

This is the peer reviewed version of the following article:

ReAG: Reasoning-Augmented Generation for Knowledge-based Visual Question Answering / Compagnoni, Alberto; Morini, Marco; Sarto, Sara; Cocchi, Federico; Caffagni, Davide; Cornia, Marcella; Baraldi, Lorenzo; Cucchiara, Rita. - (2026). (IEEE/CVF Conference on Computer Vision and Pattern Recognition Denver (CO), United States June 3-7, 2026).

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

23/04/2026 14:00

(Article begins on next page)

ReAG: Reasoning-Augmented Generation for Knowledge-based Visual Question Answering

Alberto Compagnoni^{*1,2} Marco Morini^{*1} Sara Sarto¹ Federico Cocchi^{1,2} Davide Caffagni¹

Marcella Cornia¹ Lorenzo Baraldi¹ Rita Cucchiara¹

¹University of Modena and Reggio Emilia, Italy ²University of Pisa, Italy

¹{name.surname}@unimore.it ²{name.surname}@phd.unipi.it

aimagelab.github.io/ReAG

Abstract

Multimodal Large Language Models (MLLMs) have shown impressive capabilities in jointly understanding text, images, and videos, often evaluated via Visual Question Answering (VQA). However, even state-of-the-art MLLMs struggle with domain-specific or knowledge-intensive queries, where relevant information is underrepresented in pre-training data. Knowledge-based VQA (KB-VQA) addresses this by retrieving external documents to condition answer generation, but current retrieval-augmented approaches suffer from low precision, noisy passages, and limited reasoning. To address this, we propose **ReAG**, a novel Reasoning-Augmented Multimodal RAG approach that combines coarse- and fine-grained retrieval with a critic model that filters irrelevant passages, ensuring high-quality additional context. The model follows a multi-stage training strategy leveraging reinforcement learning to enhance reasoning over retrieved content, while supervised fine-tuning serves only as a cold start. Extensive experiments on Encyclopedic-VQA and InfoSeek demonstrate that ReAG significantly outperforms prior methods, improving answer accuracy and providing interpretable reasoning grounded in retrieved evidence.

1. Introduction

Multimodal Large Language Models (MLLMs) [1, 4, 30, 31] unify tasks involving multiple modalities, such as text, images and videos [6, 44]. Many of these tasks can be framed as Visual Question Answering (VQA) [2, 14, 53], where a query may require understanding visual content, and the model must generate a faithful, correctly formatted response. Despite their broad pre-training, state-of-the-art MLLMs struggle with underrepresented, domain-specific queries [10, 36]. This problem, known as Knowledge-based VQA (KB-VQA) [35], is commonly addressed by enriching MLLMs

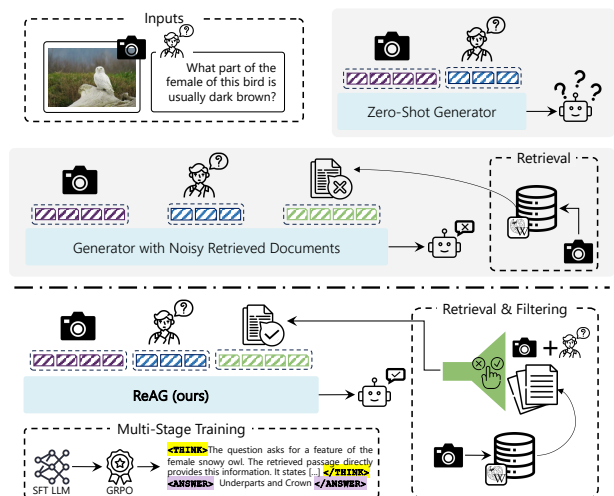


Figure 1. Comparison between Zero-Shot (ZS) MLLMs, retrieval-augmented models, and ReAG. ZS MLLMs lack specialized knowledge and fail on domain-specific queries (top). Retrieval-augmented models introduce external context but often add noisy or irrelevant passages (middle). ReAG overcomes this with a filtering stage over retrieved content and a multi-stage training strategy to enhance reasoning over passages.

with domain-specific information from external sources, *i.e.*, via Retrieval-Augmented Generation (RAG) [26].

Despite impressive results [13, 16, 54, 55, 58, 59], this setting still presents significant open challenges. One lies in information retrieval itself, as users’ queries can be extremely heterogeneous while the external knowledge-base can reach millions of documents in cardinality [10, 36] – thus lowering the recall of the retrieved results and adding noise to the MLLM input. This is further exacerbated by feature extraction and integration issues when queries and documents are multimodal, which is rapidly becoming usual [8, 29, 50]. Next, even assuming that the retrieved documents were relevant to the query, understanding them and extracting the right piece of information to generate the answer is not trivial [17].

^{*}Equal contribution.

To address these challenges, we propose **ReAG**, short for **Reasoning-Augmented Generation**, a novel multimodal retrieval-augmented generation approach that (i) mitigates low-recall and noisy retrieval by employing a multi-level retrieval pipeline followed by a critic model that effectively filters out irrelevant samples, and (ii) equips the MLLM with the capability of reasoning over retrieved results through a dedicated reinforcement learning training protocol.

During retrieval, following common practice, ReAG first employs an off-the-shelf multimodal encoder [41, 48] for coarse-grained embedding-based retrieval, which achieves high recall when retaining many documents but suffers from low precision due to noise. To improve precision, a fine-grained retrieval stage focuses on the visual regions most relevant to the question. A critic model then classifies each passage as *relevant* or *irrelevant*, ensuring that only high-quality documents are passed to the generator (Fig. 1).

Building on advances in reasoning models [15], ReAG further improves generation quality by allowing the model to produce explicit natural-language reasoning traces before the final answer. Unlike prior multimodal RAG methods that rely on implicit reflection signals [3, 13], we train the model to create a full reasoning trace in natural language, without obeying any predefined pattern. To achieve this, ReAG replaces the traditional supervised training in favor of a reinforcement learning framework inspired by GRPO [47, 57], equipped with a reward scheme tailored for KB-VQA, to refine the model ability to reason over the user query and retrieved evidence, while using supervised fine-tuning only as a cold-start to establish initial reasoning behavior.

Experimentally, we evaluate the proposed approach on Encyclopedic-VQA [36] and InfoSeek [10], two benchmarks containing question-answer pairs linked to Wikipedia-derived knowledge bases. Extensive experiments show that ReAG substantially outperforms prior methods, not only improving answer accuracy but also generating explicit reasoning traces. These traces provide insight into the usefulness of retrieved passages and the steps leading to the final answer, offering full explainability of the model predictions.

In summary, the contributions of this work are as follows:

- We propose ReAG, a novel reasoning-augmented multimodal RAG model that combines coarse- and fine-grained retrieval with a critic model to improve precision and reduce noise injected to the generator. Notably, the critic is agnostic to the retrieval backbone, making it seamlessly applicable on top of any state-of-the-art retrieval engine.
- ReAG trains the generator employing a multi-stage training strategy, leveraging SFT only as a cold start, followed by a reinforcement learning framework inspired by GRPO, with a reward scheme specifically designed for KB-VQA.
- We empirically validate ReAG on two popular and challenging KB-VQA benchmarks, Encyclopedic-VQA and InfoSeek, where ReAG reaches a new state-of-the-art.

2. Related Work

Knowledge-based VQA. The task requires models to answer questions that depend on external or specialized knowledge beyond the visual content of an image. Early datasets [35, 45, 46] targeted specialized reasoning. However, with the advent of more powerful MLLMs, these datasets have become insufficient for evaluating performance in realistic and knowledge-intensive settings. To address this, benchmarks such as Encyclopedic-VQA [36] and InfoSeek [10] introduce more challenging, Wikipedia-scale scenarios requiring fine-grained and entity-specific reasoning over external knowledge, making retrieval essential.

The RAG framework has become the standard approach for this task, retrieving relevant content from sources such as Wikipedia. One line of work focuses on enhancing the retrieval itself to obtain more accurate and less noisy results, as in WikiLLaVA [7], which integrates external multimodal knowledge through a hierarchical retrieval framework. Other focuses on handling noisy retrieval, refining visual tokens [40] or re-ranking retrieved textual passages before processing by the LLM [55]. A third direction strengthens model-level control: VLM-PRF [16] employs external tools for knowledge filtering, while ReflectiVA [13] uses control tokens to guide retrieval and knowledge assessment.

In this work, we first aim to reduce retrieval noise through effective filtering mechanisms and then empower the model with reasoning capabilities to critically evaluate retrieved knowledge before generating the final answer.

RL-based Strategies for LLMs and MLLMs. As frontier LLMs advance, RL has emerged as a key paradigm for aligning model outputs with human values and desired behaviors [24, 39, 42]. Concurrently, the quality and diversity of training data remain crucial for robust alignment and generalization [60]. Recently, GRPO [47] has emerged as a promising approach to improve sample efficiency and training stability at scale. Several variants [33, 57] revisit the underlying loss formulation to mitigate bias and enhance token-level optimization efficiency.

Inspired by the performance improvements achieved through GRPO-based methods, similar strategies have been extended to MLLMs [28, 51]. In particular, GRPO-CARE [11] enhances the coherence between intermediate reasoning traces and final outputs, leading to more reliable reasoning-grounded responses. Advanced LLM reasoning has also been explored in knowledge-intensive tasks. For example, Search-R1 [20] integrates retrieval and reasoning for complex queries, while subsequent approaches extend this paradigm to multimodal search [19, 38, 52].

In this work, we build upon these foundations by introducing a multi-stage reinforcement learning framework that operates on top of a supervised fine-tuned MLLM, enhancing its ability to reason effectively over retrieved evidence.

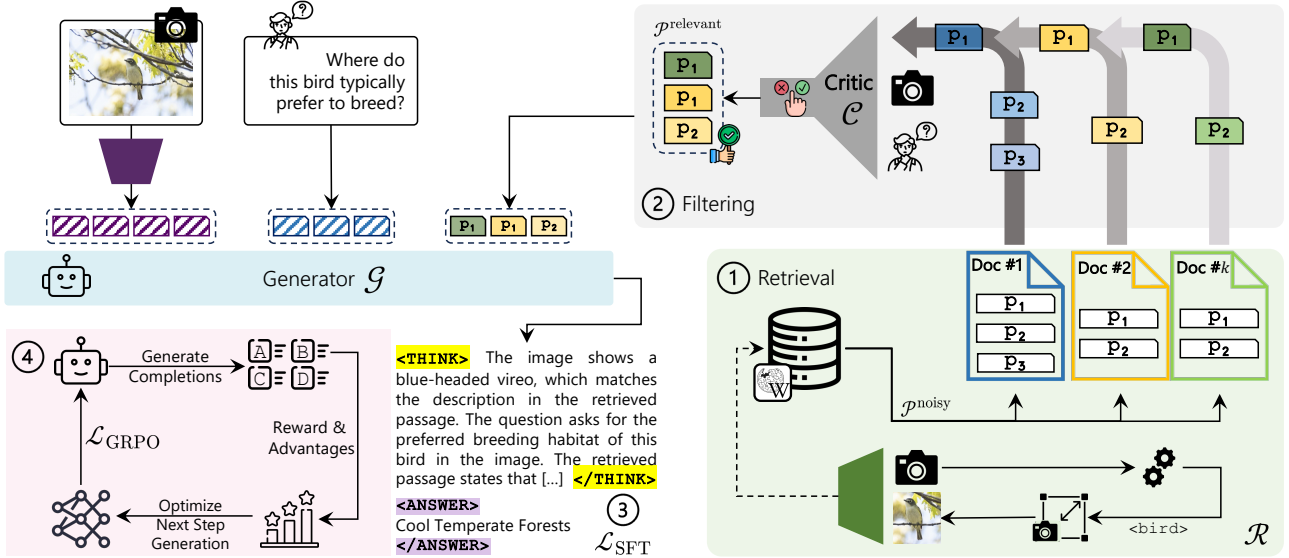


Figure 2. Overview of the proposed ReAG model. A multi-level retriever module extracts noisy passages, which are refined by a critic model. The resulting relevant passages are fed to a generator trained via SFT and a reinforcement learning stage designed for the KB-VQA task.

3. Proposed Method

Task Definition. In the standard VQA task, a multimodal LLM, referred to as the generator model \mathcal{G} , must answer a question q about an image I_q . The task requires the model to understand the visual content and provide a correct answer. While MLLMs large-scale pretraining captures general knowledge, it may be insufficient for answering highly specific or domain-specific questions.

Knowledge-based VQA (KB-VQA) extends VQA by incorporating external knowledge. In our setting, the external knowledge base \mathcal{KB} is a collection of N multimodal documents (*e.g.*, Wikipedia pages) each containing textual passages and images. Formally, the knowledge base can be represented as

$$\mathcal{KB} = \{d_1, \dots, d_N\}, \quad d_i = (\mathcal{T}_i, I_i, \mathcal{P}_i), \quad (1)$$

where \mathcal{T}_i is the metadata of the i -th document (*e.g.*, title and summary of a Wikipedia page), I_i is the associated image, if present, and \mathcal{P}_i are the textual passages of the document.

A retrieval model \mathcal{R} is employed to select the top- k relevant documents from \mathcal{KB} and their associated passages $\tilde{\mathcal{P}} = \{p_0, \dots, p_j\}$, which are then provided within the generator context window. Finally, the generator produces an answer A conditioned on both the image, the question, and the retrieved passages, as follows:

$$A \sim \mathcal{G}(A | q, I_q, \{p_0, \dots, p_j\}). \quad (2)$$

During training, the generator \mathcal{G} is optimized to maximize the likelihood of producing the correct answer given the visual and textual context. In particular, the retrieved passages $\tilde{\mathcal{P}}$ act as external conditioning signals that augment

the understanding of the model of the visual scene and the question, enabling knowledge-grounded reasoning. The objective can thus be expressed as the negative log-likelihood of the ground-truth answer tokens, averaged over the training distribution, *i.e.*

$$\mathcal{L}(\theta) = \mathbb{E}_{(I_q, q, \tilde{\mathcal{P}}) \sim \mathcal{D}} \left[\frac{1}{|y|} \sum_{t=1}^{|y|} \log \mathcal{G}_\theta(y_t | q, I_q, \tilde{\mathcal{P}}, y_{<t}) \right], \quad (3)$$

where \mathcal{D} denotes the training distribution.

Methodology Summary. To address the existing challenges of retrieval-augmented models, **ReAG** enhances KB-VQA performance by employing retrieval, filtering, and reasoning over the retrieved passages. The approach consists of two key components: a **critic model** that filters retrieval results, and a **generator** trained to reason over the filtered documents before producing the final answer. The overall pipeline is organized into four main stages: (1) a multi-level retrieval stage to gather candidate passages, (2) a filtering stage where the critic selects relevant content, (3) a cold-start supervised fine-tuning (SFT) stage to instill initial reasoning capabilities in the generator, and (4) a reinforcement learning stage to further refine reasoning and answer generation. Together, these components reduce noise and improve the ability of the generator to produce accurate, knowledge-intensive answers. An overview of our methodology is shown in Fig. 2.

3.1. Retrieval Stage

The retrieval stage identifies potentially informative passages related to the query image, which are subsequently filtered to provide the generator with relevant external knowledge for reasoning and answer generation. This process comprises

two complementary steps: a *coarse-grained* retrieval, which retrieves candidate documents based on the entire query image, and a *fine-grained* retrieval, which performs retrieval using localized cues. Notably, ReAG is agnostic to the choice of retriever, so \mathcal{R} can be any cross-modal encoder that maps the query image and either the metadata \mathcal{T}_i or the image I_i^1 associated with each document d_i into a shared embedding space. Relevance between queries and documents is then computed via cosine similarity.

Coarse-Grained Retrieval. An initial set of relevant textual passages, denoted as \mathcal{P}^{cg} , is constructed by aggregating all the passages contained in the top- k retrieved documents when using the original image I_q as query to the retriever \mathcal{R} . Since each document contains a variable number of passages, the resulting collection is represented as $\mathcal{P}^{\text{cg}} = \{p_1^{\text{cg}}, \dots, p_m^{\text{cg}}\}$, where m denotes the total number of passages gathered from the top- k documents.

Fine-Grained Retrieval. To improve retrieval recall, we introduce a fine-grained retrieval stage that focuses on the specific visual region relevant to the question. Given the input image I_q and the question q , we identify a bounding box corresponding to the subject of the question, employing an off-the-shelf detection model. If such a region is detected, we crop the image accordingly, obtaining a focused image patch I_q^{crop} . This cropped image is then used as input to the retriever model \mathcal{R} , which computes relevance scores with respect to each document d_i , as in the coarse-grained stage.

The top- k documents retrieved in this stage form the fine-grained candidate passages, denoted as $\mathcal{P}^{\text{fg}} = \{p_1^{\text{fg}}, \dots, p_l^{\text{fg}}\}$. By restricting the visual input to the region of interest, this stage allows the retriever to focus on more fine-grained visual details, yielding passages that are more likely to be relevant to the specific question.

Final Set of Retrieved Passages. The documents comprising \mathcal{P}^{cg} and \mathcal{P}^{fg} are merged and ranked by their relevance scores, and all passages contained in the top- k ranked documents are retained to form the final set $\mathcal{P}^{\text{noisy}}$ from \mathcal{KB} .

3.2. Filtering

After the retrieval steps, we obtain a set $\mathcal{P}^{\text{noisy}}$ of passages from the k retrieved documents. While increasing k generally improves recall by including more potentially relevant passages, this typically comes at the cost of a lower precision, as the probability of introducing noisy information rises as well. To mitigate this, we design a critic model \mathcal{C} to filter out irrelevant passages, resulting in a refined set of relevant passages $\mathcal{P}^{\text{relevant}}$.

Critic Model. Given a question q and its corresponding image I_q , the critic model \mathcal{C} predicts if each retrieved textual passage in $\mathcal{P}^{\text{noisy}}$ is useful for answering the question². In

ReAG, the critic model is implemented as an autoregressive MLLM fine-tuned with a next-token prediction objective optimized on an annotated dataset. Specifically, starting from a subset of samples drawn from the dataset employed in [13], we extract tuples (I_q, q, p, y) , where p is a textual passage to be evaluated and $y \in \{Yes, No\}$ indicates whether the passage is relevant. The critic model is trained to predict y conditioned on (I_q, q, p) , enabling it to robustly discriminate between relevant and irrelevant passages.

At inference time, only passages yielding a positive prediction with probability above a threshold are kept, yielding the final subset of relevant passages $\mathcal{P}^{\text{relevant}}$, defined as:

$$\mathcal{P}^{\text{relevant}} = \{p \in \mathcal{P}^{\text{noisy}} \mid \Pr(Yes|\mathcal{C}, q, I_q, p) > \text{thresh}\}. \quad (4)$$

The resulting set $\mathcal{P}^{\text{relevant}}$ is fed to the generator \mathcal{G} , which leverages these passages to produce the final answer (Eq. 2).

3.3. Generator Cold Start

Following the approach popularized by DeepSeek-R1 [15], we train our generator \mathcal{G} using a multi-stage strategy. The initial stage is designed to enhance the reasoning and zero-shot capabilities of the model, while mitigating potential instabilities during the subsequent reinforcement learning stage. Unlike standard SFT, which focuses solely on answer prediction, our cold-start phase exposes \mathcal{G} to explicit reasoning trajectories that link the visual content, retrieved passages, and the question.

Collecting Reasoning Traces. To achieve this, we fine-tune \mathcal{G} using high-quality reasoning data. Starting from the same subset used for training the critic model, we extend each tuple (I_q, q, p, y) with a reasoning trace tr . Specifically, each tuple is provided as input to an MLLM, which is prompted to generate an explicit reasoning trace that logically explains how the passage p contributes to answering the question q given the image I_q . To guide this reasoning, the prompt includes both the final answer and the relevance label y , indicating whether the reasoning should be grounded in the passage or not. By explicitly conditioning on these signals, the MLLM produces structured reasoning traces that reflect a coherent inference process from the evidence to the answer. These traces are used as supervision for the cold-start fine-tuning of the generator \mathcal{G} .

Training Protocol. Having collected the reasoning-augmented dataset, the generator \mathcal{G} is trained to optimize both its reasoning ability and answer accuracy. To guide the model towards structured reasoning behavior, we encourage a templated output format, where the reasoning trace and the final answer are delimited by special tokens which are explicitly added to the vocabulary, *i.e.*

```
<think> reasoning trace </think>
<answer> answer </answer>.
```

¹Depending on the test case, see Sec. 4.2 for details.

²The exact prompt can be found in the supplementary material.

This structure encourages the model to separate intermediate reasoning from the final prediction, improving interpretability and stability during generation. Training is performed using a next-token prediction objective over both the reasoning trace and the final answer. The overall SFT loss balances the two components as follows:

$$\mathcal{L}_{\text{SFT}} = \alpha \mathcal{L}_{\text{A}} + (1 - \alpha) \mathcal{L}_{\text{T}}, \quad (5)$$

where \mathcal{L}_{A} and \mathcal{L}_{T} denote the negative log-likelihood losses computed over the answer and reasoning trace, respectively.

3.4. Generator RL Training

While supervised fine-tuning on cold-start data equips the generator with basic reasoning skills and coherent chain-of-thought generation, we further enhance its quality and robustness through a subsequent reinforcement learning stage.

Task-specific RL with Retrieved Passages. Our generator model is optimized with a custom objective inspired by GRPO [47], incorporating several modifications from DAPO [57]. Formally, the objective is defined as follows:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{(I_q, q, p) \sim \mathcal{D}, \{o_i\}_{i=1}^N \sim \mathcal{G}_{\theta_{\text{old}}(\cdot | I_q, q, p)}} \left[\frac{1}{\sum_{i=1}^N |o_i|} \sum_{i=1}^N \sum_{t=1}^{|o_i|} \min \left(r_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(r_{i,t}(\theta), 1 - \varepsilon, 1 + \varepsilon) \hat{A}_{i,t} \right) \right], \quad (6)$$

where $\mathcal{G}_{\theta_{\text{old}}}$ is the generator initialized from the SFT cold-start phase and \mathcal{G}_{θ} the generator after optional off-policy updates. Moreover, $\{o_i\}_{i=1}^N$ are the generated completions with associated rewards R_i , and \hat{A}_i the corresponding GRPO advantages [47]. In the formula, $r_{i,t}$ is computed as follows:

$$r_{i,t}(\theta) = \frac{\mathcal{G}_{\theta}(o_{i,t} | I_q, q, p, o_{i,<t})}{\mathcal{G}_{\theta_{\text{old}}}(o_{i,t} | I_q, q, p, o_{i,<t})}. \quad (7)$$

As in our setting the updates are never off-policy, \mathcal{G}_{θ} coincides with $\mathcal{G}_{\theta_{\text{old}}}$, thus the ratio $r_{i,t}(\theta)$ is always 1. Unlike the GRPO formulation and following DAPO, we omit the KL divergence penalty, which overly constrains exploration of alternative reasoning trajectories. This also improves memory efficiency and training speed by removing the need for a reference model and an extra forward pass. Furthermore, the loss is computed at the token level, as averaging over variable-length sequences would reduce the contribution of tokens in longer sequences and weaken their updates.

At each training iteration, the generator \mathcal{G} is prompted with (I_q, q, p) to autoregressively generate a group of N completions $\{o_i\}_{i=1}^N$. Each generated completion is then evaluated by one or more rule-based reward functions, producing a reward R_i to compute the advantage as:

$$\hat{A}_{i,t} = \frac{R_i - \text{mean}(\{R_i\}_{i=1}^N)}{\text{std}(\{R_i\}_{i=1}^N)}. \quad (8)$$

Advantages guide the policy updates: completions with above-average rewards have their likelihood increased, while those below the mean are down-weighted. This exposure allows the generator to explore diverse strategies for interacting with passages, gradually steering it toward producing reasoning trajectories that yield correct answers.

Rule-Based Reward Design. In our setting, we employ two complementary rule-based binary reward functions: a task-specific *accuracy reward* and a *format reward*. The task-specific accuracy reward $R_{\text{task}}(o_i)$ verifies whether a generated completion is correct by parsing the prediction according to the question type (*i.e.*, numerical or textual, single- or multi-answer)³. The format reward $R_{\text{fmt}}(o_i)$ enforces adherence to the expected output template. Both functions return 1 in case of success and 0 otherwise. The final reward associated to a completion o_i is defined as a weighted sum of the two. Formally,

$$R_i = \gamma R_{\text{task}}(o_i) + \delta R_{\text{fmt}}(o_i), \quad (9)$$

where γ and δ are two hyperparameters.

4. Experiments

4.1. Datasets and Evaluation Metrics

Encyclopedic-VQA. The Encyclopedic-VQA [36] dataset contains 221k question-answer pairs, each linked to up to five images and covering 16.7k fine-grained entities. The questions are categorized into *single-hop* and *two-hop* types: single-hop questions can be answered using information from a single Wikipedia page, whereas two-hop questions require sequential retrieval across multiple pages. The dataset is divided into training, validation, and test splits comprising 1M, 13.6k, and 5.8k samples. All experiments are conducted on the test split, which includes 4.8k single-hop questions. The dataset contains an external knowledge base derived from Wikipedia, consisting of approximately 2M pages. Each comprises the article title, its textual sections, and associated images. In our experiments, we employ the original 2M-page knowledge base provided with the dataset.

InfoSeek. The InfoSeek dataset [10] consists of approximately 1.3M image-question-answer triplets corresponding to around 11k distinct Wikipedia pages. It is partitioned into training, validation, and test splits, containing roughly 934k, 73k, and 348k samples. Both validation and test sets include questions about unseen entities. InfoSeek provides an external knowledge base of around 6M Wikipedia entities. Following previous works [7, 13, 55], experiments are conducted using a knowledge base of 100k pages.

Evaluation Metrics. We follow the original evaluation protocols provided with each dataset. For E-VQA, generated answers are evaluated using the BERT-based matching

³Further details are provided in the supplementary material.

Table 1. VQA accuracy scores on the Encyclopedic-VQA test set and the InfoSeek validation set. The marker † represents our reproductions, while gray color indicates models tested with non-comparable knowledge bases.

Model	Generator	Retriever	E-VQA		InfoSeek		
			Single-Hop	All	Unseen-Q	Unseen-E	All
BLIP-2 [27]	-	-	12.6	12.4	12.7	12.3	12.5
LLaVA-v1.5-7B [31]	-	-	16.3	16.9	9.6	9.4	9.5
LLaVA-MORE-8B [12]	-	-	16.0	16.9	8.3	8.9	7.8
Qwen2.5-VL-3B [4]	-	-	21.9	21.9	18.9	17.7	18.3
Qwen2.5-VL-7B [4]	-	-	23.6	23.2	22.8	24.1	23.7
DPR _{v+t} [25]	Multi-passage BERT	CLIP ViT-B/32	29.1	-	-	-	12.4
RORA-VLM [40]	LLaVA-v1.5-7B	CLIP ViT-L/14	-	20.3	25.1	27.3	-
EchoSight [55]	Mistral-7B/LLaMA-3-8B	EVA-CLIP-8B	41.8	-	-	-	31.3
CoMEM [54]	Qwen2.5-VL-7B	Custom VLM	-	-	32.8	28.5	-
Wiki-LLaVA [7]	LLaVA-v1.5-7B	CLIP ViT-L/14+Contriever	17.7	20.3	30.1	27.8	28.9
EchoSight [55]	LLaMA-3.1-8B	EVA-CLIP-8B	36.3	34.2	30.0	30.7	30.4
ReflectiVA [13]	LLaVA-MORE-8B	EVA-CLIP-8B	35.5	35.5	40.4	39.8	40.1
mR ² AG [59]	LLaVA-v1.5-7B	CLIP ViT-L/14	-	-	40.6	39.8	40.2
VLM-PRF [16]	LLaVA-MORE-8B	EVA-CLIP-8B	36.3	35.5	41.3	40.6	40.8
mKG-RAG [58]	LLaVA-MORE-8B	Custom VLM	38.4	36.3	41.4	39.6	40.5
ReflectiVA [13]†	Qwen2.5-VL-3B	EVA-CLIP-8B	33.7	35.2	39.6	38.1	38.9
VLM-PRF [16]	Qwen2.5-VL-3B	EVA-CLIP-8B	31.1	32.4	39.7	38.8	39.0
ReAG (Ours)	Qwen2.5-VL-3B	EVA-CLIP-8B	41.3	42.9	43.7	42.9	43.3
			Δ+7.6	Δ+7.7	Δ+4.0	Δ+4.1	Δ+4.3
ReflectiVA [13]†	Qwen2.5-VL-7B	EVA-CLIP-8B	36.8	36.8	43.5	44.3	43.9
VLM-PRF [16]	Qwen2.5-VL-7B	EVA-CLIP-8B	37.1	36.0	43.3	42.7	42.8
VLM-PRF [16]	InternVL3-8B	EVA-CLIP-8B	40.1	39.2	43.5	42.1	42.5
ReAG (Ours)	Qwen2.5-VL-7B	EVA-CLIP-8B	44.9	47.0	48.3	46.2	47.2
			Δ+4.8	Δ+7.8	Δ+4.8	Δ+1.9	Δ+3.3

score (BEM) [5], which measures semantic similarity of predicted and ground-truth answers. For InfoSeek, evaluation depends on the question type: we employ standard VQA accuracy [14] as well as relaxed accuracy [37].

4.2. Implementation Details

Retrieval Details. To retrieve potentially informative documents for a query image, we employ EVA-CLIP-8B [48]. In the coarse-grained stage, the entire query image is encoded through EVA-CLIP to perform large-scale retrieval over the knowledge base. For InfoSeek, we use an image-to-text retrieval setup that computes similarity between the query image and document metadata (*i.e.*, the title of the page and the summary). For Encyclopedic-VQA, we adopt image-to-image retrieval, comparing the query image with the images inside Wikipedia pages. In the fine-grained stage, we extract the main visual subject mentioned in the question using the spaCy library⁴ and localize it in the image via GroundingDINO [32], whose bounding box is re-encoded through EVA-CLIP. Retrieval is done using the FAISS library [21], with the top- k results with $k = 20$ retained at each stage.

Critic Model and Dataset. Independently from the generator scale, our critic model builds upon Qwen2.5-VL-3B, fine-

tuned on a curated subset of the ReflectiVA dataset [13]⁵. The model is trained for 1 epoch with a learning rate of 2×10^{-6} and a global batch size of 32.

Generator Training. We build two versions of our generator, both based on Qwen2.5-VL [4], using the 3B and 7B model variants, and optimize them using the SFT plus RL training scheme. In the SFT phase, we use the same ReflectiVA [13] subset employed to train the critic model, collecting reasoning traces from Qwen2.5-VL-7B. We set $\alpha = 0.8$ to give more importance to final-answer tokens. The generator is trained for one epoch with SFT loss using AdamW [34], a learning rate of 2×10^{-6} and an effective batch size of 128. For RL post-training, we use the full Encyclopedic-VQA and InfoSeek sets from ReflectiVA, excluding samples from LLaVA-Instruct [31]. Each batch includes 128 prompts with 8 completions per prompt. Training is conducted with Adam [22], a learning rate 1×10^{-6} . Rewards weigh accuracy $\gamma = 1.0$ over format $\delta = 0.2$. In all our experiments, we update the MLP adapter and LLM weights while keeping the vision encoder frozen.

4.3. Comparison with the State of the Art

Main Results. We present a comprehensive comparison of ReAG on the E-VQA test set and the InfoSeek validation

⁴<https://github.com/explosion/spaCy>

⁵Further details on the training subset are given in the supplementary.

Table 2. VQA accuracy scores on Encyclopedic-VQA and InfoSeek with OMGM as retrieval modality.

Model	Generator	E-VQA		InfoSeek	
		Single-Hop	Un-Q	Un-E	All
ReflectiVA [13]	LLaVA-MORE-8B	41.8	33.8	34.5	34.1
OMGM [56]	LLaVA-v1.5-7B	50.2	43.5	43.5	43.5
ReflectiVA [13]	Qwen2.5-VL-3B	44.6	40.5	41.6	41.1
ReAG (Ours)	Qwen2.5-VL-3B	49.1	47.2	45.3	46.2
ReflectiVA [13]	Qwen2.5-VL-7B	44.0	43.3	44.0	43.6
ReAG (Ours)	Qwen2.5-VL-7B	52.5	50.3	48.2	49.2

Table 3. VQA accuracy scores on Encyclopedic-VQA and InfoSeek with oracle Wikipedia pages.

Model	Generator	E-VQA		InfoSeek	
		Single-Hop	Un-Q	Un-E	All
Qwen2.5-VL-3B [4]	Qwen2.5-3B	72.1	47.0	43.0	44.9
Qwen2.5-VL-7B [4]	Qwen2.5-7B	78.3	41.6	41.3	41.4
ReflectiVA [13]	Qwen2.5-VL-3B	72.9	53.4	53.9	53.7
ReAG (Ours)	Qwen2.5-VL-3B	79.0	55.1	53.3	54.2
Wiki-LLaVA [7]	LLaVA-v1.5-7B	38.5	52.7	50.3	51.5
ReflectiVA [13]	LLaVA-MORE-8B	75.2	57.8	57.4	57.6
ReflectiVA [13]	Qwen2.5-VL-7B	71.3	56.0	56.0	56.0
ReAG (Ours)	Qwen2.5-VL-7B	81.5	60.7	58.9	59.7

set against both zero-shot MLLMs and retrieval-augmented baselines. Specifically, we evaluate BLIP-2 [27], LLaVA-v1.5 [31], LLaVA-MORE [12], and Qwen2.5-VL [4] in a zero-shot setting, where the models receive only the query image and question as input. We further include retrieval-augmented approaches such as DPR [25], RORA-VLM [40], EchoSight [55], COMEM [54], WikiLLaVA [7], mR²AG [59], mKG-RAG [58], ReflectiVA [13], and VLM-PRF [16]. For fairness, we reproduce ReflectiVA using Qwen2.5-VL backbones at both 3B and 7B scales.

As shown in Table 1, zero-shot MLLMs, which rely solely on internal knowledge, are unable to accurately answer questions in knowledge-intensive benchmarks, underscoring the need for external retrieval. With the introduction of a retrieval pipeline, performance improves substantially. For example, on InfoSeek, the overall accuracy rises from around 20% for the zero-shot Qwen2.5-VL-7B model to roughly 40% with retrieval-augmented methods such as mKG-RAG. ReAG further advances these results, achieving state-of-the-art performance on both E-VQA and InfoSeek across model scales. Specifically, on E-VQA ReAG yields a +7.7 point gain over ReflectiVA when using Qwen2.5-VL-3B and a +7.8 point improvement over VLM-PRF when leveraging the stronger InternVL3-8B backbone. Similar gains are observed on InfoSeek, with overall improvements of +4.3 and +3.3 points for Qwen2.5-VL-3B and 7B, respectively. These consistent improvements across both datasets demonstrate the effectiveness and robustness of our approach.

Results with OMGM Retrieval Mode. We also compare against the OMGM framework [56], which adopts a coarse-

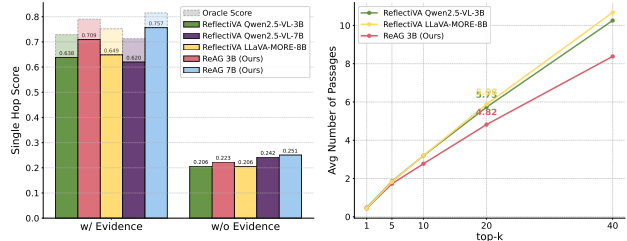


Figure 3. Comparison of performance on E-VQA with and without evidence, including oracle upper bounds (left). Analysis on average number of passages retained at different top- k values (right).

to-fine, multi-stage retrieval strategy, leveraging an image-to-text summary retriever in the first step. To ensure a fair comparison, in Table 2 we evaluate our method using the same retrieval modality. Notably, our method consistently outperforms OMGM [56] across both E-VQA and InfoSeek benchmarks. With the Qwen2.5-VL-3B generator, our approach achieves substantial improvements over ReflectiVA [13], yielding gains of +4.5 and +5.1 points on E-VQA and InfoSeek. When scaling to Qwen2.5-VL-7B, performance further increases, reaching 52.5 on E-VQA and 49.2 on InfoSeek, surpassing OMGM by 2.3 and 5.7 points, respectively. These results indicate that, even when using only the initial retrieval step of OMGM, our critic-based filtering, together with ReAG reasoning capabilities, leads to consistently higher performance compared to both prior methods and the full multi-stage retrieval of OMGM.

Results with Oracle Documents. We also experiment under an oracle setting (Table 3), where the ground-truth document (*i.e.*, the Wikipedia page corresponding to the query) is provided. We compare results from zero-shot models (Qwen2.5-VL [4] in both 3B and 7B variants), which take the entire Wikipedia pages as input, and retrieval-based methods (WikiLLaVA [7], ReflectiVA [13], and ReAG), which process retrieved passages through an additional model-specific filtering stage. In this configuration, ReAG receives all passages from the oracle document, which are then passed to the critic model for filtering before being fed to the generator.

Notably, ReAG achieves the best performance across both E-VQA and InfoSeek at all model scales. On the 3B variant, ReAG outperforms ReflectiVA by +6.1 points on E-VQA, while on Infoseek the 7B version consistently improves over ReflectiVA by +3.7 and still surpasses it by +2.1 even when ReflectiVA employs a larger generator (LLaVA-MORE-8B).

Retrieval and Generation Pipeline Analysis. The performance of RAG-based approaches strongly depends on the presence of the evidence passage in the retrieved set, and on the number of passages provided to the generator.

In Fig. 3 (left), we evaluate the ability to produce the correct answer when the evidence passage is either present or absent in the context. ReAG consistently outperforms all competitors of comparable scale, demonstrating that our

Table 4. Ablation study results on Encyclopedic-VQA and InfoSeek to validate the effectiveness of our model component.

Generator	Retrieval Pipeline			Generation Pipeline				E-VQA	InfoSeek		
	Coarse-grained	Fine-grained	Critic	SFT	Reasoning	Traces	RL	Single-Hop	Unseen-Q	Unseen-E	All
	-	-	-	-	-	-	-	21.9	18.9	17.7	18.3
	✓	-	-	-	-	-	-	19.2	10.2	10.0	10.1
	✓	-	✓	-	-	-	-	38.0	27.9	26.1	27.0
	✓	✓	✓	-	-	-	-	40.2	28.1	26.1	27.1
	✓	✓	✓	✓	-	-	-	39.3	37.9	37.1	37.5
	✓	✓	✓	✓	✓	-	-	38.1	41.9	40.6	41.3
	✓	✓	✓	✓	-	-	✓	39.5	39.8	39.4	39.6
ReAG (Ours) Qwen2.5-VL-3B	✓	✓	✓	✓	✓	✓	✓	41.3	43.7	42.9	43.3
	✓	✓	✓	-	-	-	-	41.7	29.3	27.8	28.5
	✓	✓	✓	✓	✓	-	-	42.0	42.0	40.7	41.4
ReAG (Ours) Qwen2.5-VL-7B	✓	✓	✓	✓	✓	✓	✓	44.9	48.3	46.2	47.2

Q: What is the closest upper taxonomy of this bird?



Qwen2.5-VL-7B (ZS) [4]:
The closest taxonomy of this bird is the family Laridae. ✗
ReflectiVA [13]:
Sterna ✗
ReAG (Ours):
Thalasseus ✓

Q: Which road, railway or canal does this bridge carry?



Qwen2.5-VL-7B (ZS) [4]:
The bridge in the picture is Blackfriars Bridge. ✗
ReflectiVA [13]:
Road ✗
ReAG (Ours):
A201 ✓

Q: What is the brand of this vehicle?



Qwen2.5-VL-7B (ZS) [4]:
The vehicle in the picture is a Ford. This can be determined [...] ✗
ReflectiVA [13]:
Ford ✗
ReAG (Ours):
Isuzu ✓

Figure 4. Qualitative results on InfoSeek image-question pairs comparing ReAG, ReflectiVA [13], and the corresponding zero-shot model.

reasoning-enhanced approach is robust even in the absence of direct evidence. Each model is also accompanied by its oracle performance, clearly showing that ReAG consistently gets closer to the oracle upper bound than other approaches.

In Fig. 3 (right), we instead report the average number of passages passed to the generation at different k values, comparing the filtering behavior of ReflectiVA [13] with that of our critic model. ReAG reduces noise introduced in the generator context by achieving a reduction of 18.0% and 15.9% in the number of passages compared to ReflectiVA based on LLaVA-MORE-8B and Qwen2.5VL-3B respectively, further emphasizing the advantage of our filtering pipeline.

Qualitative Results. In Fig. 4, we present a qualitative comparison on image-question pairs from InfoSeek. Notably, the zero-shot model tends to produce longer and detailed answers, whereas ReflectiVA and ReAG generate responses that follow the dataset-specific format. Overall, the results consistently demonstrate that ReAG answers accurate and outperforms competing approaches.

4.4. Ablation Studies

We finally perform an ablation study by progressively enabling key components of the final architecture. Table 4 reports results at each stage with k equal to 20, evaluating the effect of the critic and fine-grained retrieval modules, as well as design choices in the generation pipeline.

We first examine the zero-shot setup under different retrieval configurations, using the 3B-scale model. As shown in the first two rows of Table 4, directly passing all passages from the top-20 documents into the generator severely de-

grades performance due to excessive noise. Introducing the critic model (third row) effectively filters irrelevant passages, yielding more than a twofold gain. Adding the fine-grained retriever (fourth row) further improves retrieval recall and yields additional gains on both datasets.

Fixing the retrieval pipeline to its final configuration, we then assess the impact of different generation strategies. As shown in [15], introducing a cold-start phase can help prepare the model for reasoning by exposing it to intermediate traces before full supervision. Results show that applying reinforcement learning after this cold-start phase outperforms standard SFT, indicating that the cold-start phase effectively prepares the model for multi-step reasoning, allowing the RL algorithm to operate on a model already prepared for structured thinking. A similar trend is observed with the 7B variant, where both training phases contribute significantly to the final performance, further validating the robustness of the proposed ReAG pipeline.

5. Conclusion

We have presented ReAG, a multimodal retrieval-augmented approach designed for the KB-VQA task. Our method demonstrates that incorporating reasoning and filtering retrieved passages significantly improves answer quality by reducing the noise introduced by irrelevant passages and producing a structured reasoning process when generating the final answer. Extensive experimental results on both Encyclopedic-VQA and Infoseek show that ReAG consistently outperforms existing methods, confirming the effectiveness of our retrieval and reasoning strategies.

Acknowledgments

We acknowledge the CINECA award under the IS CRA initiative, for the availability of high-performance computing resources. This work has been supported by the EU Horizon project “ELLIOT - European Large Open Multi-Modal Foundation Models For Robust Generalization On Arbitrary Data Streams” (No. 101214398), by the EuroHPC JU project “MINERVA” (GA No. 101182737), and by the PRIN 2022-PNRR project “MU CES” (CUP E53D23016290001) funded by the EU - NextGenerationEU.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a Visual Language Model for Few-Shot Learning. In *NeurIPS*, 2022. 1
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual Question Answering. In *ICCV*, 2015. 1
- [3] Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hananeh Hajishirzi. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. In *ICLR*, 2024. 2
- [4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2.5-VL Technical Report. *arXiv preprint arXiv:2502.13923*, 2025. 1, 6, 7, 8, 12, 13, 17, 18
- [5] Jannis Bulian, Christian Buck, Wojciech Gajewski, Benjamin Boerschinger, and Tal Schuster. Tomayto, Tomahto. Beyond Token-level Answer Equivalence for Question Answering Evaluation. *arXiv preprint arXiv:2202.07654*, 2022. 6
- [6] Davide Caffagni, Federico Cocchi, Luca Barsellotti, Nicholas Moratelli, Sara Sarto, Lorenzo Baraldi, Lorenzo Baraldi, Marcella Cornia, and Rita Cucchiara. The Revolution of Multimodal Large Language Models: A Survey. In *ACL Findings*, 2024. 1
- [7] Davide Caffagni, Federico Cocchi, Nicholas Moratelli, Sara Sarto, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Wiki-LLaVA: Hierarchical Retrieval-Augmented Generation for Multimodal LLMs. In *CVPR Workshops*, 2024. 2, 5, 6, 7
- [8] Davide Caffagni, Sara Sarto, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Recurrence Meets Transformers for Universal Multimodal Retrieval. *arXiv preprint arXiv:2509.08897*, 2025. 1
- [9] Lluís Castrejon, Thomas Mensink, Howard Zhou, Vittorio Ferrari, Andre Araujo, and Jasper Uijlings. HAMMR: Hierarchical MultiModal React agents for generic VQA. *arXiv preprint arXiv:2404.05465*, 2024. 13
- [10] Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, Alan Ritter, and Ming-Wei Chang. Can Pre-trained Vision and Language Models Answer Visual Information-Seeking Questions? In *EMNLP*, 2023. 1, 2, 5
- [11] Yi Chen, Yuying Ge, Rui Wang, Yixiao Ge, Junhao Cheng, Ying Shan, and Xihui Liu. Grpo-care: Consistency-aware reinforcement learning for multimodal reasoning. *arXiv preprint arXiv:2506.16141*, 2025. 2
- [12] Federico Cocchi, Nicholas Moratelli, Davide Caffagni, Sara Sarto, Lorenzo Baraldi, Marcella Cornia, and Rita Cucchiara. LLaVA-MORE: A Comparative Study of LLMs and Visual Backbones for Enhanced Visual Instruction Tuning. In *ICCV Workshops*, 2025. 6, 7
- [13] Federico Cocchi, Nicholas Moratelli, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Augmenting Multimodal LLMs with Self-Reflective Tokens for Knowledge-based Visual Question Answering. In *CVPR*, 2025. 1, 2, 4, 5, 6, 7, 8, 12, 13, 14, 15, 16, 17, 18
- [14] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in VQA Matter: Elevating the Role of Image Understanding in Visual Question Answering. In *CVPR*, 2017. 1, 6
- [15] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv preprint arXiv:2501.12948*, 2025. 2, 4, 8
- [16] Yuyang Hong, Jiaqi Gu, Qi Yang, Lubin Fan, Yue Wu, Ying Wang, Kun Ding, Shiming Xiang, and Jieping Ye. Knowledge-based Visual Question Answer with Multimodal Processing, Retrieval and Filtering. In *NeurIPS*, 2025. 1, 2, 6, 7
- [17] Tianze Hua, Tian Yun, and Ellie Pavlick. How do vision-language models process conflicting information across modalities? *arXiv preprint arXiv:2507.01790*, 2025. 1
- [18] Hugging Face. Open R1: A fully open reproduction of DeepSeek-R1, 2025. 12
- [19] Dongzhi Jiang, Renrui Zhang, Ziyu Guo, Yanmin Wu, Jiayi Lei, Pengshuo Qiu, Pan Lu, Zehui Chen, et al. MMSearch: Benchmarking the Potential of Large Models as Multi-modal Search Engines. In *ICLR*, 2024. 2
- [20] Bowen Jin, Hansi Zeng, Zhenrui Yue, Jinsung Yoon, Serkan Arik, Dong Wang, Hamed Zamani, and Jiawei Han. Search-r1: Training llms to reason and leverage search engines with reinforcement learning. In *COLM*, 2025. 2
- [21] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-Scale Similarity Search with GPUs. *IEEE Trans. on Big Data*, 7(3): 535–547, 2019. 6
- [22] Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015. 6
- [23] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *ACM SIGOPS*, 2023. 12
- [24] Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, et al. Rlaif vs. rlhf: Scaling reinforcement learning from human feedback with ai feedback. In *ICML*, 2023. 2
- [25] Paul Lerner, Olivier Ferret, and Camille Guinaudeau. Cross-modal Retrieval for Knowledge-based Visual Question Answering. In *ECIR*, 2024. 6, 7

- [26] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *NeurIPS*, 2020. 1
- [27] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping Language-Image Pre-training with Frozen Image Encoders and Large Language Models. In *ICML*, 2023. 6, 7
- [28] Zongxia Li, Wenhao Yu, Chengsong Huang, Rui Liu, Zhenwen Liang, Fuxiao Liu, Jingxi Che, Dian Yu, et al. Self-rewarding vision-language model via reasoning decomposition. *arXiv preprint arXiv:2508.19652*, 2025. 2
- [29] Weizhe Lin, Jingbiao Mei, Jinghong Chen, and Bill Byrne. PreFLMR: Scaling Up Fine-Grained Late-Interaction Multimodal Retrievers. In *ACL*, 2024. 1
- [30] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual Instruction Tuning. In *NeurIPS*, 2023. 1
- [31] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved Baselines with Visual Instruction Tuning. In *CVPR*, 2024. 1, 6, 7
- [32] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection. *arXiv preprint arXiv:2303.05499*, 2023. 6
- [33] Zichen Liu, Changyu Chen, Wenjun Li, Penghui Qi, Tianyu Pang, Chao Du, Wee Sun Lee, and Min Lin. Understanding R1-Zero-Like Training: A Critical Perspective. In *COLM*, 2025. 2, 15
- [34] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. In *ICLR*, 2019. 6
- [35] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-VQA: A Visual Question Answering Benchmark Requiring External Knowledge. In *CVPR*, 2019. 1, 2
- [36] Thomas Mensink, Jasper Uijlings, Lluís Castrejon, Arushi Goel, Felipe Cadar, Howard Zhou, Fei Sha, André Araujo, and Vittorio Ferrari. Encyclopedic VQA: Visual Questions About Detailed Properties of Fine-Grained Categories. In *ICCV*, 2023. 1, 2, 5
- [37] Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. PlotQA: Reasoning over Scientific Plots. In *WACV*, 2020. 6
- [38] Kartik Narayan, Yang Xu, Tian Cao, Kavya Nerella, Vishal M Patel, Navid Shiee, Peter Grasch, Chao Jia, Yinfei Yang, and Zhe Gan. DeepMMSearch-R1: Empowering Multimodal LLMs in Multimodal Web Search. *arXiv preprint arXiv:2510.12801*, 2025. 2
- [39] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training Language Models to Follow Instructions with Human Feedback. In *NeurIPS*, 2022. 2
- [40] Jingyuan Qi, Zhiyang Xu, Rulin Shao, Yang Chen, Jing Di, Yu Cheng, Qifan Wang, and Lifu Huang. RoRA-VLM: Robust Retrieval-Augmented Vision Language Models. *arXiv preprint arXiv:2410.08876*, 2024. 2, 6, 7
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning Transferable Visual Models from Natural Language Supervision. In *ICML*, 2021. 2
- [42] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*, 2023. 2
- [43] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *IEEE SC20*, 2020. 12
- [44] Sara Sarto, Marcella Cornia, and Rita Cucchiara. Image Captioning Evaluation in the Age of Multimodal LLMs: Challenges and Future Perspectives. In *IJCAI*, 2025. 1
- [45] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-OKVQA: A Benchmark for Visual Question Answering Using World Knowledge. In *ECCV*, 2022. 2
- [46] Sanket Shah, Anand Mishra, Naganand Yadati, and Partha Pratim Talukdar. KVQA: Knowledge-aware Visual Question Answering. In *AAAI*, 2019. 2
- [47] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024. 2, 5
- [48] Quan Sun, Jinsheng Wang, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, and Xinlong Wang. EVA-CLIP-18B: Scaling CLIP to 18 Billion Parameters. *arXiv preprint arXiv:2402.04252*, 2024. 2, 6
- [49] Leandro von Werra, Younes Belkada, Lewis Tunstall, Edward Beeching, Tristan Thrush, Nathan Lambert, Shengyi Huang, Kashif Rasul, and Quentin Galloudec. TRL: Transformer Reinforcement Learning. <https://github.com/huggingface/trl>, 2020. 12
- [50] Cong Wei, Yang Chen, Haonan Chen, Hexiang Hu, Ge Zhang, Jie Fu, Alan Ritter, and Wenhui Chen. UniIR: Training and Benchmarking Universal Multimodal Information Retrievers. In *ECCV*, 2024. 1
- [51] Lai Wei, Yuting Li, Chen Wang, Yue Wang, Linghe Kong, Weiran Huang, and Lichao Sun. Unsupervised post-training for multi-modal llm reasoning via grpo. In *NeurIPS*, 2025. 2
- [52] Jinming Wu, Zihao Deng, Wei Li, Yiding Liu, Bo You, Bo Li, Zejun Ma, and Ziwei Liu. MMSearch-R1: Incentivizing LMMs to Search. *arXiv preprint arXiv:2506.20670*, 2025. 2
- [53] Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton Van Den Hengel. Visual question answering: A survey of methods and datasets. *CVIU*, 163:21–40, 2017. 1
- [54] Wenyi Wu, Zixuan Song, Kun Zhou, Yifei Shao, Zhiting Hu, and Biwei Huang. Towards General Continuous Memory for Vision-Language Models. *arXiv preprint arXiv:2505.17670*, 2025. 1, 6, 7
- [55] Yibin Yan and Weidi Xie. EchoSight: Advancing Visual-Language Models with Wiki Knowledge. In *EMNLP Findings*, 2024. 1, 2, 5, 6, 7

- [56] Wei Yang, Jingjing Fu, Rui Wang, Jinyu Wang, Lei Song, and Jiang Bian. OMGM: Orchestrate Multiple Granularities and Modalities for Efficient Multimodal Retrieval. In *ACL*, 2025. [7](#)
- [57] Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. DAPO: An Open-Source LLM Reinforcement Learning System at Scale. *arXiv preprint arXiv:2503.14476*, 2025. [2](#), [5](#)
- [58] Xu Yuan, Liangbo Ning, Wenqi Fan, and Qing Li. mKG-RAG: Multimodal Knowledge Graph-Enhanced RAG for Visual Question Answering. *arXiv preprint arXiv:2508.05318*, 2025. [1](#), [6](#), [7](#)
- [59] Tao Zhang, Ziqi Zhang, Zongyang Ma, Yuxin Chen, Zhonggang Qi, Chunfeng Yuan, Bing Li, Junfu Pu, Yuxuan Zhao, Zehua Xie, et al. mR²AG: Multimodal Retrieval-Reflection-Augmented Generation for Knowledge-Based VQA. *arXiv preprint arXiv:2411.15041*, 2024. [1](#), [6](#), [7](#), [13](#)
- [60] Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. LIMA: Less Is More for Alignment. In *NeurIPS*, 2023. [2](#)

ReAG: Reasoning-Augmented Generation for Knowledge-based Visual Question Answering

Supplementary Material

A. Additional Implementation Details

This section contains detailed descriptions of our implementation, including model architectures details, training setup, and evaluation protocols.

Retrieval Details. As mentioned in the main paper, to retrieve candidate documents for each query image, we adopt a two-stage retrieval pipeline based on EVA-CLIP-8B. In the coarse-grained stage, the full query image is encoded and matched against the knowledge base. Instead, in the fine-grained retrieval phase, the subject of the question, corresponding to the visual entity in the image, is extracted using SpaCy⁶. According to the structure of query questions, extraction prioritizes noun phrases starting with demonstratives like “this” or “these”, followed by nouns serving as objects of prepositions (e.g., “of”, “by”, “in”, “from”, etc.). If neither pattern is found, the last noun in the question is used as a fallback. This approach ensures that the most relevant entity is reliably identified for retrieval and reasoning.

Critic Model and Dataset. The critic model builds upon Qwen2.5-VL-3B and is fine-tuned on a curated subset of the ReflectiVA dataset [13]. Specifically, we select 1M samples, balanced between InfoSeek and Encyclopedic-VQA, where each sample is paired with a passage labeled as relevant or irrelevant. To encourage more robust discrimination, irrelevant passages are further divided into *soft negatives* (i.e., semantically related but unhelpful passages) and *hard negatives* (i.e., irrelevant passages that exhibit high similarity to the relevant ones), in proportions of 30% and 70% respectively. During training, we fine-tune only the visual projector and the LLM, while keeping the vision encoder frozen.

At inference time, given (I_q, q, p) (where I_q and q are the input image and question, and p is the passage to evaluate), we apply a *yes*-probability threshold equal to 0.1 (cf. Eq. 4 of the main paper). This conservative threshold ensures that the critic \mathcal{C} filters out only those passages for which it is highly confident they are irrelevant.

Generator Training. Our generator is based on Qwen2.5-VL [4], employing the 3B and 7B scale. As mentioned in the main paper, we employ a multi-stage training strategy: first an SFT stage that serves as cold start for a subsequent RL-based training stage. Specifically, for both the cold-start and RL stages, we apply a cosine learning rate schedule with 450 and 150 warm-up steps, respectively. Weight decay is set to 0.01 during cold-start and removed during RL. SFT

fine-tuning of the 3B and 7B models requires roughly 24 hours on 16 and 64 NVIDIA A100 64GB GPUs, respectively. In the RL stage, completions are generated with a temperature equal to 1.0 and a repetition penalty of 1.05, using vLLM [23].

During training with the custom GRPO loss (cf. Eq. 6 of the main paper), \mathcal{G}_θ and $\mathcal{G}_{\theta_{\text{old}}}$ share the same weights, although they operate as separate models. $\mathcal{G}_{\theta_{\text{old}}}$ runs inside the vLLM worker and remains frozen, and gradients are applied only to \mathcal{G}_θ . After each backward pass, we synchronize $\mathcal{G}_{\theta_{\text{old}}}$ with the updated weights of \mathcal{G}_θ . Fine-tuning the 3B and 7B models with our RL strategy takes roughly 48 hours on 32 and 64 NVIDIA A100 64GB GPUs, respectively. We select the best checkpoint based on the best task-specific accuracy on a held-out validation split. All runs employ DeepSpeed ZeRO-3 [43] and gradient checkpointing. Our training codebase builds on Open-R1 [18] and TRL [49].

B. Reward Design

As discussed in Sec. 3.4 of the main paper, we employ a task-specific accuracy reward. The reward function evaluates only the final answer rather than the intermediate reasoning. To extract the predicted answer, we first search for content enclosed within the `<answer></answer>` tags. If no such content is found, we extract all text following the first `<answer>` tag. If this is unsuccessful, we instead use the text following the `</think>` tag. When none of these patterns appear, the entire model output is used.

In every case, format-specific special tokens are removed. We then apply the same normalization procedure used in the InfoSeek and Encyclopedic-VQA evaluations, including the removal of articles, punctuation, extra whitespace, and capitalization, along with standardization of digits and contractions. The final task-specific reward depends on the source dataset and task type, and is computed as follows:

$$R_{\text{task}}(\tilde{o}_i, o_i^*, \tau_i) = \begin{cases} \mathbb{1}[\Psi_{\text{num}}(\tilde{o}_i, o_i^*)], & \text{if } \tau_i = \text{numerical}, \\ \mathbb{1}[\text{IoU}(\tilde{o}_i, o_i^*) \geq 0.5], & \text{if } \tau_i = \text{multi}, \\ \mathbb{1}[\tilde{o}_i = o_i^*], & \text{otherwise.} \end{cases}$$

where \tilde{o}_i , o_i^* and τ_i denote respectively the extracted prediction, the ground-truth answer and the task type of the i -th sample, and Ψ_{num} evaluates success or failure in numerical match. When multiple alternative ground-truths are provided for a sample, we compute the reward with respect to each and take the maximum. For samples from InfoSeek, we use

⁶In particular, we use the `en_core_web_sm` model available at <https://spacy.io/models/en>.




Figure 5. Task-specific accuracy reward progression across training iterations of the ReAG 7B generator.

exact string matching for *entity* and *time* questions, while *numerical* questions are evaluated with Ψ_{num} :

$$\Psi_{\text{num}}(\tilde{o}, o^*) = \begin{cases} |\tilde{o} - o^*| \leq 0.1, & \text{if } \text{is_scalar}(\tilde{o}) \\ & \wedge \text{is_scalar}(o^*), \\ \tilde{o} \in o^*, & \text{if } \text{is_scalar}(\tilde{o}) \\ & \wedge \text{is_interval}(o^*), \\ \text{IoU}(\tilde{o}, o^*) \geq 0.5, & \text{if } \text{is_interval}(\tilde{o}) \\ & \wedge \text{is_interval}(o^*). \end{cases}$$

For samples from Encyclopedic-VQA dataset, we adopt exact match scoring for *single-answer* questions, while for *multi-answer* questions the prediction is rewarded as correct only if intersection-over-union between predicted and ground-truth items reaches or surpasses 0.5.

The evolution of task-specific accuracy reward during training is reported in Fig. 5.

C. Additional Experimental Results

In this section, we provide additional experiments and analyses that complement the results reported in the main paper.

C.1. Results with Google Lens Retriever

We further extend the analysis in the main paper by evaluating all methods under an alternative retrieval setup. Specifically, in Table 5, we employ the Wikipedia pages retrieved by Google Lens⁷ when provided with the query image for each question of Encyclopedic-VQA, which have been officially released along with the dataset. Even though ReAG does not use the fine-grained retriever in this setting, it consistently outperforms ReflectiVA [13] across different generator scales. Notably, ReAG at the 3B scale performs comparable to ReflectiVA at the 8B scale and HAMMR [9] at 55B. In addition, the stronger reasoning capabilities of ReAG allow it to benefit substantially from improved retrieval quality, improving single-hop accuracy from 48.0 (3B) to 55.5 (7B), showing a notable gain of +7.5 points.

⁷A visual recognition service by Google, available at <https://lens.google.com/>.

Table 5. VQA accuracy scores on the Encyclopedic-VQA test set with Google Lens employed as retriever.

Model	Generator	E-VQA	
		Single-Hop	All
Qwen2.5-VL-3B [4]	-	35.2	-
Qwen2.5-VL-7B [4]	-	44.4	-
HAMMR [9]	PaLI-X-55B	47.8	-
mR ² AG [59]	LLaVA-v1.5-7B	55.9	-
ReflectiVA [13]	LLaVA-MORE-8B	48.2	46.1
ReflectiVA [13]	Qwen2.5-VL-3B	45.8	43.2
ReAG (Ours)	Qwen2.5-VL-3B	48.0	48.0
ReflectiVA [13]	Qwen2.5-VL-7B	48.3	46.8
ReAG (Ours)	Qwen2.5-VL-7B	55.5	56.9

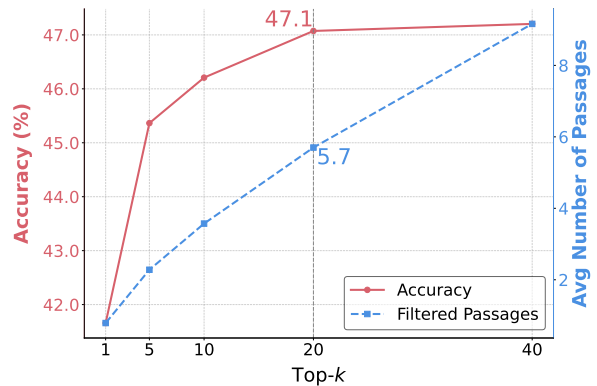


Figure 6. Performance of ReAG 7B (red) and average number of filtered passages (blue) when varying the number k of retrieved documents. Accuracy and number of passages are computed as the average of E-VQA and InfoSeek scores.

C.2. Varying the Number of Retrieved Documents

In Fig. 6, we analyze the effect of varying the number of retrieved documents k on the overall performance and on the average number of filtered passages fed to the generator. As shown, the model achieves the best results around $k = 20$, which represents the optimal trade-off between coverage and noise, and is therefore adopted as the default retrieval depth in our pipeline. Retrieving too few documents results in insufficient contextual evidence, causing a drop in recall and limiting the ability of the model to access the necessary information. Conversely, increasing k beyond this point does not yield meaningful performance gains while substantially inflating the computational cost of the filtering stage.

C.3. Effectiveness of the Critic Model

Critic Analysis. In Fig. 7, we provide a detailed analysis on the effectiveness of the proposed critic model, employed in ReAG to filter relevant passages. Specifically, the plot reports the average number of passages retained after the filtering performed by the critic model when varying the number k of retrieved documents. Across all retrieval sizes, the critic model substantially reduces the number of retained

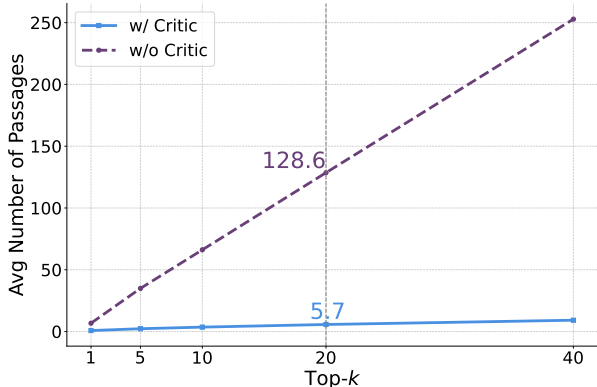


Figure 7. Comparison of the average number of passages fed to the generator with and without the critic filtering.

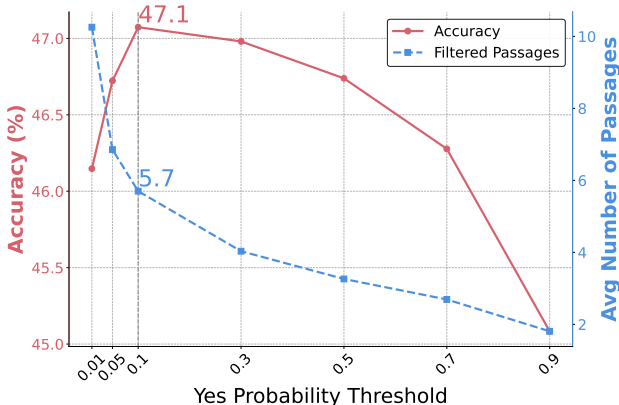


Figure 8. Performance of ReAG 7B (red) and average number of filtered passages (blue) when varying the *yes*-probability threshold in our critic model. Accuracy and number of passages are computed as the average of E-VQA and InfoSeek scores.

passages (e.g., from an average of 128.6 to 5.7 at $k = 20$), while preserving answer-relevant information. This highlights the strong ability of the critic model to discard noisy or off-topic passages, leading to a more compact and semantically aligned evidence set for multimodal reasoning.

Critic Threshold. In Fig. 8, we report how the performance varies as a function of the *yes*-probability threshold used in our critic model (cf. Eq. 4). The results show that instead of simply letting the fine-tuned MLLM decide if the passage is relevant or not (i.e., $\text{thresh} = 0.5$), leveraging the confidence of the model in predicting the “Yes” token allows us to gain more control over the filtering phase. A threshold $\text{thresh} = 0.1$ yields the best trade-off between precision and recall in retrieving relevant passages. This setting enables the critic model to reliably filter out passages for which it is most confident of their non-relevance to the query.

Filtering Analysis. To evaluate the quality of our filtering stage, we report passage-level statistics measuring evidence retention (Recall) and irrelevant-passage recognition (Specificity) in Table 6. Our Critic retains relevant evidence with higher recall while filtering irrelevant passages with higher

Table 6. Recall and specificity for passage filtering on E-VQA.

Model		Recall	Specificity
ReflectiVA [13]	LLaVA-MORE-8B	89.6	93.8
ReflectiVA [13]	Qwen2.5-VL-3B	91.3	93.8
ReflectiVA [13]	Qwen2.5-VL-7B	92.3	95.1
ReAG (Ours)	Qwen2.5-VL-3B	94.6	95.7

specificity compared to ReflectiVA. As a result, the generator receives a cleaner context with fewer passages on average, reducing noise while preserving useful evidence.

C.4. Sources of Performance Gains

Impact of Evidence Quality. As shown in the ablation study (Table 4), equipping the generator with our retrieval pipeline – composed of the Critic and fine-grained retrieval – already yields substantial gains in a Zero-Shot setting, reaching 40.2 on E-VQA and 27.1 on InfoSeek. This indicates that a significant portion of the improvements stems from higher-quality evidence selection rather than parametric memorization.

Training for Visual Robustness. The SFT+RL procedure further improves performance by teaching the generator to leverage retrieved evidence when available and rely on visual reasoning otherwise. To encourage this behavior, the training data includes both relevant passages and distractors (30% soft negatives and 70% hard negatives), forcing the model to distinguish useful evidence from noise.

Finally, the relatively small differences in the *w/o evidence* setting (Fig. 3, left) are partly due to a design choice in our retrieval pipeline. Unlike [13], which inserts a random passage when all evidence is filtered out, our method provides no fallback context, avoiding the risk of conditioning the generator on misleading information.

D. Critic and Generator Prompts

Critic Prompt Design. The prompt illustrates the instruction used to query the critic model. For each image-question pair (I_q, q) and a candidate passage p , the critic is explicitly asked to determine whether p contains any information that could help answer the question. The formulation uses a minimal, binary response space (i.e., “Yes”/“No”), which simplifies supervision and ensures consistent outputs across diverse samples. This concise design encourages the model to focus on relevance estimation rather than generative reasoning, enabling more stable fine-tuning and robust filtering of noisy passages during retrieval.

Critic System Prompt

You are a multimodal reasoning assistant specialized in Knowledge-Based Visual Question Answering (KB-VQA). Your task is to evaluate whether a given text passage provides useful and relevant

information for answering a question about an image.

You will be given:

- **Image:** a visual scene containing entities, actions, and context.

- **Question:** a natural-language question that refers to the image.

- **Text Passage:** an external knowledge snippet retrieved from a database.

You must analyze the semantic alignment between the text, the image, and the question. Follow these steps carefully before giving your final decision:

1. Understand the visual scene: Identify the key objects, people, actions, and context visible in the image.

2. Interpret the question: Determine what information the question seeks (e.g., factual, reasoning, counting, attribute-based).

3. Analyze the text passage: Extract the main claims, facts, and entities mentioned in the text.

Compare for relevance: Assess whether the information in the text:

- Contains at least one sentence that supports answering the question about the image, OR

- Provides background knowledge needed to interpret or reason about the image-question pair.

Important:

- If even a single sentence in the passage is relevant or useful, consider the entire passage as relevant and answer "Yes".

- If no part of the passage contributes meaningfully to answering the question, answer "No".

Output only one word:

"Yes" -> if the text provides relevant or useful information for answering the question.

"No" -> if the text is irrelevant or unhelpful.

Critic User Prompt

Here is the question on the image above:

{Question}

Here is the text passage to analyze:

{Passage}

Does the text passage contain at least one sentence that may have some information useful to answer the user question?

"Yes"/"No" answer:

Generator Prompt Design. This prompt defines the instruction for the generator model, which receives the image, question, and textual context. During training, the model is provided with only the single passage associated with the current example from the ReflectiVA dataset [13], whereas at inference it takes the subset of passages selected by the critic model. The generator is prompted to synthesize a fi-

nal answer grounded in both visual and textual evidence. At inference, providing critic-filtered passages as input encourages concise, evidence-based reasoning, reduces the impact of irrelevant or noisy information, and improves factual grounding in multimodal responses. The generator system prompt is adapted from Dr. GRPO [33]. Notably, when the critic model filters all retrieved passages (*i.e.*, when $j = 0$) the user prompt is changed and only the question with the image is fed to the generator.

Generator System Prompt

A conversation between User and Assistant. The user asks a question, and the Assistant solves it. The assistant first thinks about the reasoning process and then provides the user with the answer. The reasoning process and answer are enclosed within <think> </think> and <answer> </answer> tags, respectively, *i.e.*, <think>reasoning process here</think><answer>short answer here</answer>.

Generator User Prompt

{Question}

The following paragraphs may contain useful information to help answer the question correctly:

<paragraph>{Passage₁}</paragraph>

...

<paragraph>{Passage_j}</paragraph>.

Reasoning-Trace Prompt Design. To extract reasoning traces used during the SFT training stage (cf. Sec. 3.3 of the main paper), we employ a structured prompting strategy that elicits explicit, step-by-step inference from a teacher MLLM. The system prompt instructs the model to analyze the image, the question, and one retrieved passage, then produce a hidden reasoning trace (<think>...</think>) that (i) grounds its steps in visual evidence (*e.g.*, objects, attributes, spatial relations), (ii) evaluates the content of the passage and explicitly states whether it is relevant or irrelevant, and (iii) connects visual and textual cues via a logical chain. The user prompt supplies the question, the retrieved passage with its relevance tag, and the correct answer; the model must output the trace plus the final answer in a strict schema. Collected reasoning traces are used to initialize the generator with explicit reasoning trajectories that link the image, retrieved evidence, and the question, thereby strengthening its reasoning capabilities before the RL stage.

Reasoning Traces Generation System Prompt

You are a multimodal reasoning assistant. Your goal is to analyze the image, the question, and the retrieved passage, and

then produce a hidden reasoning trace that logically leads to the given answer. The reasoning must be step-by-step, plausible, and based on both the visual evidence and the retrieved text passage. You MUST explicitly state, within your reasoning trace, whether the passage is relevant or not according to the information provided (i.e., if it is labeled as "irrelevant", your reasoning must clearly and logically explain why it is not relevant, and if it is labeled as "relevant", your reasoning must logically support its relevance).

Do not mention, restate, or hint at the correct answer in the reasoning trace. Your reasoning trace should include:

- Description of relevant visual evidence (objects, spatial relations, attributes).
- Analysis of the retrieved passage (what it states, whether it supports or contradicts the image/question, and its relevance).
- Logical deduction that connects the visual and textual evidence to reach a conclusion.

Here you have two good examples of reasoning traces:

EXAMPLE 1 with Relevant passage:
 Question: {Example 1 Question}
 Retrieved **Relevant** Passage:
 {Example 1 Relevant Passage}
 Correct answer: {Example 1 Answer}
 Output:
 <think> {Example 1 Reasoning Trace} </think>
 <answer> {Example 1 Answer} </answer>

EXAMPLE 2 with Irrelevant passage:
 Question: {Example 2 Question}
 Retrieved **Irrelevant** Passage:
 {Example 2 Irrelevant Passage}
 Correct answer: {Example 2 Answer}
 Output:
 <think> {Example 2 Reasoning Trace} </think>
 <answer> {Example 2 Answer} </answer>

Output your reasoning and the correct answer using the exact format below:
 <think> [your reasoning trace here] </think>
 <answer> [the provided answer] </answer>

Reasoning Traces Generation User Prompt

Question: {Question}
 Retrieved **{Relevant}** Passage: {Passage}
 Correct Answer: {Answer}
 Please produce a reasoning trace that could logically lead to this answer, based on both the image and the retrieved passage if relevant.
 Do not mention or hint at the answer explicitly in your reasoning.
 Concentrate on providing a coherent explanation that supports the indicated relevance or irrelevance of the passage

in the reasoning trace, integrating both textual and visual evidence.
 Make sure to insert the correct answer between the answer tags.

E. Additional Qualitative Results

Reasoning Traces. To further interpret the behavior of our model, we visualize qualitative examples of the reasoning traces generated by ReAG in Fig. 9. The zero-shot baseline produces partial reasoning but lacks a consistent structure and does not adhere to the output format defined by the evaluation datasets. In contrast, ReflectiVA follows the correct answer format but fails to generate explicit reasoning traces, limiting interpretability. In contrast, ReAG generates coherent, well-structured traces that reveal the step-by-step logic behind its predictions. These examples highlight the ability of the proposed solution to integrate visual and textual cues, assess the relevance of retrieved passages, and maintain consistent reasoning even under noisy or irrelevant evidence, where baselines often over-rely on passages or hallucinate unsupported details.

KB-VQA Qualitative Results. Fig. 10 presents additional qualitative examples from the InfoSeek and Encyclopedic-VQA benchmarks, comparing the responses of ReAG, ReflectiVA [13]. As shown, ReAG produces answers that remain aligned with both the visual content and the retrieved evidence, benefiting from its critic-guided filtering and structured reasoning. These examples further highlight the robustness of ReAG in handling complex, knowledge-driven VQA scenarios.

F. Limitations and Impact

While ReAG demonstrates strong performance across standard benchmarks, it still faces some limitations. First, the generator produces a detailed reasoning trace, which improves the explainability of the final answer but may also increase latency, as more tokens must be generated before producing the answer. Second, the quality of ReAG depends on the reliability of the retrieved evidence. Although the critic effectively filters irrelevant passages, retrieval failures or missing knowledge can still lead to incomplete or incorrect reasoning. Moreover, the model may occasionally over-structure its explanations, producing reasoning that is correct in format but not perfectly aligned with human logic.

Despite these limitations, the explicit separation of evidence filtering from reasoning and answer generation enables ReAG to achieve strong performance while promoting greater transparency and explainability, potentially inspiring future research on modular and trustworthy multimodal reasoning frameworks.



Q: Controversy surrounds the origin of what in the city from which one would drive to reach this nature reserve?

Qwen2.5-VL-7B (ZS) [4]:

Based on the information provided, the city from which one would drive to reach the nature reserve is Port Augusta West. The text mentions that Port Augusta West is a suburb of Port Augusta, South Australia, and lists it as the starting point for driving to the nature reserve ✗

ReflectiVA [13]:

The Town ✗

ReAG (Ours):

<think> The image shows a nature reserve with a tree silhouette against a sunset sky, which matches the setting described in the retrieved passages. The question specifically asks for the origin of something in a city from which one would drive to reach this particular nature reserve. None of the retrieved passages provide any relevant information whatsoever regarding the question asked. Therefore, the final logical is that these passages must be disregarded insofar as the question is concerned. </think> <answer> Willi Stoph ✓ </answer>



Q: What genus are this plant and schefflera actinophylla in?

Qwen2.5-VL-7B (ZS) [4]:

Not enough information is provided to determine the genus in question ✗

ReflectiVA [13]:

Schefflera ✗

ReAG (Ours):

<think> The image shows a plant that matches the description of Schefflera actinophylla in the retrieved passage. The question specifically asks for the genus of this particular plant in the image and Schefflera actinophylla. Both plants are identified as belonging to the same genus in the retrieved passage. The question directly corresponds to the information presented in the retrieved passage without any need to refer to the other retrieved passages. Therefore, the final verdict is that the final question must ignore these retrieved passages entirely. </think> <answer> Heptapleurum ✓ </answer>



Q: Who designed this dock?

Qwen2.5-VL-7B (ZS) [4]:

Not enough information is provided to determine who designed the dock ✗

ReflectiVA [13]:

Jesse Hartley ✗

ReAG (Ours):

<think> The image shows the Albert Dock in Liverpool, which matches the question asked. The question specifically asks for the designer of this particular dock. The retrieved passage directly states that the dock in question, the Albert Dock, was designed by Jesse Hartley. Since the question specifically matches the dock in the image with the one mentioned in the passage, the retrieved passage directly answers the question without any need to refer to the question itself or the retrieved passage. </think> <answer> Royal Albert Dock ✓ </answer>

Figure 9. Qualitative examples from Encyclopedic-VQA where we show the ability of ReAG to reason about the question-image pair and the retrieved passages. For reference, we also report the generation from ReflectiVA [13] and the corresponding zero-shot model.

Q: What did the zuni people use the bark of this plant to dye black?



Qwen2.5-VL-7B (ZS) [4]:
The plant you are referring to is the Desert Paintbrush [...] ✗
ReflectiVA [13]:
Dyes ✗
ReAG (Ours):
Deerskin ✓

Q: What part of a tree is this plant often found under?



Qwen2.5-VL-7B (ZS) [4]:
Polystichum is often found under trees in disturbed habitats such as road cuts, talus slopes, and stream banks ✗
ReflectiVA [13]:
Trunk ✗
ReAG (Ours):
Canopy ✓

Q: In what german city is this landmark the most famous?



Qwen2.5-VL-7B (ZS) [4]:
The landmark in the picture is the Brunswick Lion ✗
ReflectiVA [13]:
Braunschweig ✗
ReAG (Ours):
Brunswick ✓

Q: What type of bird is the Contopus cooperi?



Qwen2.5-VL-7B (ZS) [4]:
Contopus cooperi (Olive-sided Flycatcher) ✗
ReflectiVA [13]:
Passerine ✗
ReAG (Ours):
Flycatcher ✓

Q: Who dies without a male heir in the netherlands the year this plant was introduced to canada?



Qwen2.5-VL-7B (ZS) [4]:
Queen Beatrix ✗
ReflectiVA [13]:
The dutch royal family ✗
ReAG (Ours):
King William III ✓

Q: In what month does it not snow in this lake?



Qwen2.5-VL-7B (ZS) [4]:
The lake is not affected by snowfall, [...] ✗
ReflectiVA [13]:
July ✗
ReAG (Ours):
July and August ✓

Q: Who designed this palace?



Qwen2.5-VL-7B (ZS) [4]:
Not enough information is provided to determine [...] ✗
ReflectiVA [13]:
Peter Frederick Robinson ✗
ReAG (Ours):
John Nash ✓

Q: What type of hawksbeard is this plant commonly known as?



Qwen2.5-VL-7B (ZS) [4]:
The plant in the image is likely rough hawksbeard ✗
ReflectiVA [13]:
Mouse-ear ✗
ReAG (Ours):
Narrowleaf ✓

Q: What monument was placed under state control in 1882?



Qwen2.5-VL-7B (ZS) [4]:
Newgrange was placed under state control in 1882 ✗
ReflectiVA [13]:
Newgrange ✗
ReAG (Ours):
Newgrange and Knowth and Dowth ✓

Q: What is the habitat of this animal?



Qwen2.5-VL-7B (ZS) [4]:
In the upper intertidal zone on open coast [...] ✗
ReflectiVA [13]:
Salt marshes and quiet bays ✗
ReAG (Ours):
Intertidal zone on the open coast ✓

Q: What month of meteorological winter in the northern hemisphere did hungary purchase this building?



Qwen2.5-VL-7B (ZS) [4]:
According to the information provided, the U.S. Civil Service Building was built [...] in 1911 ✗
ReflectiVA [13]:
December ✗
ReAG (Ours):
Third ✓

Q: Who was the duke of milan when Edward Cowper began leading the monks at this abbey?



Qwen2.5-VL-7B (ZS) [4]:
The Duke of Brittany, William de Lis, was the owner of the patronage of Roche Abbey ✗
ReflectiVA [13]:
Charles I ✗
ReAG (Ours):
Maximilian Sforza ✓

Q: What is the religious order of this monastery?



Qwen2.5-VL-7B (ZS) [4]:
Bremian Knighthood ✗
ReflectiVA [13]:
Premonstratensian ✗
ReAG (Ours):
Order of the Holy Cross ✓

Q: In which country or region does this animal live?



Qwen2.5-VL-7B (ZS) [4]:
It is native to southern Canada, the contiguous USA, and Mexico ✗
ReflectiVA [13]:
Eurasia ✗
ReAG (Ours):
North America ✓

Q: What did Alexander Wilson do to this bird in 1811?




Qwen2.5-VL-7B (ZS) [4]:
Alexander Wilson did not harm the bird in any way [...] ✗
ReflectiVA [13]:
I don't see any information about a bird species ✗
ReAG (Ours):
Discovered ✓

Q: What type of habitat does this plant prefer?



Qwen2.5-VL-7B (ZS) [4]:
It cannot be determined from the given information alone, but it is likely to [...] ✗
ReflectiVA [13]:
Dry ✗
ReAG (Ours):
Chaparral ✓

Q: Which teams played against each other in the first match ever held in this stadium?



Qwen2.5-VL-7B (ZS) [4]:
The first match ever held in FirstEnergy Stadium, [...] ✗
ReflectiVA [13]:
Minnesota vikings ✗
ReAG (Ours):
Browns and Minnesota Vikings ✓

Q: What wars ended the year the fichtelberg cable car was built on this mountain?




Qwen2.5-VL-7B (ZS) [4]:
Not enough information is provided to determine [...] ✗
ReflectiVA [13]:
World war I ✗
ReAG (Ours):
The American Indian Wars ✓

Q: For what football cup final is this stadium the venue?



Qwen2.5-VL-7B (ZS) [4]:
The football cup final for which this stadium is the venue is the Taça de Portugal ✗
ReflectiVA [13]:
Uefa women's championship ✗
ReAG (Ours):
Portuguese ✓

Q: What is the surface area of this lake?



Qwen2.5-VL-7B (ZS) [4]:
Ełk Lake is a fresh water lake in the Masurian Lake District of Poland's Warmia-Mazury Province ✗
ReflectiVA [13]:
3.82 ✗
ReAG (Ours):
3.82 square kilometres ✓

Q: How many eggs does this bird typically lay?



Qwen2.5-VL-7B (ZS) [4]:
Three eggs ✗
ReflectiVA [13]:
Three ✗
ReAG (Ours):
Three to four ✓

Figure 10. Qualitative results on InfoSeek and Encyclopedic-VQA image-question pairs comparing ReAG, ReflectiVA [13], and the corresponding zero-shot model.