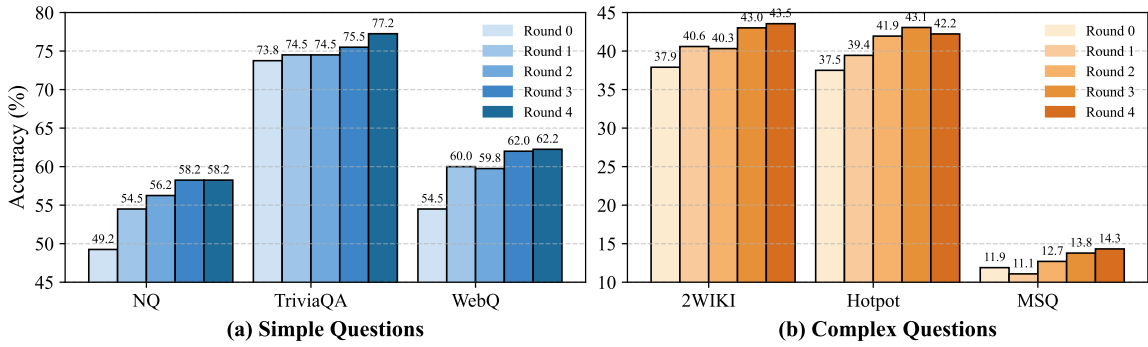


Figure 5: Changes in the model’s retrieve-then-read capability. (a) Results on simpler datasets (NQ, TriviaQA, WebQ), (b) Results on more complex datasets (2Wiki, HotpotQA, MSQ), where accuracy shows consistent growth with each iteration, even in challenging scenarios.

	Consistency with GPT-4o (%)		
	2WIKI	HotpotQA	MSQ
Round 1	74.30	64.70	60.00
Round 2	72.67	66.30	76.00
Round 3	79.67	77.33	<b>79.33</b>
Round 4	<b>80.67</b>	<b>79.33</b>	78.00

Table 3: Consistency analysis with GPT-4o on each datasets. The results show progressive improvements in consistency with GPT-4o, highlighting the model’s enhanced self-critique ability through iterative training.

process in RISE effectively enhances the model’s self-criticism capability. (For more experiment details see Appendix A.2.3.)

### 4.3 Ablation Study

To evaluate the impact of each synthesized training dataset on the model’s performance, we conduct an ablation study. As shown in Table 4, the experiment uses the same three MHQA datasets as before and the three training datasets generated in the round1,

with accuracy as the primary evaluation metric.

Removing the question decomposition dataset leads to accuracy drop of 3.5% on 2Wiki, highlighting its importance in enabling effective multi-hop reasoning. Excluding the retrieve-then-read dataset causes accuracy declines on HotpotQA (2.77%) and Musique (2.43%), highlighting the importance of this dataset in synthesizing evidence from diverse sources to mitigate the impact of noise. The removal of the self-critique dataset results in consistent accuracy reductions across all three datasets, emphasizing its pivotal function in refining reasoning paths processes. These results demonstrate the complementary and indispensable contributions of the question decomposition, retrieve-then-read, and self-critique datasets to the model’s performance.

Furthermore, we conduct separate training for the three tasks (Separate), where three LLMs are individually trained for decomposition, retrieve-then-read, and self-critique tasks. Compared to joint training (RISE), the accuracy of separate training is consistently lower across all datasets.

	2WIKI	Hotpot	MSQ
	Acc	Acc	Acc
w/o Decomp	37.63	33.89	11.08
w/o R-t-R	40.59	33.06	9.46
w/o Critique	38.98	33.89	10.27
Separate	40.86	34.72	10.54
RISE	<b>41.13</b>	<b>35.83</b>	<b>11.89</b>

Table 4: Ablation study on 2WIKI, HotpotQA, and MSQ, showing the impact of removing individual tasks (*Question Decomposition*, *Retrieve-then-Read*, and *Self-Critique*) and comparing joint training (*RISE*) with separate training (*Separate*) of individual tasks.

## 5 Related Works

**Multi-hop Question Answering:** MHQA tasks address questions that require integrating information from multiple sources and performing multi-step reasoning to produce a complete answer (Zhang et al., 2024; Li and Du, 2023). Question decomposition has been a pivotal approach for understanding and solving multi-hop questions, some works (Wei et al., 2022a; Wang et al., 2023b; Zhou et al., 2023; Shi et al., 2024b) leverage LLMs to divide complex questions into simpler single-hop sub-questions that are solved sequentially. Self-Ask (Press et al., 2023) uses LLMs to generate and resolve follow-up sub-questions with an external search engine. However, the effectiveness of these approaches depends on LLM’s inherent question decomposition capability, which constrained by hallucinations.

**Retrieval-Augmented Generation for MHQA:** RAG (Gua et al., 2020; Lewis et al., 2020; Izacard et al., 2022; Nakano et al., 2021; Asai et al., 2023; Ma et al., 2023; Yu et al., 2024; Shi et al., 2024a) integrates retrieval with generation to solve knowledge-intensive tasks (Zhu et al., 2024; Feng et al., 2024). The original RAG framework excels at single-hop QA but faces significant challenges in handling multi-hop QA and complex reasoning tasks (Lewis et al., 2020; Xu et al., 2024).

To address these challenges, various methods have been proposed. Chain of Thought (CoT) (Wei et al., 2022b) and Tree of Thought (ToT) (Yao et al., 2024) are integrated with RAG to enable multi-step reasoning and iterative retrieval (Press et al., 2023; Yao et al., 2023; Zhou et al., 2023; Khattab et al., 2023), allowing the model to incorporate a broader range of external knowledge and improve its reasoning capabilities. However, existing retrieval-augmented systems are inevitably affected by the

limitations of retrievers, often introducing irrelevant or noisy information (Yin et al., 2023; Xu et al., 2024; Ma et al., 2023). Enhancing the model’s reasoning capabilities to filter noise and focus on critical evidence is essential for accurate summaries, which our method achieves through reasoning decomposition, improving both logical reasoning and QA performance.

### Self-Improvement in Large Language Models:

Self-improvement refers to the process by which models generate and utilize their own output data to enhance performance (Zelikman et al., 2024; Singh et al., 2024; Gülçehre et al., 2023). Existing approaches, such as self-training (Du et al., 2021) and self-play (Yuan et al., 2024; Chen et al., 2024), leverage pseudo-label generation and iterative policy optimization to improve the utilization of unlabeled data and enhance decision-making capabilities. Self-Rewarding (Yuan et al., 2024) employs the LLM-as-Judge paradigm to strengthen reasoning abilities, while Self-Refine (Madaan et al., 2024) iteratively optimizes generated outputs through self-feedback mechanisms.

In complex tasks like code generation and agent-based learning, self-improvement proves effective. Methods such as Self-Evolve (Jiang et al., 2023), NExT (Ni et al., 2024), and AutoAct (Qiao et al., 2024) leverage self-feedback, self-guided tracking, and self-planning to enhance performance. However, the application of self-iterative techniques in RAG scenarios remains underexplored. Our method addresses this gap by integrating self-exploration into RAG to generate diverse training data, enabling continuous model evolution and enhancing performance in complex tasks.

## 6 Conclusion

We propose RISE, a framework that addresses two key errors in MHQA tasks: Evidence Aggregation and Reasoning Decomposition. Through self-exploration, RISE continuously enhances reasoning capabilities. Additionally, RISE integrates self-iterative paradigm with RAG framework, bridging the gap in applying self-iterative strategies to MHQA scenarios without requiring manual intervention or reliance on large models, thereby offering a cost-effective solution. Experimental results on MHQA benchmarks demonstrate significant improvements in reasoning accuracy and task performance, highlighting RISE’s robustness and adaptability in tackling complex reasoning challenges.



## Limitation

While RISE achieves strong performance in complex reasoning tasks, there remain opportunities for further enhancement. The current framework relies on external retrieval mechanisms without explicit optimization, which may limit the quality of evidence for downstream reasoning. Future work could explore self-improvement across the entire pipeline—spanning question decomposition, retrieval, generation, and reflection—to achieve more seamless integration and efficiency.

## References

Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2023. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *The Twelfth International Conference on Learning Representations*.

Jonathan Berant, Andrew Chou, Roy Frostig, and Percy Liang. 2013. Semantic parsing on freebase from question-answer pairs. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1533–1544.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language models are few-shot learners*. In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, et al. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *Transactions on Machine Learning Research*.

Zixiang Chen, Yihe Deng, Huizhuo Yuan, Kaixuan Ji, and Quanquan Gu. 2024. Self-play fine-tuning converts weak language models to strong language models. In *Forty-first International Conference on Machine Learning*.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113.

Jingfei Du, Edouard Grave, Beliz Gunel, Vishrav Chaudhary, Onur Celebi, Michael Auli, Veselin Stoyanov, and Alexis Conneau. 2021. *Self-training improves pre-training for natural language understanding*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5408–5418, Online. Association for Computational Linguistics.

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 e-prints*, pages arXiv–2407.

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.

Çağlar Gülçehre, Tom Le Paine, Srivatsan Srinivasan, Ksenia Konyushkova, Lotte Weerts, Abhishek Sharma, Aditya Siddhant, Alex Ahern, Miaosen Wang, Chenjie Gu, et al. 2023. Reinforced self-training (rest) for language modeling. *CoRR*.

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.

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.

Shuyang Jiang, Yuhao Wang, and Yu Wang. 2023. *Self-evolve: A code evolution framework via large language models*. *Preprint*, arXiv:2306.02907.

Mandar Joshi, Eunsol Choi, Daniel S Weld, and Luke Zettlemoyer. 2017. Triviaqa: A large scale distantly supervised challenge dataset for reading comprehension. *arXiv preprint arXiv:1705.03551*.

Omar Khattab, Keshav Santhanam, Xiang Lisa Li, David Hall, Percy Liang, Christopher Potts,



739	Jonathan Uesato, Nate Kushman, Ramana Kumar, Francis Song, Noah Siegel, Lisa Wang, Antonia Creswell, Geoffrey Irving, and Irina Higgins. 2022. <a href="#">Solving math word problems with process- and outcome-based feedback</a> . <i>Preprint</i> , arXiv:2211.14275.	795
740		796
741		797
742		798
743		799
744	Tianduo Wang, Shichen Li, and Wei Lu. 2024. Self-training with direct preference optimization improves chain-of-thought reasoning. In <i>Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 11917–11928.	800
745		801
746		802
747		803
748		804
749		
750	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023a. Self-consistency improves chain of thought reasoning in language models. In <i>The Eleventh International Conference on Learning Representations</i> .	805
751		806
752		807
753		808
754		
755		
756	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023b. <a href="#">Self-consistency improves chain of thought reasoning in language models</a> . In <i>The Eleventh International Conference on Learning Representations</i> .	809
757		810
758		811
759		812
760		813
761		814
762	Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023c. Self-instruct: Aligning language models with self-generated instructions. In <i>The 61st Annual Meeting Of The Association For Computational Linguistics</i> .	815
763		816
764		817
765		818
766		819
767		820
768	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022a. <a href="#">Chain-of-thought prompting elicits reasoning in large language models</a> . In <i>Advances in Neural Information Processing Systems</i> , volume 35, pages 24824–24837. Curran Associates, Inc.	821
769		822
770		823
771		824
772		825
773		826
774		
775	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022b. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	827
776		828
777		829
778		830
779		831
780	Shicheng Xu, Liang Pang, Huawei Shen, Xueqi Cheng, and Tat-Seng Chua. 2024. <a href="#">Search-in-the-chain: Interactively enhancing large language models with search for knowledge-intensive tasks</a> . In <i>Proceedings of the ACM Web Conference 2024</i> , WWW '24, page 1362–1373, New York, NY, USA. Association for Computing Machinery.	832
781		833
782		834
783		835
784		
785		
786		
787	Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. <a href="#">HotpotQA: A dataset for diverse, explainable multi-hop question answering</a> . In <i>Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing</i> , pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.	836
788		837
789		838
790		839
791		840
792		
793		
794		
	Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. Tree of thoughts: Deliberate problem solving with large language models. <i>Advances in Neural Information Processing Systems</i> , 36.	841
		842
		843
		844
		845
		846
		847
		848
	Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In <i>The Eleventh International Conference on Learning Representations</i> .	
	Hongbin Ye, Tong Liu, Aijia Zhang, Wei Hua, and Weiqiang Jia. 2023. Cognitive mirage: A review of hallucinations in large language models. <i>arXiv preprint arXiv:2309.06794</i> .	
	Xunjian Yin, Baizhou Huang, and Xiaojun Wan. 2023. <a href="#">ALCUNA: Large language models meet new knowledge</a> . In <i>Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing</i> , pages 1397–1414, Singapore. Association for Computational Linguistics.	
	Wenhao Yu, Dan Iter, Shuohang Wang, Yichong Xu, Mingxuan Ju, S Sanyal, Chenguang Zhu, Michael Zeng, and Meng Jiang. 2023. Generate rather than retrieve: Large language models are strong context generators. In <i>International Conference on Learning Representations</i> .	
	Yue Yu, Wei Ping, Zihan Liu, Boxin Wang, Jiaxuan You, Chao Zhang, Mohammad Shoeybi, and Bryan Catanzaro. 2024. Rankrag: Unifying context ranking with retrieval-augmented generation in llms. In <i>The Thirty-eighth Annual Conference on Neural Information Processing Systems</i> .	
	Weizhe Yuan, Richard Yuanzhe Pang, Kyunghyun Cho, Xian Li, Sainbayar Sukhbaatar, Jing Xu, and Jason E Weston. 2024. Self-rewarding language models. In <i>Forty-first International Conference on Machine Learning</i> .	
	Eric Zelikman, Eliana Lorch, Lester Mackey, and Adam Tauman Kalai. 2024. Self-taught optimizer (stop): Recursively self-improving code generation. In <i>OPT 2023: Optimization for Machine Learning</i> .	
	Aohan Zeng, Xiao Liu, Zhengxiao Du, Zihan Wang, Hanyu Lai, Ming Ding, Zhuoyi Yang, Yifan Xu, Wendi Zheng, Xiao Xia, et al. 2022. Glm-130b: An open bilingual pre-trained model. In <i>The Eleventh International Conference on Learning Representations</i> .	
	Jiahao Zhang, Haiyang Zhang, Dongmei Zhang, Liu Yong, and Shen Huang. 2024. <a href="#">End-to-end beam retrieval for multi-hop question answering</a> . In <i>Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 1718–1731, Mexico City, Mexico. Association for Computational Linguistics.	

Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. 2023. [Least-to-most prompting enables complex reasoning in large language models](#). In *The Eleventh International Conference on Learning Representations*.

Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Haonan Chen, Zheng Liu, Zhicheng Dou, and Ji-Rong Wen. 2024. [Large language models for information retrieval: A survey](#). *Preprint*, arXiv:2308.07107.



## A Appendix

### A.1 Prompts

#### A.1.1 Self-Exploration Prompts

We designed detailed prompts for the three tasks in the self-exploration phase: question decomposition (Figure 6), retrieve-then-read (Figure 7), and self-critique (Figure 8). The examples used in the decomposition prompt are inspired by self-ask (Press et al., 2023).

#### A.1.2 Self-Decomposition Evaluation Prompt

In this paper, the evaluation of the question decomposition capability is conducted using GPT-4o with prompt as shown in Figure 9. The analysis involves assessing and scoring the decomposition results of different iterations across multiple dimensions, ultimately leading to a comparative analysis of the two models. The dimensions of the analysis include:

- **Conciseness:** Whether the decomposition avoids redundancy while ensuring comprehensiveness.
- **Rationality:** Whether the decomposed sub-problems are closely related to the original problem.
- **Sequencing:** Whether the decomposition of sub-problems follows a logical order and facilitates the problem-solving process.
- **Goal Orientation:** Whether the decomposition is clearly centered around addressing the main problem’s objective. Are the sub-problems closely aligned with the core goal of the main problem? Does it avoid redundant issues that deviate from the primary objective?

### A.2 Experiment detail

#### A.2.1 Implementation Details

We conduct all experiments on a server equipped with four NVIDIA A800 80G GPUs. For the experimental setup, we use the following hyperparameters: learning rate of  $1 \times 10^{-4}$ , batch size of 64, and cut-off length of 8192. Furthermore, for the weighting parameters  $\alpha$ ,  $\beta$ , and  $\gamma$  in the overall loss function, values of 1 are uniformly adopted in this research.

Datasets	$\mathcal{D}_d$	$\mathcal{D}_r$	$\mathcal{D}_c$
Round 1	3276	2501	3925
Round 2	8309	6311	8074
Round 3	4858	2106	2312
Round 4	6913	4759	5307

Table 5: Number of samples accumulated in datasets  $\mathcal{D}_d$ ,  $\mathcal{D}_r$ , and  $\mathcal{D}_c$  after each round of self-exploration.

	2WIKI	HotpotQA	MSQ
Round 1	223	194	180
Round 2	218	199	228
Round 3	239	232	238
Round 4	242	238	234
Total	300	300	300

Table 6: Number of instances in each round’s self-critique capability evaluation that aligned with GPT-4o.

#### A.2.2 Datasets

The cold-start dataset  $Q^0$  consists of 800 randomly sampled instances from the training sets of 2WikiMultiHopQA, HotpotQA, and MuSiQue, totaling 2,400 cold-start samples. Table 5 provides detailed information on the training datasets constructed during each round of self-exploration. The evaluation datasets we used 2WikiMultiHopQA, HotpotQA, and MuSiQue each contain 1,000 samples, Nature Questions, Web Questions, and TriviaQA each contain 400 samples.

#### A.2.3 Self-Critique Capability Experiments Details

To demonstrate the improvement in the self-critique capability of the model across iterations, we sampled 300 instances from the generated  $\mathcal{D}_c$  at each round and compared them with GPT-4o. The responses from GPT-4o were used as ground truth to calculate the self-critique accuracy of our model. In Table 6, we present the number of instances in each round’s self-critique capability evaluation that aligned with GPT-4o.

### A.3 Additional Experiments

#### A.3.1 RISE Robustness

To further verify the robustness of our experimental conclusions, we conducted additional experiments using the Qwen2.5-7B model. Specifically, we performed four rounds of self-exploration following the same experimental setup as in our original RISE framework. The results consistently demonstrate the effectiveness of RISE, with performance

	2WIKI			HotpotQA			MSQ		
	Acc	F1	EM	Acc	F1	EM	Acc	F1	EM
Round 0	38.98	23.28	12.10	31.94	21.72	10.28	7.30	8.98	2.70
Round 1	45.97	46.49	32.50	35.83	<b>44.55</b>	<b>32.50</b>	10.54	18.65	8.92
Round 2	46.24	46.33	30.83	36.11	41.58	30.00	11.89	16.19	8.92
Round 3	<b>47.31</b>	<b>47.10</b>	<b>34.41</b>	<b>38.33</b>	42.82	31.39	<b>13.24</b>	<b>20.70</b>	<b>10.81</b>

Table 7: Performance of the Qwen2.5-7B model after four rounds of self-exploration on different datasets, showing improvements in accuracy, F1, and EM scores across 2WIKI, HotpotQA, and MSQ.

Methods	Average Input Token	Average Output Tokens	Number of LLMs Calls
Naive LLM	18	244	1
CoT	24	284	1
CoT-SC	240	2840	10
GenRead	307	233	2
Naive RAG	440	104	1
Self-Ask	782	172	2
WebGLM	436	103	1
Self-RAG	866	294	2
RRR	496	86	2
GenGround	2449	167	5
RISE	2881	192	5

Table 8: Token consumption comparison between RISE and other baseline methods, showing average input tokens, average output tokens, and the number of LLMs calls required for each approach. RISE demonstrates a higher input token consumption due to its multi-step reasoning, but maintains efficient reasoning performance.

improvements observed across multiple datasets after each iteration. This confirms that RISE maintains strong generalization capabilities and stable performance even when applied to different large language models.

### A.3.2 Token Consumption Details

In addition to performance evaluation, we analyzed the token consumption of RISE compared to other baseline methods. We measured both the average input token consumption and the average output token length, as well as the number of model calls required in each approach. The results reveal that while RISE consumes more input tokens due to its multi-step reasoning process, it achieves higher efficiency in output generation and overall reasoning effectiveness. This analysis highlights the trade-off between token usage and model performance, demonstrating that RISE achieves a balanced optimization in complex reasoning tasks.

## A.4 Supply Metrics of Main Results

This section provides additional details to supplement the main results, including comprehensive Exact Match (EM) and F1 scores across six QA datasets: 2WikiMultiHopQA, HotpotQA, MuSiQue, Natural Questions, Web Questions, and TriviaQA. We compare RISE (Ours) with both prompting-based and fine-tuning-based methods, under settings with and without retrieval. The results offer a deeper understanding of RISE’s performance, highlighting its consistent improvements over baseline models.

Method	Model	MHQA			SHQA		
		2WIKI	Hotpot	MSQ	NQ	WebQ	Trival
w/o retrieval							
Naive LLM	LLaMA-3.1-8B	0.00	0.30	0.00	0.25	0.00	2.25
	GPT-3.5-turbo	0.50	5.50	0.20	2.50	0.75	26.75
CoT (Wei et al., 2022b)	LLaMA-3.1-8B	0.00	0.00	0.00	0.00	0.00	0.25
CoT-SC* (Wang et al., 2023a)	GPT-3.5-turbo	18.40	31.90	9.60	30.75	22.75	71.75
GenRead (Yu et al., 2023)	LLaMA-3.1-8B	17.00	25.60	7.60	35.75	25.50	61.75
w retrieval							
Naive RAG	LLaMA-3.1-8B	0.30	1.70	0.20	0.00	0.00	2.25
Self-Ask (Press et al., 2023)	LLaMA-3.1-8B	17.00	25.60	9.60	28.50	21.50	58.75
WebGLM (Liu et al., 2023)	LLaMA-3.1-8B	0.00	0.30	0.00	0.00	0.00	0.00
Self-RAG* (Asai et al., 2023)	LLaMA2-7B	2.90	3.90	0.80	0.00	0.00	2.00
	LLaMA2-13B	3.40	2.50	0.40	0.00	0.00	4.50
RRR (Ma et al., 2023)	LLaMA-3.1-8B	0.00	0.40	0.00	0.75	0.75	2.75
	GPT-3.5-turbo	3.20	2.20	0.30	2.50	2.00	20.00
GenGround (Shi et al., 2024a)	LLaMA-3.1-8B	23.50	24.30	10.20	20.50	18.50	60.25
RISE(Ours)	LLaMA-3.1-8B	32.70	30.50	11.70	28.50	19.25	59.50

Table 9: **EM** metrics of RISE with other methods on the 2WikiMultiHopQA, HotpotQA, MuSiQue, Natural Questions, Web Questions and TriviaQA. Methods marked with asterisk (\*) involve specific considerations: CoT-SC uses GPT-3.5 due to LLaMA-3.1’s limitations in adhering to instructions, and Self-RAG employs publicly released model weights because its dataset is unavailable. All other methods are reproduced with LLaMA-3.1-8B. ■ represent *Prompting-based Methods*, while ■ represent *Finetuning-based Methods*.

Method	Model	MHQA			SHQA		
		2WIKI	Hotpot	MSQ	NQ	WebQ	Trival
w/o retrieval							
Naive LLM	LLaMA-3.1-8B	7.99	5.49	2.67	4.09	5.46	17.55
	GPT-3.5-turbo	13.38	17.66	8.10	16.53	19.14	44.11
CoT (Wei et al., 2022b)	LLaMA-3.1-8B	2.95	2.31	1.53	1.77	2.19	4.93
CoT-SC* (Wang et al., 2023a)	GPT-3.5-turbo	24.31	39.59	15.36	40.98	36.54	80.65
GenRead (Yu et al., 2023)	LLaMA-3.1-8B	22.39	34.51	15.26	47.90	42.55	69.59
w retrieval							
Naive RAG	LLaMA-3.1-8B	6.39	9.60	4.02	9.02	9.72	19.37
Self-Ask (Press et al., 2023)	LLaMA-3.1-8B	23.42	36.14	19.67	40.35	39.61	67.02
WebGLM (Liu et al., 2023)	LLaMA-3.1-8B	8.27	6.03	3.76	5.57	6.68	9.30
Self-RAG* (Asai et al., 2023)	LLaMA2-7B	17.38	15.44	6.02	25.56	21.45	11.34
	LLaMA2-13B	14.82	13.41	6.70	15.65	9.34	15.34
RRR (Ma et al., 2023)	LLaMA-3.1-8B	5.35	3.21	1.48	4.46	6.68	9.56
	GPT-3.5-turbo	13.70	16.76	6.08	22.80	23.16	43.68
GenGround (Shi et al., 2024a)	LLaMA-3.1-8B	34.33	34.53	19.81	36.29	28.93	67.20
RISE(Ours)	LLaMA-3.1-8B	43.01	42.00	22.87	42.97	36.22	72.70

Table 10: **F1** metrics of RISE with other methods on the 2WikiMultiHopQA, HotpotQA, MuSiQue, Natural Questions, Web Questions and TriviaQA. Methods marked with asterisk (\*) involve specific considerations: CoT-SC uses GPT-3.5 due to LLaMA-3.1’s limitations in adhering to instructions, and Self-RAG employs publicly released model weights because its dataset is unavailable. All other methods are reproduced with LLaMA-3.1-8B. ■ represent *Prompting-based Methods*, while ■ represent *Finetuning-based Methods*.

### Decomposition Prompt

Instruction: Please answer the following questions according to the given format. Strictly follow each format specification, as this will ensure consistency and clarity in your response.

- Only add follow-up questions if additional details are needed to arrive at the final answer.
- For each follow-up question, use exactly this format: 'Follow up: question'
- Ensure each follow-up question is direct and structured to be easily searchable, focusing on key information for efficient search engine retrieval.
- For each answer to a follow-up question, use exactly this format: 'Intermediate answer: answer'
- Do not repeat or alter any previously generated follow-up questions or intermediate answers.
- Conclude with the final answer using this exact format: 'So the final answer is: final answer' if no further questions are needed.

Use the examples below to understand the expected structure, and follow this format without deviating from these instructions.

Question: Who lived longer, Muhammad Ali or Alan Turing?  
Are follow up questions needed here: Yes.  
Follow up: How old was Muhammad Ali when he died?  
Intermediate answer: Muhammad Ali was 74 years old when he died.  
Follow up: How old was Alan Turing when he died?  
Intermediate answer: Alan Turing was 41 years old when he died.  
So the final answer is: Muhammad Ali.

Question: When was the founder of craigslist born?  
Are follow up questions needed here: Yes.  
Follow up: Who was the founder of craigslist?  
Intermediate answer: Craigslist was founded by Craig Newmark.  
Follow up: When was Craig Newmark born?  
Intermediate answer: Craig Newmark was born on December 6, 1952.  
So the final answer is: December 6, 1952.

... ..

Now, **\*\*continue the response\*\*** using the following question and information provided below.  
Only add follow-up questions if necessary to reach the final answer.  
**\*\*Ensure all follow-up questions are optimized for search engine queries, making each question concise, direct, and easily searchable. Avoid modifying or repeating any existing content.\*\***

Question (ORIGINAL): {question}  
Are follow up questions needed here: Yes.

Figure 6: Question Decomposition prompt template.



### Retrieve-then-Read Prompt

#Question-Answering-in-Reference-Task#

Instruction:

- Use the references provided to answer the question as specifically and completely as possible.
- If the references do not directly answer the question, combine relevant information from multiple references to create a well-supported answer.
- When references are Null or insufficient, use your own knowledge to provide a clear and relevant answer.
- When a direct answer cannot be determined, list any information in the references that could be relevant or provide partial insights related to the question. Avoid responses such as 'I don't know' or 'more information is needed.'
- Always prioritize specificity and relevance in your answer, providing helpful context or details that approach a complete answer.

Reference [1]

Reference [2]

...

Question: {question}

Figure 7: Retrieve-then-Read prompt template.

### Self-Critique Prompt

Main Question: {question}

Below is a list of previously generated subquestions and their intermediate answers, created as part of a multi-step reasoning process to answer the main question.

Your task is to evaluate whether the information in the current subquestion is necessary and contributes incrementally towards solving the main question.

Previously generated subquestions and answers:

{previous subquestions}

Current subquestion and answer candidate:

{subquestion and intermediate answer }

Instruction:

- Step 1: Check for Redundancy. Check if the current subquestion or answer repeats information already provided in previous subquestions. If it does, return 'flag = False' as this information is redundant.
- Step 2: Assess Relevance. If the information is not a duplicate, analyze its relevance to the main question. Determine whether it provides new, relevant information that helps move closer to solving the main question, even if it only provides indirect context or background. Note that information does not need to directly answer the main question to be considered relevant; it can also support understanding or provide necessary context. Mark it as 'flag = True'.
- Step 3: Based on your analysis, provide a final judgment in the following format:

**\*\*Final Judgment\*\***: [flag = True or flag = False]

Examples:

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

- Follow up: "How old was Muhammad Ali when he died?" (Flag = True, relevant for lifespan comparison.)
- Follow up: "How old was Alan Turing when he died?" (Flag = True, completes lifespan comparison.)
- Redundant Example: "How old was Muhammad Ali when he passed?" (Flag = False, redundant with earlier subquestion.)

Main Question: "Are both the directors of Jaws and Casino Royale from the same country?"

- Follow up: "Who directed Jaws?" (Flag = True, needed for director identification.)
- Follow up: "Where is Steven Spielberg from?" (Flag = True, relevant to nationality check.)
- Irrelevant Example: "What is Steven Spielberg's favorite genre?" (Flag = False, not relevant to nationality.)

Reminder: Use "flag = True" for any subquestion that provides useful information or context toward solving the main question, even if indirectly. Set "flag = False" only if it is redundant or entirely irrelevant.

Figure 8: Self-Critique prompt template.

### Self-Decomposition Evaluation Prompt

You are given two problem decomposition results for the same complex problem. Your task is to compare these results from Conciseness, Rationality, Sequencing and Goal Orientation. Analyze the two decomposition results using the criteria above. Clearly explain which approach is more effective for solving the problem and why, while highlighting the strengths and weaknesses of each approach in detail.

# Scoring Criteria:

- Score each dimension on a scale of 1-5, where:
- 1: Poor
- 2: Needs Improvement
- 3: Average
- 4: Good
- 5: Excellent

# The output follows the format below. Do not add any additional text: {

"Conciseness": {

"Result 1 Score": X,

"Result 2 Score": Y,

"Explanation": "How effectively does each decomposition avoid unnecessary complexity while still addressing all relevant aspects of the problem? Is the explanation clear and straightforward?"

},

"Rationality": {

"Result 1 Score": X,

"Result 2 Score": Y,

"Explanation": "Are the identified components logical and directly related to the problem? Do the solutions align well with the identified components?"

},

"Sequencing": {

"Result 1 Score": X,

"Result 2 Score": Y,

"Explanation": "Is the order of steps or components logical and easy to follow? Does the sequence facilitate efficient problem-solving?"

},

"Goal Orientation": {

"Result 1 Score": X,

"Result 2 Score": Y,

"Explanation": "Do the sub-questions stay aligned with the core goal of the main problem? Are there any redundant sub-questions that deviate from the primary objective?"

},

"Result": "Decomposition Results 1 Decomposition Results 2 Tie"

}

# Problem:

{problem} # Decomposition Results to Compare:

- Decomposition Results 1:

{result1}

- Decomposition Results 2:

{result2} # Output:

Figure 9: GPT-4o decomposition prompt template.

### Multi-Hop Question Generation Prompt

#Multi-Hop-Question-Generation-in-{Task}#

Instruction:

- You are an AI assistant tasked with generating multi-hop questions similar to those in the {task} dataset.
- These questions require combining multiple pieces of information to reach the answer.
- Typically, these questions involve indirect references and nested relationships.

Examples from the {task} dataset:

Example 1: Question: [Insert Example Here]

Example 2: Question: [Insert Example Here]

Example 3: Question: [Insert Example Here]

Example 4: Question: [Insert Example Here]

Example 5: Question: [Insert Example Here]

Example 6: Question: [Insert Example Here]

Example 7: Question: [Insert Example Here]

Example 8: Question: [Insert Example Here]

Please generate a **\*\*new and unique multi-hop question\*\*** that meets the following criteria:

- **\*\*Requires reasoning across multiple facts or entities\*\***.
- **\*\*Asks for only one piece of information or answer\*\*** without multiple sub-questions.
- **\*\*Is a single, coherent question\*\*** that requires multi-step reasoning to answer.
- **\*\*Includes indirect references or nested relationships\*\***.
- **\*\*Matches the complexity and structure\*\*** of questions in the {task} dataset.
- **\*\*Is concise\*\*** (one sentence) and clearly worded.
- **\*\*Covers new topics\*\*** or involves different entities from the examples above.
- **\*\*Avoids being a duplicate\*\*** of the examples above.
- **\*\*Avoids using conjunctions like 'and', 'or', or commas\*\*** to ask more than one thing.

Your response should **\*\*only contain the generated question\*\*** without any extra text and follow the format:

Question: {question}

Figure 10: Multi-hop question generation prompt template.