

MMRAG-RFT: Two-stage Reinforcement Fine-tuning for Explainable Multi-modal Retrieval-augmented Generation

Shengwei Zhao^{1*}, Jingwen Yao^{1*}, Sitong Wei¹, Linhai Xu¹, Yuying Liu¹, Dong Zhang¹, Zhiqiang Tian², Shaoyi Du^{1†}

¹National Key Laboratory of Human-Machine Hybrid Augmented Intelligence, National Engineering Research Center for Visual Information and Applications, and Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University

²School of Software Engineering, Xi'an Jiaotong University

{zhaoshengwei, yaoooo, weist0422}@stu.xjtu.edu.cn, xlh@mail.xjtu.edu.cn, yuyingliu@xjtu.edu.cn, dongzhangcv@gmail.com, {zhiqiangtian, dushaoyi}@xjtu.edu.cn

Abstract

Multi-modal Retrieval-Augmented Generation (MMRAG) enables highly credible generation by integrating external multi-modal knowledge, thus demonstrating impressive performance in complex multi-modal scenarios. However, existing MMRAG methods fail to clarify the reasoning logic behind retrieval and response generation, which limits the explainability of the results. To address this gap, we propose to introduce reinforcement learning into multi-modal retrieval-augmented generation, enhancing the reasoning capabilities of multi-modal large language models through a two-stage reinforcement fine-tuning framework to achieve explainable multi-modal retrieval-augmented generation. Specifically, in the first stage, rule-based reinforcement fine-tuning is employed to perform coarse-grained point-wise ranking of multi-modal documents, effectively filtering out those that are significantly irrelevant. In the second stage, reasoning-based reinforcement fine-tuning is utilized to jointly optimize fine-grained list-wise ranking and answer generation, guiding multi-modal large language models to output explainable reasoning logic in the MMRAG process. Our method achieves state-of-the-art results on WebQA and MultimodalQA, two benchmark datasets for multi-modal retrieval-augmented generation, and its effectiveness is validated through comprehensive ablation experiments.

Introduction

In recent years, multi-modal large language models have attracted extensive attention due to their comprehensive understanding capabilities and powerful interaction abilities in complex multi-modal scenarios. Existing multi-modal large language models, such as GPT-4V (Achiam et al. 2023) and LLaVA (Liu et al. 2023), have achieved impressive performance in tasks like visual question answering, image captioning, and open-vocabulary recognition by implicitly storing world knowledge in billions of parameters of neural networks. However, similar to large language models, multi-modal large language models still face challenges in updating world knowledge and eliminating hallucinations (Ye and

*These authors contributed equally.

†Corresponding author.

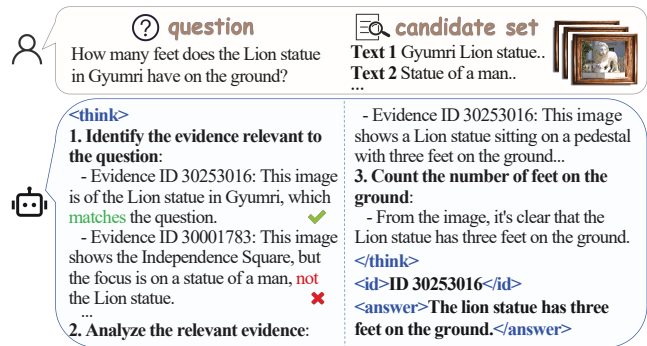


Figure 1: An example of explainable multi-modal retrieval-augmented generation. The reinforcement fine-tuned model obtains the most relevant document and the final answer based on clear reasoning logic.

Durrett 2022). As a solution, Retrieval-Augmented Generation (RAG) integrates external knowledge references to update the latest knowledge and ensure high-credibility generation (Guu et al. 2020; Lewis et al. 2020; Feng et al. 2025). Recent studies (Chen et al. 2022; Yang et al. 2023a,b; Bai, Grant, and Wang 2025) have proposed applying Retrieval-Augmented Generation to the multi-modal domain. The question-related reference documents retrieved from external multi-modal sources are utilized to assist in generating more accurate responses.

However, existing multi-modal retrieval-augmented generation methods (MMRAG) only stop at presenting the retrieval results and question responses, and fail to clarify the reasoning logic of retrieving specific documents and outputting corresponding responses accordingly, which is of great importance for the explainability of the generation results. To this end, a crucial question emerges: How can we conduct multi-modal explainable reasoning in the process of multi-modal retrieval-augmented generation as shown in Figure 1? To address this issue, a straightforward approach is to perform supervised fine-tuning (SFT) on multi-modal large language models based on the data with the labeled

multi-modal chain-of-thought reasoning process. Nevertheless, limited by the high cost of data annotation, in the field of multi-modal retrieval-augmented generation, it is extremely difficult to obtain high-quality multi-modal chain-of-thought reasoning data. Therefore, there is an urgent need to explore the potential of multi-modal retrieval-augmented generation to develop reasoning capabilities in the absence of chain-of-thought supervision data.

Fortunately, the latest advancements in the field of natural language processing (NLP), such as DeepSeek-R1 (Guo et al. 2025), have revealed that even without supervised fine-tuning, large-scale reinforcement learning can be employed to significantly enhance the reasoning capabilities of large language models, especially in the fields of science (e.g., mathematics) and code generation. Inspired by the remarkable success of reinforcement learning in the field of NLP, We integrate reinforcement learning into multi-modal retrieval-augmented generation, where the reasoning capabilities of multi-modal large language models are enhanced via a two-stage reinforcement fine-tuning framework, thus facilitating explainable multi-modal retrieval-augmented generation. Specifically, considering the limited context length of multi-modal large models, point-wise ranking is more efficient than list-wise ranking in large-scale document ranking. In the first stage, rule-based reinforcement fine-tuning is first employed to guide the multi-modal large language model in performing coarse-grained point-wise multi-modal document ranking, thereby initially filtering out multi-modal documents that are significantly irrelevant to the query. Meanwhile, since point-wise ranking only considers the relevance score of individual documents and ignores the relationships between them, its discriminative ability is limited. Therefore, in the second stage, reasoning-based reinforcement fine-tuning is utilized to jointly optimize fine-grained list-wise ranking and the answer generation, based on the multi-modal document list after coarse-grained filtering. Among them, carefully designed preset prompts are leveraged to steer the multi-modal large language model to explore the multi-modal chain of thought, thereby explaining the justifications for fine-grained list-wise ranking and the reasoning logic for answering questions based on the most relevant documents.

Furthermore, we have noticed that existing multi-modal retrieval-augmented datasets suffer from issues such as annotation errors and uneven category distribution. Therefore, we have constructed the Mini-WebQA training dataset containing 5,000 high-quality samples. To demonstrate the effectiveness of our proposed approach, we conduct extensive experiments on multiple multi-modal retrieval-augmented generation datasets, including WebQA dataset (Chang et al. 2022), MultimodalQA dataset (Talmor et al. 2021), and the Mini-WebQA dataset we constructed. Compared with previous methods, our proposed approach in multi-modal retrieval-augmented generation tasks achieves state-of-the-art results, and even yields highly competitive results when trained solely on the Mini-WebQA dataset. Our main contributions can be summarized as follows:

- We integrate reinforcement learning into multi-modal

retrieval-augmented generation, enabling multi-modal large language models to enhance their reasoning capabilities in the absence of chain-of-thought supervision data and thereby facilitating explainable multi-modal retrieval-augmented generation.

- We designed a two-stage reinforcement fine-tuning framework for multi-modal large language models, which first performs coarse-grained document filtering via efficient point-wise ranking, then conducts joint optimization of fine-grained list-wise ranking and answer generation, effectively balancing retrieval efficiency and response quality while enabling explainable reasoning.
- Our method achieves state-of-the-art results on WebQA and MultimodalQA, two benchmark datasets for multi-modal retrieval-augmented generation, with comprehensive ablation experiments validating its effectiveness.

Related Work

Multi-Modal Retrieval-Augmented Generation

In recent years, multi-modal retrieval-augmented generation has attracted increasing attention, as it boosts open-world question answering by incorporating relevant multi-modal information from various knowledge sources. Early methods build structured knowledge representations via encoder models to select multi-modal documents, then use questions and retrieved documents jointly for decoder generation. Examples include MuRAG (Chen et al. 2022), which integrates multi-modal encoders for retrieval-enhanced encoding, SKURG (Yang et al. 2023a), which performs multi-modal multi-hop QA via entity-centric fusion encoder, and PERQA (Yang et al. 2023b), which selects key documents through iterative retrieval. However, these studies overlook the potential of multi-modal large models in document understanding and ranking. Recently, RAMQA (Bai, Grant, and Wang 2025) applies SFT to guide multi-modal large language models in point-wise ranking, and VLMT (Lim et al. 2025) constructs a multi-modal large model for retrieval-augmented generation via supervised fine-tuning. Though they perform well, both ignore reinforcement learning’s role in enhancing reasoning ability and the importance of explainable reasoning. Thus, we designed a two-stage reinforcement fine-tuning framework to unlock the potential of multi-modal large models and enable explainable reasoning.

Large Language Model Reasoning via Reinforcement Learning

Recent years have seen breakthroughs in using reinforcement learning (RL) to enhance the reasoning capabilities of large language models (LLMs): OpenAI-O1 (Jaech et al. 2024), the first reasoning-focused LLM trained via large-scale RL, has achieved state-of-the-art performance across various benchmarks. Deepseek-R1-Zero (Guo et al. 2025) further demonstrated that pure RL frameworks can unlock the inherent reasoning potential of large language models without supervised training data, leveraging the Group Relative Policy Optimization (GRPO) algorithm (Shao et al. 2024) combined with a rule-based reward mechanism.

While remarkable achievements have been made in text reasoning, the potential of reinforcement learning in multi-modal reasoning also attracts significant attention. Driven by the rapid development of the multimodal field in recent years (Zhao et al. 2023), emerging studies (Shen et al. 2025; Liu et al. 2025) have extended reinforcement fine-tuning to large vision-language models, empowering them in cross-modal tasks like mathematical reasoning and open-vocabulary object detection. However, multi-modal retrieval-augmented generation based on reinforcement learning has not received specific attention in the above studies, which is precisely the core focus of our research.

Methodology

In this section, we elaborate on our two-stage reinforcement fine-tuning framework, whose overall structure is shown in Fig.2. We first introduce the problem formulation of multi-modal retrieval-augmented generation. Then, as the first stage of the framework, we specify the details of how to perform coarse-grained point-wise document ranking. Subsequently, as the second stage, we describe in detail how to jointly optimize fine-grained list-wise document ranking and answer generation based on the document list filtered by coarse-grained ranking. Finally, we present the construction process of the Mini-WebQA dataset.

Problem Formulation

In this work, we aim to address the task of multi-modal retrieval-augmented generation, which is to retrieve multi-modal documents relevant to the question and generate the answer based on the retrieved documents. Formally, suppose we have a set of N training quadruplets, denoted as $\xi = \{(Q^q, D^q, D^t, A^t)_i\}_{i=1}^N$, where Q^q represents the query question, $D^q = \{d_i\}_{i=1}^{n_c}$ represents n_c candidate multi-modal documents, and $D^t = \{d_i\}_{i=1}^{n_t}$ represents n_t target multi-modal documents related to the question, and A^t represents the target answer to the question Q^q . Given a query question Q^q and multiple candidate multi-modal documents D^q , our goal is to first establish a coarse-grained correlation evaluation function $f_c(Q^q, d_i)$ between the query question Q^q and a candidate multi-modal document d_i , and retrieve the top-k multi-modal documents $D^k = \{d_i\}_{i=1}^k$ related to the query question based on the point-wise ranking method $S_{pw}(Q^q, D^q, f_c)$. Subsequently, we seek to jointly optimize the fine-grained list-wise multi-modal document ranking method $S_{lw}(Q^q, D^k)$ and the answer generation method $G(Q^q, D^k)$ to ensure that the generative model can give accurate answers based on optimal multi-modal reference documents. This can be formally expressed as:

$$\begin{aligned} S_{lw}(Q^q, S_{pw}(Q^q, D^q, f_c)) &\rightarrow D^t \\ G(Q^q, S_{lw}(Q^q, S_{pw}(Q^q, D^q, f_c))) &\rightarrow A^t \end{aligned}$$

Coarse-Grained Point-Wise Document Ranking

In the first stage, given the high efficiency of point-wise ranking in large-scale document ranking, coarse-grained point-wise document ranking is employed to obtain comparable relevance scores between the query and each multi-modal candidate document, thereby preliminarily filtering

out significantly irrelevant documents. To this end, we explored the point-wise ranking method for multi-modal documents based on multi-modal large language models, which encompasses two core components: rule-based reinforcement fine-tuning and point-wise probabilistic reasoning.

Rule-Based Reinforcement Fine-Tuning The rule-based reinforcement fine-tuning aims to endow the multi-modal large language model with stronger semantic understanding ability, enabling it to effectively distinguish the degree of relevance between documents in different modalities and the query question, and thus optimize the effect of coarse-grained point-wise multi-modal document ranking. First, the initial training data $\xi = \{(Q^q, D^q, D^t, A^t)_i\}_{i=1}^N$ is processed into the required format for reinforcement fine-tuning. Specifically, the target multi-modal documents $\{d_i\}_{i=1}^{n_t}$ are taken as positive samples and their labels are set to "Yes", while the complement of the target multi-modal documents in the candidate documents is taken as negative samples $\{d_i\}_{i=1}^{n_c} \setminus \{d_i\}_{i=1}^{n_t}$ and their labels are set to "No". The sample in the negative sample set with the largest semantic similarity to the query question is defined as a hard negative sample d_{neg} . Collectively, the positive and hard negative samples $\{(Q^q_i, d_1), \dots, (Q^q_i, d_{n_t}), (Q^q_i, d_{neg})\}_{i=1}^N$ are input into the multi-modal large language model as the overall query $\{x_k\}_{k=1}^{N \times (n_t+1)}$. During training, the overall query is combined with the corresponding "Yes/No" judgments $\{J_k\}_{k=1}^{N \times (n_t+1)}$ to jointly optimize the ability of the multi-modal large language model to discern semantic boundaries in relevance judgment.

Then, rule-based group relative strategy optimization algorithm is used to guide the multi-modal large language model to correct the model parameters according to the training data. Specifically, for each training sample input x_k , we first sample G different responses $\{o_j\}_{j=1}^G$ based on the old model policy $\pi_{\theta_{old}}$. Correspondingly, each response will obtain a corresponding reward value $\{r_j\}_{j=1}^G$ based on the format reward function $R_{format}(o_j)$ and the relevance judgment reward function $R_{judge}(o_j)$. The rule-based reward evaluation method is defined as follows:

$$R_{format}(o_j) = \begin{cases} 1, & \text{if } o_j = \text{Yes or } o_j = \text{No}, \\ 0, & \text{otherwise.} \end{cases}$$

$$R_{judge}(o_j) = \begin{cases} 1, & \text{if } o_j = J_k, \\ 0, & \text{otherwise.} \end{cases}$$

$$r_j = R_{format}(o_j) + R_{judge}(o_j)$$

Therefore, the reward value r_j corresponding to each response o_j can be obtained. By calculating the reward mean and variance of G different responses $\{o_j\}_{j=1}^G$, the relative advantage A_j of the current response can be defined as:

$$A_j = \frac{r_j - \text{mean}(\{r_1, r_2, \dots, r_G\})}{\text{std}(\{r_1, r_2, \dots, r_G\})}$$

Based on these relative advantages, we aim to optimize the policy of current model π_{θ} . It updates the model parameters θ by maximizing the expected cumulative reward. The

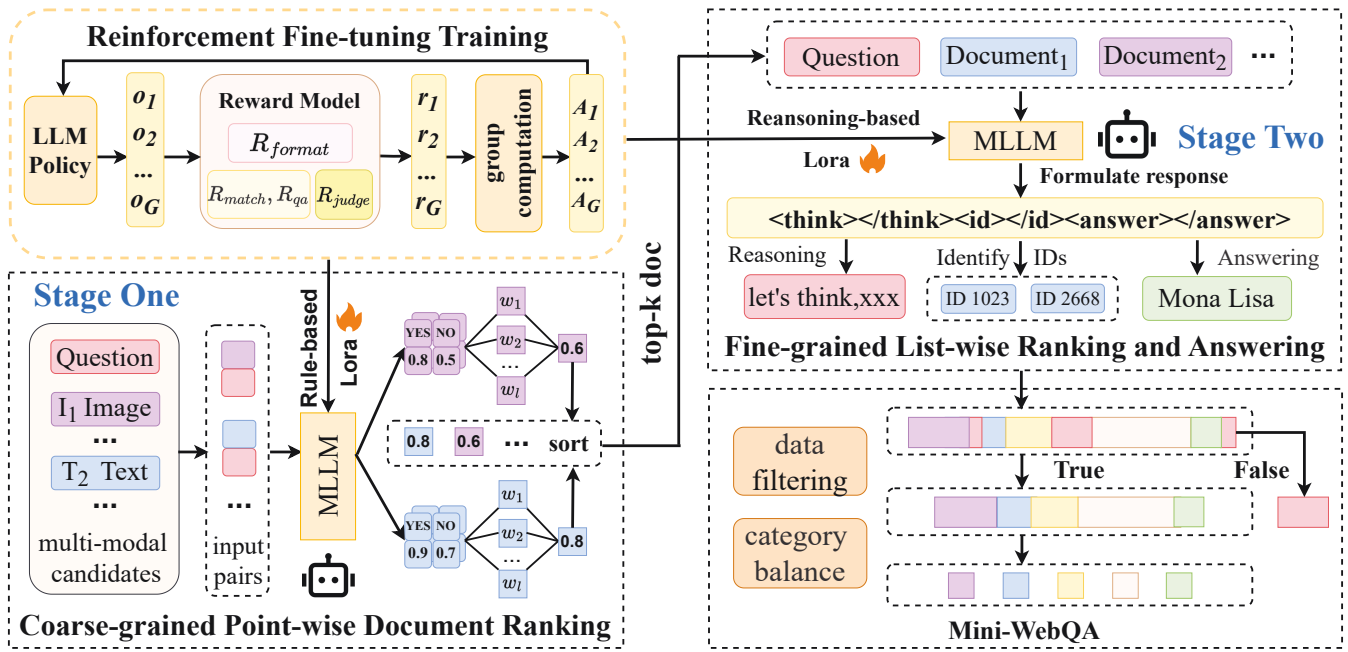


Figure 2: Illustration of our proposed two-stage reinforcement fine-tuning framework.

objective function $\mathcal{J}^{coarse}(\theta)$ for coarse-grained point-wise ranking optimization is formulated as follows:

$$\mathcal{J}^{coarse}(\theta) = \mathbb{E}[x \sim P(X), \{o_j\}_{j=1}^G \sim \pi_{\theta_{old}}(O|x)]$$

$$\frac{1}{G} \sum_{j=1}^G (\min(p_{j,\theta} A_j, \text{clip}(p_{j,\theta}, 1 + \epsilon, 1 - \epsilon) A_j))$$

$$p_{j,\theta} = \pi_{\theta}(o_j|x) / \pi_{\theta_{old}}(o_j|x)$$

where ϵ is a clipping-related hyper-parameter introduced in proximal policy optimization for stabilizing training.

Point-Wise Probabilistic Inference After rule-based reinforcement fine-tuning, point-wise probabilistic inference is used for testing, which aims to estimate the similarity between a candidate multi-modal document d_i^{test} and the query question Q^{test} based on the answer (Yes/No) of the multi-modal large language model. Specifically, the test query Q^{test} and each document d_i^{test} in the list of candidate multi-modal documents $D^{test} = \{d_i^{test}\}_{i=1}^{n_{test}}$ are reconstructed as a test samples $\{Q^{test}, d_i^{test}\}$ for input into the multi-modal large language model. The probability $\text{pro}_s(\text{Yes/No}|Q^{test}, d_i^{test})$ of "Yes" or "No" generated by the multi-modal large language model is used to evaluate the similarity $\text{Sim}_s(Q^{test}, d_i^{test})$ between the test query and a candidate document. At the same time, considering the randomness of the large model output, L responses are generated based on the current test input sampling, and the adaptive probability aggregation strategy is used to comprehensively evaluate the similarity $\text{Sim}_s(Q^{test}, d_i^{test})$ between the test query and a given document. Thus, we established the point-wise relevance evaluation function $f_c(Q^{test}, d_i^{test})$ between the query question and the multi-modal document,

which can be defined as follows:

$$\text{logit}(\text{Yes}) = \exp(\text{pro}_s(\text{Yes}|Q^{test}, d_i^{test}))$$

$$\text{logit}(\text{No}) = \exp(\text{pro}_s(\text{No}|Q^{test}, d_i^{test}))$$

$$\text{Sim}_s(Q^{test}, d_i^{test}) = \frac{\text{logit}(\text{Yes})}{\text{logit}(\text{Yes}) + \text{logit}(\text{No})}$$

$$w_s = \frac{\exp(\text{Sim}_s(Q^{test}, d_i^{test}))}{\sum_{s=1}^L \exp(\text{Sim}_s(Q^{test}, d_i^{test}))}$$

$$f_c(Q^{test}, d_i^{test}) = \sum_{s=1}^L w_s \text{Sim}_s(Q^{test}, d_i^{test})$$

Accordingly, the point-wise relevance evaluation function is used for coarse-grained multi-modal document ranking. Specifically, the correlation score $f_c(Q^{test}, d_i^{test})$ between the test query Q^{test} and each document d_i^{test} in the candidate multi-modal document list $D^{test} = \{d_i^{test}\}_{i=1}^{n_{test}}$ is used to construct the point-wise ranking method S_{pw} . After the scores are sorted in descending order, the top-k multi-modal documents D_{topk}^{test} related to the query are obtained, which can be defined as follows:

$$S_{pw}(Q^{test}, D^{test}, f_c) = \text{argsort}(f_c(Q^{test}, d_1^{test}),$$

$$f_c(Q^{test}, d_2^{test}), \dots, f_c(Q^{test}, d_{n_{test}}^{test}))$$

$$D_{topk}^{test} = D^{test}[S_{pw}(Q^{test}, D^{test})[:, \text{topk}]]$$

Fine-Grained List-Wise Document Ranking and Answer Generation

In the second stage, addressing the limited discriminative power of point-wise ranking and the poor explainability of retrieval-augmented generation, we apply reasoning-based reinforcement fine-tuning to jointly optimize fine-grained list-wise ranking and answer generation. Based on

the coarsely filtered document list, the multi-modal large language model after reinforcement fine-tuning retrieves relevant document IDs and corresponding answers to the query through explainable reasoning logic. The method consists of two core components: reasoning-based reinforcement fine-tuning and explainable chain-of-thought inference.

Reasoning-Based Reinforcement Fine-Tuning The reasoning-based reinforcement fine-tuning aims to endow the multi-modal large language model with long-form reasoning thought chain capabilities, enabling it to generate structured semantic analyses by integrating query questions and top-k relevant documents from point-wise ranking, and thus jointly optimize the accuracy of fine-grained list-wise multi-modal document ranking and answer generation. Accordingly, similar to point-wise ranking, the training samples $\{(Q^q, D_{topk}^q)_i\}_{i=1}^N$ are input into the multi-modal large language model as the overall query $\{x_i\}_{i=1}^N$. And D_{topk}^q represents top-k multi-modal relevant documents obtained from the coarse-grained point-wise ranking. Meanwhile, the reasoning-based group relative strategy optimization algorithm is employed for the joint optimization of fine-grained ranking and answer generation by sampling G different answers $\{o_j\}_{j=1}^G$. The rewards of the reasoning-based group relative policy optimization algorithm include format reward R_{format} , fine-grained list-wise ranking reward R_{match} and answer generation reward R_{qa} . First, the format reward aims to guide the multi-modal large language model to output responses with predefined structures after performing thought chain reasoning, which can be defined as follows:

$$\begin{aligned} pat_1 &= \langle think \rangle .*? \langle /think \rangle \\ pat_2 &= \langle id \rangle .*? \langle /id \rangle \\ pat_3 &= \langle answer \rangle .*? \langle /answer \rangle \\ R_{format}(o_j) &= \begin{cases} 1, & \text{if } pat_i \text{ in } o_j, i = 1, 2, 3 \\ 0, & \text{otherwise.} \end{cases} \end{aligned}$$

Therefrom, the thinking and reasoning process of multi-modal large language models is enclosed within the $\langle think \rangle \langle /think \rangle$ tags, the answer of the most relevant document ID is enclosed within the $\langle id \rangle \langle /id \rangle$ tags, and the question answer is enclosed within the $\langle answer \rangle \langle /answer \rangle$ tags. Next, the fine-grained list-wise ranking reward is designed to guide the multi-modal large language model to directly give the IDs of the most relevant documents to the query based on the query and the top-k relevant documents obtained from the point-wise ranking, which can be defined as follows:

$$R_{match}(o_j) = \frac{|ID^{list} \cap ID^{target}|}{2n_{list}} + \frac{|ID^{list} \cap ID^{target}|}{2n_t}$$

Among them, $ID^{list} = \{id_t\}_{t=1}^{n_{list}}$ represents the most relevant document IDs generated by the model between the $\langle id \rangle \langle /id \rangle$ tags of o_j , $ID^{target} = \{id_t\}_{t=1}^{n_t}$ represents IDs of target multi-modal documents related to the question, and $|ID^{list} \cap ID^{target}|$ represents the number of elements in the intersection of the sets ID^{list} and ID^{target} . Finally, answer generation reward is intended to steer the multi-modal

large language model to generate accurate answers by leveraging the most relevant documents derived from list-wise ranking, which can be defined as follows:

$$R_{qa}(o_j) = \frac{\exp(Bartscore(A_i^t, o_j))}{\exp(Bartscore(A_i^t, A_i^t))}$$

where $Bartscore$ represents an evaluation metric that comprehensively assesses the factuality and fluency of answers (Yuan, Neubig, and Liu 2021). Accordingly, each response o_j will obtain a corresponding reward value r_j based on the overall reward function, which can be defined as follows:

$$r_j = R_{format}(o_j) + R_{match}(o_j) + R_{qa}(o_j)$$

Similar to the coarse-grained point-wise ranking in the first stage, the objective function $\mathcal{J}^{fine}(\theta)$ for fine-grained list-wise ranking and answering optimization can be defined as:

$$\begin{aligned} \mathcal{J}^{fine}(\theta) &= \mathbb{E}[x \sim P(X), \{o_i\}_{i=1}^G \sim \pi_{\theta_{old}}(O|x)] \\ &= \frac{1}{G} \sum_{i=1}^G (\min(p_{i,\theta} A_i, \text{clip}(p_{i,\theta}, 1 + \epsilon, 1 - \epsilon) A_i)) \end{aligned}$$

Explainable Chain-of-Thought Inference After reasoning-based reinforcement fine-tuning, explainable thought chain reasoning is used for testing, aiming to guide the multi-modal large model to generate the most relevant document ID and question answer through explainable reasoning logic. Specifically, the test query Q^{test} and the top-k multi-modal documents D_{topk}^{test} obtained by point-wise probability inference are jointly input into the multi-modal large language model after reasoning-based reinforcement fine-tuning. Under the guidance of preset prompts, the multi-modal large language model first performs multi-step chain reasoning between the $\langle think \rangle \langle /think \rangle$ tags to evaluate the semantic relevance between the query and each candidate document, and then generates a ranked list of document IDs $S_{lw}(Q^{test}, S_{pw})$ between the $\langle id \rangle \langle /id \rangle$ tags to identify the documents most relevant to the query question, and finally generates accurate answers $G(Q^{test}, S_{lw})$ between the $\langle answer \rangle \langle /answer \rangle$ tags based on the query and its most relevant documents. Therefore, we established the fine-grained list multi-modal document ranking method S_{lw} and the answer generation method G , which are defined as follows:

$$\begin{aligned} S_{lw}(Q^{test}, S_{pw}(Q^{test}, D^{test}, f_c)) &= \\ MLLM_{\langle id \rangle} &(Q^{test}, D_{topk}^{test}) \rightarrow D^t \\ G(Q^{test}, S_{lw}(Q^{test}, D^{test}, f_c)) &= \\ MLLM_{\langle ans \rangle} &(Q^{test}, D_{topk}^{test}) \rightarrow A^t \end{aligned}$$

Mini-WebQA

During the experiment, we noticed that the existing multi-modal retrieval-augmented generation datasets have problems with annotation errors and uneven category distribution. Therefore, we set out to construct the Mini-WebQA training dataset, which contains 5,000 high-quality samples. To eliminate low-quality samples from the training data of

Model	Type	Retr	QA-FL	QA-Acc	QA
VLP (Zhou et al. 2020)	w/o LLM/MLLM	68.9	42.6	36.7	22.6
VLP + VinVL (Chang et al. 2022)	w/o LLM/MLLM	70.9	44.2	38.9	24.1
MuRAG (Chen et al. 2022)	w/o LLM/MLLM	74.6	55.7	54.6	36.1
SKURG (Yang et al. 2023a)	w/o LLM/MLLM	88.2	55.4	57.1	37.7
PERQA (Yang et al. 2023b)	w/o LLM/MLLM	89.6	61.7	63.9	44.4
VLMT (Lim et al. 2025)	MLLM-Based	87.8	64.0	66.7	47.6
RAMQA (Bai, Grant, and Wang 2025)	MLLM-Based	88.4	64.1	66.6	48.1
AETGA (Zhang et al. 2024)	LLM-Based	88.6	<u>68.3</u>	<u>72.8</u>	<u>54.1</u>
Our Method (Trained on Mini-WebQA)	MLLM-Based	78.6	67.3	71.1	53.8
Our Method	MLLM-Based	<u>89.1</u>	70.8	76.3	58.3

Table 1: WebQA official test-set results indicated on the leaderboard. We achieve the highest result on QA-FL, QA-ACC, QA score. Bold numbers denote the best scores, and underlined numbers indicate the second-best scores.

the WebQA dataset, we adopted two steps. First, we removed question-answer pairs with obvious grammatical errors or unrecognizable characters. Second, we used a model fully trained on the complete WebQA dataset to generate answers for each question, estimating sample clarity and representativeness via answer accuracy, and retaining questions with estimated correct answers to avoid negative impacts of incorrectly labeled samples. Following these steps, 15,000 samples were retained. Additionally, we noticed that the filtered samples still suffered from the category imbalance problem. Therefore, we analyzed the quantity distribution across categories such as color, shape, yes/no, number, and multiple-choice questions, and removed samples with severe category bias. After achieving category balance, 5,000 high-quality samples were retained.

Experiments

Datasets

WebQA The WebQA dataset (Chang et al. 2022) is a multi-modal, multi-hop question-answering dataset designed to evaluate the capabilities of models in complex cross-modal reasoning and generation tasks. It contains 34.2k training question-answer pairs, 5k validation question-answer pairs, and 7.5k test question-answer pairs. Each question requires retrieving a variable number of relevant documents from approximately 40 multi-modal candidate documents for multi-step reasoning to obtain answers in the form of free-form sentences. The dataset employs two key evaluation metrics: the Retr score assesses the accuracy of models in retrieving relevant evidence documents, while the QA score evaluates the quality of answer generation by models. The QA score is a composite of the QA-FL score and the QA-Acc score. The former measures the fluency between the generated answer and the reference answer, and the latter evaluates the overlap of key entities between the generated answer and the reference answer. Among the various metrics used for WebQA evaluation, the QA score holds the most significance.

MultimodalQA MultimodalQA (Talmor et al. 2021) is a multi-modal question-answering dataset with 16 distinct question types, designed to evaluate the multi-modal, multi-hop reasoning capabilities of models. Each query is accom-

panied by approximately 20 multi-modal distractors, and models must distinguish these distractors from correct document. Following the experimental setup in RAMQA (Bai, Grant, and Wang 2025), we focus on the query subset defined as "ImageQ" related to image information. The performance metrics used for evaluation include exact match score (EM) and average F1 score, which measure the accuracy of the generated answers relative to the reference answers in terms of exact string matching and token overlap.

Implementation Details

We conducted experiments on two datasets: WebQA and MultimodalQA. Unlike previous methods that separate retrieval and generation into two modules, our proposed approach uses Qwen-2.5-VL-7B (Bai et al. 2025) as the backbone architecture throughout the entire process of multi-modal retrieval-augmented generation. We perform parameter-efficient reinforcement fine-tuning via group-relative policy optimization combined with low-rank adaptation (LoRA) (Hu et al. 2022). And CLIP-ViT-L-14 is used offline for hard negative sample selection. During the reinforcement fine-tuning process, group-relative policy optimization samples 4 responses for each question every time, sets ϵ to 0.28 to ensure stable training, and does not use KL divergence constraints to encourage the model to explore diverse responses independently. The training configuration of LoRA includes a rank of 64, a scaling factor of 128, a dropout rate of 0.05. During the training process, to improve efficiency and stability, we adopt bfloat16 mixed-precision training, FlashAttention-2 and gradient checkpointing (Chen et al. 2016) to accelerate model training. The generated sequence length is limited to 1024 tokens to avoid training collapse, the learning rate is set to $1e-5$, the generation number L is set to 4, and the data seed is set to 42 to ensure the reproducibility of the experiment.

Main Result

Tables 1 and 2 present a comparative analysis of the multi-modal retrieval-augmented generation results achieved by our proposed method against other leading methods on the WebQA dataset and MultimodalQA dataset. On the WebQA dataset, our method achieves a score of 0.708 in QA-

Model	EM	F1
Q-only (Lewis et al. 2019)	11.0	15.6
AutoRouting (Talmor et al. 2021)	37.8	37.8
MuRAG (Chen et al. 2022)	58.2	58.2
PERQA (Yang et al. 2023b)	63.9	64.1
RAMQA (Bai, Grant, and Wang 2025)	67.0	67.0
Our Method	71.7	79.7

Table 2: Experiment results on MultimodalQA dataset.

FL, 0.763 in QA-Acc, and 0.583 in overall QA performance, surpassing all baselines on these metrics. Even when trained solely on the Mini-WebQA dataset, our method still yields highly competitive results, which we attribute to the strong adaptability of our proposed framework to high-quality small-scale data. On the MultimodalQA dataset, compared with the previous state-of-the-art RAMQA, our method achieves absolute improvements of 4.7% and 12.7% in EM score and F1 score. The outstanding performance highlights the superiority of our proposed two-stage reinforcement fine-tuning framework in multi-modal semantic understanding and high-quality answer generation. It is worth noting that LLM-Based methods generally outperform traditional methods based on structural knowledge representation in terms of the quality of generated answers, but their retrieval performance is inferior to the state-of-the-art vector retrieval-based methods. We attribute this to the inherent disadvantage of generative models in retrieval capabilities relative to discriminative models.

Ablation Experiments

In this section, we conduct comprehensive ablation experiments to clarify the impact of each module in our method. All ablation experiments are based on Qwen-2.5-VL-3B as the backbone architecture, which is trained on the Mini-WebQA dataset and tested on the WebQA official test set.

Effect of Model Scale and Training Stage As shown in Tab. 3, we explored the impact of model scales and the training stages within the two-stage reinforcement fine-tuning framework on model performance. Among them, the large-scale Qwen2.5vl-7B model achieves better results than the basic Qwen2.5vl-3B model in most metrics. This is because a larger model scale enhances the ability of the model to understand multi-modal information and perform complex reasoning. Regarding the training stages, Stage1 achieves an absolute improvement of 15.0% in Retr, while Stage2 achieves an absolute improvement of 64.1% in Retr and 26.5% in QA respectively. This may be due to the fact that each stage of reinforcement fine-tuning enhances the multi-modal large language model in different aspects, and the reinforcement fine-tuning across different stages can promote each other.

Effect of Fine-Tuning Method To further investigate the impact of different fine-tuning methods on the multi-modal retrieval-augmented generative model, we carefully designed four models with distinct stage 1 and stage 2 training configurations (RFT/SFT). As shown in Tab. 4, the best

Model	Retr	QA-FL	QA-Acc	QA
Total w/ 7B model	78.6	67.3	71.1	53.8
Total	76.9	67.3	66.5	50.9
Total w/o Stage1	61.9	65.6	64.6	49.5
Total w/o Stage2	12.8	30.2	52.4	24.4

Table 3: Ablation experimental results for model scales and training stages. Stage 1:coarse-grained point-wise ranking, Stage 2: fine-grained list-wise ranking and answering.

Stage1	Stage2	Retr	QA-FL	QA-Acc	QA
RFT	RFT	76.9	67.3	66.5	50.9
SFT	RFT	51.5	63.8	61.5	45.4
RFT	SFT	63.4	49.7	53.0	34.9
SFT	SFT	43.6	46.8	49.9	31.7

Table 4: Ablation experimental results for fine-tuning methods. RFT: reinforcement fine-tuning method, SFT: supervised fine-tuning method.

performance was achieved when both stages were trained using reinforcement fine-tuning method (RFT). When Stage 1 was switched to supervised fine-tuning (SFT) while Stage 2 remained RFT, the Retr value dropped by 25.4%, indicating that the first stage of reinforcement fine-tuning training is crucial to improving the retrieval ability of the model. Meanwhile, configuring Stage 1 as RFT and Stage 2 as SFT led to substantial declines in multiple indicators for evaluating answer quality (QA-FL, QA-ACC, QA), highlighting the critical role of reinforcement fine-tuning in ensuring the accuracy and fluency of generated answers. In addition, the model that uses supervised fine-tuning in both stages performs the worst on all metrics. In summary, the reinforcement fine-tuning method significantly outperforms the traditional supervised fine-tuning method across various experimental settings. We attribute this superiority to the fact that the supervised fine-tuning forces the model to overfit specific task patterns, whereas the reinforcement fine-tuning enables the multi-modal large language model to evolve autonomously without compromising the integrity of the pre-trained knowledge architecture.

Conclusion

In this paper, we propose integrating reinforcement learning into multi-modal retrieval-augmented generation to enhance reasoning capabilities of multi-modal large language models, thereby enabling explainable multi-modal retrieval-augmented generation. We design a two-stage reinforcement fine-tuning framework: the first stage involves coarse-grained point-wise ranking of multi-modal documents for efficient filtering, while the second conducts joint optimization of fine-grained list-wise ranking and answer generation, aiming for explainable reasoning. Extensive experiments on datasets demonstrate the effectiveness of our framework.

Acknowledgments

This work was supported by the National Natural Science Foundation of China under Grant Nos. U24A20252 and 62088102, the Key Research and Development Program of Shaanxi Province of China under Grant Nos. 2024PT-ZCK-66 and 2024CY2-GJHX-48.

References

- Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altschmidt, J.; Altman, S.; Anadkat, S.; et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Bai, Y.; Grant, C. E.; and Wang, D. Z. 2025. RAMQA: A unified framework for retrieval-augmented multi-modal question answering. *arXiv preprint arXiv:2501.13297*.
- Chang, Y.; Narang, M.; Suzuki, H.; Cao, G.; Gao, J.; and Bisk, Y. 2022. Webqa: Multihop and multimodal qa. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16495–16504.
- Chen, T.; Xu, B.; Zhang, C.; and Guestrin, C. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*.
- Chen, W.; Hu, H.; Chen, X.; Verga, P.; and Cohen, W. W. 2022. Murag: Multimodal retrieval-augmented generator for open question answering over images and text. *arXiv preprint arXiv:2210.02928*.
- Feng, Y.; Hu, H.; Hou, X.; Liu, S.; Ying, S.; Du, S.; Hu, H.; and Gao, Y. 2025. Hyper-RAG: Combating LLM Hallucinations using Hypergraph-Driven Retrieval-Augmented Generation. *arXiv preprint arXiv:2504.08758*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Guu, K.; Lee, K.; Tung, Z.; Pasupat, P.; and Chang, M. 2020. Retrieval augmented language model pre-training. In *International conference on machine learning*, 3929–3938. PMLR.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Jaech, A.; Kalai, A.; Lerer, A.; Richardson, A.; El-Kishky, A.; Low, A.; Helyar, A.; Madry, A.; Beutel, A.; Carney, A.; et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2019. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Lewis, P.; Perez, E.; Piktus, A.; Petroni, F.; Karpukhin, V.; Goyal, N.; Küttler, H.; Lewis, M.; Yih, W.-t.; Rocktäschel, T.; et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33: 9459–9474.
- Lim, Q. Z.; Lee, C. P.; Lim, K. M.; and Anbananthen, K. S. M. 2025. VLMT: Vision-Language Multimodal Transformer for Multimodal Multi-hop Question Answering. *arXiv preprint arXiv:2504.08269*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Liu, Z.; Sun, Z.; Zang, Y.; Dong, X.; Cao, Y.; Duan, H.; Lin, D.; and Wang, J. 2025. Visual-rft: Visual reinforcement fine-tuning. *arXiv preprint arXiv:2503.01785*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Shen, H.; Liu, P.; Li, J.; Fang, C.; Ma, Y.; Liao, J.; Shen, Q.; Zhang, Z.; Zhao, K.; Zhang, Q.; et al. 2025. Vlm-r1: A stable and generalizable r1-style large vision-language model. *arXiv preprint arXiv:2504.07615*.
- Talmor, A.; Yoran, O.; Catav, A.; Lahav, D.; Wang, Y.; Asai, A.; Ilharco, G.; Hajishirzi, H.; and Berant, J. 2021. Multimodalqa: Complex question answering over text, tables and images. *arXiv preprint arXiv:2104.06039*.
- Yang, Q.; Chen, Q.; Wang, W.; Hu, B.; and Zhang, M. 2023a. Enhancing multi-modal multi-hop question answering via structured knowledge and unified retrieval-generation. In *Proceedings of the 31st ACM International Conference on Multimedia*, 5223–5234.
- Yang, S.; Wu, A.; Wu, X.; Xiao, L.; Ma, T.; Jin, C.; and He, L. 2023b. Progressive evidence refinement for open-domain multimodal retrieval question answering. *arXiv preprint arXiv:2310.09696*.
- Ye, X.; and Durrett, G. 2022. The unreliability of explanations in few-shot prompting for textual reasoning. *Advances in neural information processing systems*, 35: 30378–30392.
- Yuan, W.; Neubig, G.; and Liu, P. 2021. Bartscore: Evaluating generated text as text generation. *Advances in neural information processing systems*, 34: 27263–27277.
- Zhang, Q.; Lv, H.; Liu, J.; Chen, Z.; Duan, J.; Wang, H.; He, L.; and Xu, M. 2024. An entailment tree generation approach for multimodal multi-hop question answering with mixture-of-experts and iterative feedback mechanism. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 4814–4822.
- Zhao, S.; Xu, L.; Liu, Y.; and Du, S. 2023. Multi-grained representation learning for cross-modal retrieval. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2194–2198.
- Zhou, L.; Palangi, H.; Zhang, L.; Hu, H.; Corso, J.; and Gao, J. 2020. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 13041–13049.