

CoRe-MMRAG: Cross-Source Knowledge Reconciliation for Multimodal RAG

Yang Tian[‡], Fan Liu[†], Jingyuan Zhang[¶], Victoria W.[¶], Yupeng Hu^{‡*}, Liqiang Nie^{§*}

[‡] School of Software, Shandong University

[†] National University of Singapore, [¶] Independent Researcher

[§] Harbin Institute of Technology, Shenzhen

{tianyangchn,liufancs,nieliqiang}@gmail.com, huyupeng@sdu.edu.cn

Abstract

Multimodal Retrieval-Augmented Generation (MMRAG) has been introduced to enhance Multimodal Large Language Models by incorporating externally retrieved multimodal knowledge, but it introduces two challenges: Parametric-Retrieved Knowledge Inconsistency (PRKI), where discrepancies between parametric and retrieved knowledge create uncertainty in determining reliability, and Visual-Textual Knowledge Inconsistency (VTKI), where misalignment between visual and textual sources disrupts entity representation. To address these challenges, we propose **Cross-source knowledge Reconciliation for MultiModal RAG (CoRe-MMRAG)**, a novel end-to-end framework that effectively reconciles inconsistencies across knowledge sources. CoRe-MMRAG follows a four-stage pipeline: it first generates an internal response from parametric knowledge, then selects the most relevant multimodal evidence via joint similarity assessment, generates an external response, and finally integrates both to produce a reliable answer. Additionally, a specialized training paradigm enhances knowledge source discrimination, multimodal integration, and unified answer generation. Experiments on KB-VQA benchmarks show that CoRe-MMRAG achieves substantial improvements over baseline methods, achieving 5.6% and 9.3% performance gains on InfoSeek and Encyclopedic-VQA, respectively. We release code and data at <https://github.com/TyangJN/CoRe-MMRAG>.

1 Introduction

Recent advances in Multimodal Large Language Models (MLLMs) (Alayrac et al., 2022; Li et al., 2023a; Liu et al., 2023, 2024b; Achiam et al., 2023; Reid et al., 2024) have significantly improved multimodal reasoning and generation tasks by leveraging joint vision-language representations. How-

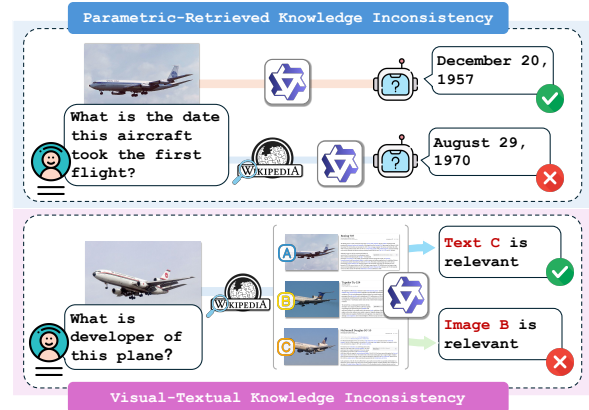


Figure 1: Two types of knowledge inconsistency in MMRAG: (1) Parametric-Retrieved Knowledge Inconsistency, where parametric and retrieved external knowledge generate conflicting answers to the same query. (2) Visual-Textual Knowledge Inconsistency, where misalignments between visual and textual sources disrupt entity representation.

ever, these models inherently suffer from hallucination (Bai et al., 2024) and knowledge limitations (Caffagni et al., 2024), as their parametric knowledge is frozen after pretraining and cannot dynamically adapt to external information.

Multimodal Retrieval-Augmented Generation (MMRAG) has emerged as a promising approach to enhance MLLMs by incorporating retrieved textual and visual knowledge during inference (Yan and Xie, 2024; Qi et al., 2024; Zhang et al., 2024b). By accessing external information, MMRAG helps mitigate knowledge gaps and improves factual grounding. However, integrating retrieved knowledge into MLLMs presents two key challenges. First, **Parametric-Retrieved Knowledge Inconsistency (PRKI)**. Since MLLMs rely on frozen pretraining knowledge, retrieved text and images may contradict, extend, or refine this information. Moreover, retrieved content can be incomplete, noisy, or misleading, introducing biases or factual errors. Without effective reconciliation, the model may struggle to balance reliability between internal para-

* Corresponding author

metric and external retrieved knowledge, leading to incorrect responses. As shown in Figure 1, for the aircraft’s first flight date, introducing noisy retrieval information (August 29, 1970) overrides the model’s reliable parametric knowledge (December 20, 1957). Second, **Visual-Textual Knowledge Inconsistency (VTKI)**. Since each modality captures different aspects of entity representation (Li et al., 2023a), retrieved images and textual documents often provide non-overlapping and misaligned information. For instance, an image may visually relate to the query while its paired text describes a different aspect or interpretation. These inconsistencies disrupt knowledge integration, making it difficult for the model to determine which information to prioritize.

To address the above-mentioned challenges, we propose **Cross-source knowledge Reconciliation for MultiModal RAG (CoRe-MMRAG)**, a novel end-to-end framework designed to mitigate the inconsistencies between different knowledge sources. CoRe-MMRAG follows a four-stage pipeline: it first generates an initial response using only the model’s internal knowledge, then performs joint similarity assessment to select the most relevant multimodal evidence, followed by generating a response grounded in the retrieved content, and finally integrates both sources to produce a coherent and reliable answer. This structured process enables the model to effectively reconcile diverse knowledge inputs and leverage complementary information from multiple modalities.

To further enhance knowledge reconciliation, CoRe-MMRAG incorporates a specialized training paradigm with three key objectives: knowledge source selection that learns to identify the most reliable information source between internal parametric and external retrieved knowledge, multimodal selection that optimizes the joint understanding of visual-textual pairs, and unified answer generation that ensures consistent and accurate responses. Through this comprehensive training strategy, CoRe-MMRAG develops robust capabilities in handling knowledge inconsistencies between different sources.

We conduct comprehensive experiments on two knowledge-based VQA benchmarks (Mensink et al., 2023; Chen et al., 2023) to evaluate our approach. Using Qwen2-VL-7B (Wang et al., 2024) as the base MLLM, CoRe-MMRAG demonstrates substantial improvements in both zero-shot and fine-tuned setting. Specifically, our method

achieves performance gains of 5.6% over the baseline on the InfoSeek, while also surpassing the baseline on the Encyclopedic-VQA benchmark by 9.3%. The contributions of this work are summarized as follows:

- We identify and formalize two fundamental challenges in multimodal RAG: Parametric-Retrieved Knowledge Inconsistency and Visual-Textual Knowledge Inconsistency. To address these issues, we propose CoRe-MMRAG, a novel end-to-end framework that effectively reconciles inconsistencies across different knowledge sources.
- We design a specialized training paradigm with three targeted objectives that enhance MLLMs in knowledge source discrimination, multimodal integration, and unified answer generation, enabling effective knowledge inconsistency mitigation.
- Extensive experiments demonstrate that CoRe-MMRAG achieves significant performance gains over previous SOTAs on multiple Knowledge-based VQA benchmarks.

2 Related Work

Multimodal RAG. Recent advancements in MLLMs, such as LLaVA-family (Liu et al., 2023; Li et al., 2024), Qwen2-VL (Wang et al., 2024), MiniCPM-V (Yao et al., 2024), and Intern-VL (Chen et al., 2024), have demonstrated remarkable performance across various multimodal tasks. However, these models inherently suffer from hallucination issues (Li et al., 2023b; Caffagni et al., 2024). One effective approach to mitigating hallucination is the incorporation of multimodal data (Liu et al., 2024a), which provides complementary knowledge beyond text-based sources. By integrating multimodal information, models can ground their responses in more diverse and concrete evidence, reducing the likelihood of hallucination. Another approach is inspired by RAG (Kandpal et al., 2023; Li et al., 2025; Gao et al., 2023; Asai et al., 2024), several multimodal RAG frameworks (Lin et al., 2023; Hu et al., 2025; Adjali et al., 2024) have been proposed, including Wiki-LLaVA (Caffagni et al., 2024), EchoSight (Yan and Xie, 2024), RoRA-VLM (Qi et al., 2024) and LLaVA-mR²AG (Zhang et al., 2024b). These methods typically follow a three-stage pipeline of retrieval,

reranking, and generation. However, during reranking, they primarily rely on text-based similarity measures between retrieved passages and questions, which may lead to incorrect passage selection due to the inherent cross-modal discrepancy, ultimately affecting the final answer quality.

Multimodal Understanding. Multimodal understanding (Zhang et al., 2024a; Yin et al., 2023; Wang et al., 2025; Liu et al., 2018) aims to integrate and interpret information across multiple modalities, such as vision and language, to enhance reasoning and decision-making. A key challenge in this field is incorporating external knowledge to support tasks that require information beyond what is directly present in one modality. One major research direction is knowledge-based Visual Question Answering (VQA), where models answer questions requiring external factual knowledge. Early benchmarks like OK-VQA (Marino et al., 2019) and A-OKVQA (Schwenk et al., 2022) introduced commonsense and general knowledge reasoning, while ViQuAE (Lerner et al., 2022) expanded to entity-specific queries. More recent datasets, such as InfoSeek (Chen et al., 2023) and Encyclopedic VQA (Mensink et al., 2023), enforce both visual grounding and knowledge retrieval, exposing the limitations of existing models, including LLM-based approaches, in handling multimodal knowledge integration. Beyond VQA, multimodal understanding extends to tasks such as image-text retrieval, captioning, and reasoning, where aligning visual and textual representations is critical.

3 Methodology

In this section, we first formalize the problem of Multimodal Retrieval-Augmented Generation (MMRAG) and introduce two types of inconsistency (§3.1) that arise when applying MMRAG to KB-VQA. We then present our framework (§3.2) and detail our training approach (§3.3), both of which are designed to enhance the capacity of MLLMs to resolve knowledge inconsistency.

3.1 Problem Formulation

To assess the effectiveness of our proposed framework, we conduct evaluations on Knowledge-based Visual Question Answering (KB-VQA) tasks. Given an input image-question pair $Q = (Q^v, Q^t)$ from the question set \mathbf{Q} with support from an external knowledge base \mathbf{K} , where each knowledge entry \mathbf{K}_i comprises a visual component V_i

and the associated textual article T_i . A typical MM-RAG pipeline addresses KB-VQA through three stages: retrieval, reranking, and generation. In the retrieval stage, a frozen CLIP (Radford et al., 2021) encodes both the query image Q^v and knowledge base images $\{V_i\}_{i=1}^{|\mathbf{K}|}$ into a shared embedding space, where the relevance is measured via cosine similarity: $\text{sim}(Q^v, V_i)$. The top- k entries $\{(V_i, T_i)\}_{i=1}^k$ are retrieved based on these similarity scores. Then, the retrieved entries are reranked by the multimodal model \mathcal{M} based on the semantic relevance between Q and the retrieved articles $\{T_i\}_{i=1}^k$. Finally, \mathcal{M} processes Q along with the most relevant entry to generate the final answer.

During this process, we identify two types of issues. The first problem arises from the inconsistency between the model’s parametric knowledge and the retrieved external knowledge, referred to as **Parametric-Retrieved Knowledge Inconsistency (PRKI)**. Formally, let $\mathcal{M}(Q)$ denote the output based on the model’s parametric knowledge, and let $\mathcal{M}(Q, \mathbf{P})$ represent the response when augmented with the retrieved knowledge set $\mathbf{P} = \{(V_i, T_i)\}_{i=1}^k$. The PRKI occurs when:

$$\mathcal{M}(Q) \neq \mathcal{M}(Q, \mathbf{P}). \quad (1)$$

The second issue is **Visual-Textual Knowledge Inconsistency (VTKI)**, which arises when textual and visual modalities of the retrieved entries $\{(T_i, V_i)\}_{i=1}^k$ yield inconsistent relevance rankings. Formally, the VTKI manifests when:

$$\text{argmax } r^{\mathcal{M}}(Q, \mathbf{T}) \neq \text{argmax } r^{\mathcal{M}}(Q, \mathbf{V}), \quad (2)$$

where $\text{argmax } r^{\mathcal{M}}(\cdot, \cdot)$ represents model \mathcal{M} ’s prediction of the most relevant entry ID in each modality. Both PRKI and VTKI can significantly impact model performance. PRKI occurs when noisy external knowledge overrides reliable parametric knowledge, resulting in erroneous outputs. Meanwhile, VTKI becomes especially problematic during the reranking stage. Due to the inconsistency between the textual and visual knowledge, relying on unimodal knowledge increases the risk of introducing irrelevant information to \mathcal{M} , which can propagate errors from the reranking stage to answer generation, potentially leading to PRKI and a reduction in model performance.

3.2 CoRe-MMRAG Framework

To mitigate the PRKI and VTKI problems described in §3.1, we propose CoRe-MMRAG, a

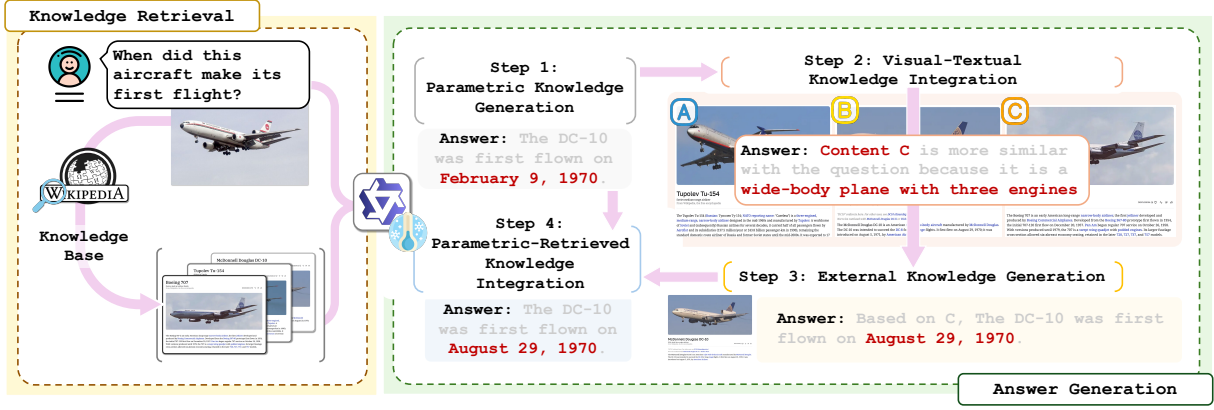


Figure 2: Overview of the CoRe-MMRAG Framework. CoRe-MMRAG processes a multimodal query and retrieved knowledge in four stages: (1) generating an initial response based solely on parametric knowledge; (2) jointly integrating visual and textual information to identify the most relevant retrieved content; (3) generating an external response based on the selected content from step (2); and (4) reconciling discrepancies between parametric and retrieved knowledge to produce the final answer.

framework that effectively reconciles inconsistencies across knowledge sources. As shown in Figure 2, given a query Q and its retrieved knowledge entries \mathbf{P} , the model \mathcal{M} is prompted to generate responses across four distinct stages in an end-to-end manner. Below, we detail each stage of the framework.

Step 1: Parametric Knowledge Generation.

Although external knowledge \mathbf{P} is available in the input, we first prompt the model to generate y^{int} based solely on its parametric knowledge:

$$y^{\text{int}} = \mathcal{M}(Q). \quad (3)$$

This generation establishes a reference point for detecting potential conflicts with retrieved knowledge in Step 4.

Step 2: Visual-Textual Knowledge Integration. Following the parametric knowledge generation, the model evaluates the relevance between query Q and knowledge entries in \mathbf{P} . Considering the potential VTKI manifested as:

$$\begin{cases} I^v = \underset{i \in \{1, \dots, k\}}{\operatorname{argmax}} r^{\mathcal{M}}(Q, \{V_i\}_{i=1}^k), \\ I^t = \underset{j \in \{1, \dots, k\}}{\operatorname{argmax}} r^{\mathcal{M}}(Q, \{T_j\}_{j=1}^k), \\ I^v \neq I^t. \end{cases} \quad (4)$$

We propose a joint similarity assessment that utilizes the complementary nature of visual and textual modalities:

$$I^{tv} = \underset{i \in \{1, \dots, k\}}{\operatorname{argmax}} r^{\mathcal{M}}(Q, \{V_i, T_i\}_{i=1}^k), \quad (5)$$

where I^{tv} denotes the most relevant entry ID based on multimodal knowledge. This unified ranking ap-

proach jointly leverages abstract semantic concepts from textual descriptions and detailed visual characteristics, eliminating the bias introduced by separate unimodal evaluations, which enables a more robust relevance assessment and resolves VTKI.

Step 3: External Knowledge Generation. After obtaining the most relevant knowledge entry $\mathbf{P}_{I^{tv}}$, the model is prompted to generate a response based on the retrieved external knowledge:

$$y^{\text{ext}} = \mathcal{M}(Q, (V_{I^{tv}}, T_{I^{tv}})), \quad (6)$$

where y^{ext} denotes the model’s response that explicitly considers the retrieved visual-textual knowledge pair $(V_{I^{tv}}, T_{I^{tv}})$.

Step 4: Parametric-Retrieved Knowledge Integration. Given responses from parametric knowledge y^{int} and retrieved knowledge y^{ext} , the parametric-retrieved knowledge inconsistencies may arise. The model is prompted to resolve these inconsistencies and generate the final response:

$$y^* = \mathcal{M}(Q, y^{\text{int}}, y^{\text{ext}}, (V_{I^{tv}}, T_{I^{tv}})), \quad (7)$$

where y^* represents the final response, which is determined by comparing the credibility of parametric knowledge and retrieved external knowledge. This process enables the model to leverage both knowledge sources while ensuring the reliable generation of the final answer.

3.3 Inconsistency-Aware Multimodal Training

Inspired by the self-training mechanism in STaR (Zelikman et al., 2022), we propose a fine-tuning paradigm that leverages the model’s self-generated outputs under different knowledge conditions. The

model learns to resolve PRKI and mitigate VTKI through three specialized training objectives, thus improving its ability to generate accurate answers based on retrieved knowledge.

Parametric-Retrieved Knowledge Selection.

We begin by generating answers based solely on the model’s parametric knowledge, $\hat{y}^{\text{int}} = \mathcal{M}(Q)$, and re-evaluate the same question with retrieved external knowledge, obtaining $\hat{y}^{\text{ext}} = \mathcal{M}(Q, \mathbf{P})$. After generating both outputs, we filter questions where the model produces correct answers exclusively using either internal or external knowledge, forming fine-tuning datasets \mathbf{D}^{int} and \mathbf{D}^{ext} :

$$\begin{cases} \hat{y}^{\text{int}} \neq \hat{y}^{\text{ext}}, \\ \mathbf{D}^{\text{int}} = \{(Q_j, \mathbf{P}_j) | \hat{y}_j^{\text{int}} = \hat{y}_j^{\text{gt}}\}_{j=1}^{|\mathbf{Q}|}, \\ \mathbf{D}^{\text{ext}} = \{(Q_i, \mathbf{P}_i) | \hat{y}_i^{\text{ext}} = \hat{y}_i^{\text{gt}}\}_{i=1}^{|\mathbf{Q}|}. \end{cases} \quad (8)$$

Then, we fine-tune the model \mathcal{M} on \mathbf{D}^{int} and \mathbf{D}^{ext} , the training is guided by loss function $\mathcal{L}_{\text{PRKI}}$:

$$\mathcal{L}_{\text{PRKI}} = - \left[\sum_{(Q_j, \mathbf{P}_j) \sim \mathbf{D}^{\text{int}}} \log \mathcal{M}(\hat{y}^{\text{int}} | Q_j, \mathbf{P}_j) + \sum_{(Q_i, \mathbf{P}_i) \sim \mathbf{D}^{\text{ext}}} \log \mathcal{M}(\hat{y}^{\text{ext}} | Q_i, \mathbf{P}_i) \right], \quad (9)$$

which encourages the model to prioritize the knowledge source that leads to correct answer generation, ensuring robustness in handling PRKI.

Visual-Textual Knowledge Selection. The model computes the most relevant entry IDs independently using visual knowledge $\hat{I}^v = \arg\max r^{\mathcal{M}}(Q, \mathbf{V})$ and textual knowledge $\hat{I}^t = \arg\max r^{\mathcal{M}}(Q, \mathbf{T})$. The training datasets \mathbf{D}^v and \mathbf{D}^t are constructed by selecting samples as follows:

$$\begin{cases} \hat{I}^v \neq \hat{I}^t, \\ \mathbf{D}^v = \{(Q_m, \mathbf{P}_m) | \hat{I}_m^v = I_m^{\text{gt}}\}_{m=1}^{|\mathbf{Q}|}, \\ \mathbf{D}^t = \{(Q_n, \mathbf{P}_n) | \hat{I}_n^t = I_n^{\text{gt}}\}_{n=1}^{|\mathbf{Q}|}. \end{cases} \quad (10)$$

Here, I^{gt} denotes the ground-truth index for the most relevant entry. The model is then fine-tuned on \mathbf{D}^v and \mathbf{D}^t using:

$$\mathcal{L}_{\text{VTKI}} = - \left[\sum_{(Q_m, \mathbf{P}_m) \sim \mathbf{D}^v} \log \arg \max_{j \in \{1, \dots, k\}} r^{\mathcal{M}}(\hat{I}^v | Q_m, \mathbf{P}_m^j) + \sum_{(Q_n, \mathbf{P}_n) \sim \mathbf{D}^t} \log \arg \max_{i \in \{1, \dots, k\}} r^{\mathcal{M}}(\hat{I}^t | Q_n, \mathbf{P}_n^i) \right], \quad (11)$$

Dataset	Train	Val	Test	KB
Enc-VQA	1M	13.6K	5.8K	2M Wiki
InfoSeek	934K	73K	348K	100K Wiki

Table 1: Statistics of datasets, including sample splits and their corresponding knowledge base sizes.

where the loss function $\mathcal{L}^{\text{VTKI}}$ enables the model to evaluate the reliability of visual and textual modalities and prioritize the more confident one, thereby mitigating VTKI induced by unimodal bias.

Unified Answer Generation. We continue training the model on \mathbf{D}^v and \mathbf{D}^t , applying Supervised Fine-Tuning (SFT) with the loss function:

$$\mathcal{L}_{\text{SFT}} = - \left[\sum_{(Q_m, \mathbf{P}_m) \sim \mathbf{D}^v} \log \mathcal{M}(\hat{y}^{\text{gt}} | Q_m, \mathbf{P}_m^{I^v}) + \sum_{(Q_n, \mathbf{P}_n) \sim \mathbf{D}^t} \log \mathcal{M}(\hat{y}^{\text{gt}} | Q_n, \mathbf{P}_n^{I^t}) \right], \quad (12)$$

where \mathcal{L}_{SFT} is used to fine-tune the model, encouraging it to generate accurate answers based on the ground-truth external knowledge.

4 Experiments

4.1 Datasets

We evaluate our proposed CoRe-MMRAG on two large-scale knowledge-based VQA benchmarks: Encyclopedic VQA (Mensink et al., 2023) and InfoSeek (Chen et al., 2023). Both datasets contain diverse visual-textual queries requiring fine-grained entity knowledge, with explicit knowledge bases to ensure answer verifiability. Encyclopedic VQA (Enc-VQA) contains 221K $(Q^t, \hat{y}^{\text{gt}})$ unique question-answer pairs distributed across 16.7K fine-grained entities, where each question-answer pair is associated with up to five diverse instance images, resulting in 1M image-question-answer $(Q^v, Q^t, \hat{y}^{\text{gt}})$ triplets, while InfoSeek contains 1.3M $(Q^v, Q^t, \hat{y}^{\text{gt}})$ triplets corresponding to approximately 11K distinct entities. Detailed statistics for both data sets, including sample splits and knowledge base sizes, are shown in Table 1. Following standard practice in EchoSight (Yan and Xie, 2024), we evaluate on Enc-VQA’s test set after excluding two-hop questions, resulting in 4.7K test triplets. For InfoSeek, we report results on the validation split, which contains unseen entities (Unseen-E) and novel questions (Unseen-Q).

4.2 Metrics

Metrics for Retrieval. We adopt $\text{Recall}@k$ as the standard metric to evaluate the retrieval performance (Liu et al., 2021). This metric examines whether the ground-truth article appears within the top- k retrieved results. Following EchoSight, the evaluation criterion considers an article correct only when its URL exactly matches the ground-truth URL, ensuring precise assessment of retrieval accuracy.

Metrics for Answer Generation. We employ dataset-specific metrics following standard practices in knowledge-based VQA. For Enc-VQA, we use BEM (Zhang et al., 2019), while for InfoSeek (Chen et al., 2023), we adopt both VQA accuracy (Marino et al., 2019) and Relaxed Accuracy (Methani et al., 2020; Masry et al., 2022).

4.3 Implementation Details

Retriever. For external knowledge retrieval, Enc-VQA utilizes a knowledge base consisting of 2M Wikipedia articles and up to 5M associated images. In contrast, InfoSeek employs a filtered subset of 100K Wikipedia articles with approximately 350K images, following the setup in (Yan and Xie, 2024). Visual features are extracted using a frozen Eva-CLIP-8B encoder (Sun et al., 2023), where we use the pooled embeddings from the last layer to compute the cosine similarity between reference images and candidate images. We construct a visual feature index using the FAISS library for efficient similarity search and retrieve the top-5 most relevant entries as external knowledge. Retrieval performance is reported in Table 2.

Zero-shot Settings. We employ Qwen2-VL-7B (Wang et al., 2024) as our base model, leveraging its 32K token input capacity to accommodate both visual and textual knowledge. To ensure a fair comparison with existing MMRAG approaches, including Wiki-LLaVA (Caffagni et al., 2024), EchoSight (Yan and Xie, 2024), RoRA-VLM (Qi et al., 2024), and LLaVA-mR²AG (Zhang et al., 2024b), we reimplement these pipelines using Qwen2-VL-7B as a unified backbone. We consider the following configurations for zero-shot evaluation: (1) **Qwen2-VL-Param**, which relies solely on the model’s internal parametric knowledge without any external context; (2) **Qwen2-VL-Oracle**, which takes the ground-truth wiki entry as input and serves as an upper-bound; (3) **Qwen2-VL-1-Stage**, which replicates Wiki-LLaVA’s one-stage pipeline by encod-

Datasets	Recall@ k (%)			
	$k=1$	$k=2$	$k=5$	$k=10$
InfoSeek	45.6	54.8	67.1	73.0
Enc-VQA	13.3	16.9	31.3	41.0

Table 2: The retrieval performance of Eva-CLIP-8B on two datasets.

ing all retrieved entries for answer generation; (4) **Qwen2-VL-2-Stage**, which follows the two-stage architecture of LLaVA-mR²AG and EchoSight by reranking retrieved entries and generating answers based on the top-ranked one; and (5) **Qwen2-VL-MMSTaR**, which incorporates all retrieved entries into a Chain-of-Thought reasoning process following the STaR framework (Zelikman et al., 2022).

Fine-tuned Setting. We fine-tune the Qwen2-VL variants with task-specific supervision to enhance both retrieval accuracy and answer generation. **Qwen2-VL-1-Stage** is trained using a standard supervised objective \mathcal{L}_{SFT} with ground-truth wiki entries as input, enhancing the model’s ability to generate correct answers from these references. **Qwen2-VL-2-Stage** is optimized with a combination of a selection objective $\mathcal{L}_{\text{PRKI}}$ and generation objective \mathcal{L}_{SFT} , enabling the model to better identify the correct entry from the top-5 retrieved candidates and generate more accurate answers based on the selected context.

Our proposed CoRe-MMRAG is trained with three objectives, including $\mathcal{L}_{\text{VTKI}}$, $\mathcal{L}_{\text{PRKI}}$, and \mathcal{L}_{SFT} , which jointly enhance parametric-retrieved knowledge selection, visual-textual knowledge selection, and final answer generation. We sample approximately 30K $(Q^v, Q^t, \hat{y}^{\text{gt}})$ triplets from the training set of each benchmark to construct the training set. The model is fine-tuned using LoRA with a rank of 8, applied to all layers. Training is conducted for 3 epochs with a learning rate of 1×10^{-4} and a batch size of 1×2 , using 8 H100 GPUs. The full training process takes approximately 10 hours.

4.4 Main Results

Table 3 presents comprehensive comparisons of our method against current SOTA approaches on Enc-VQA and InfoSeek benchmarks.

Zero-shot Setting. Qwen2-VL-Oracle, which accesses ground-truth Wikipedia entries, establishes an upper-bound performance of 51.2% on Enc-VQA and 57.9% on InfoSeek. When using retrieved entries instead of gold references, our

Model	LLM	KB	Enc-VQA	InfoSeek		
			Single-Hop	Unseen-Q	Unseen-E	All
Zero-shot Models						
LLaVA-1.5	Vicuna-7B	-	16.3	13.0	10.3	12.2
Qwen2-VL-Param	Qwen2-7B	-	12.7	23.1	21.8	22.1
Qwen2-VL-Oracle	Qwen2-7B	Wiki	51.2	57.9	57.9	57.9
Qwen2-VL-1-Stage†	Qwen2-7B	Wiki	17.9	40.8	40.9	40.9
Qwen2-VL-2-Stage†	Qwen2-7B	Wiki	17.0	34.9	35.1	35.0
Qwen2-VL-MMSTaR†	Qwen2-7B	Wiki	16.9	33.4	34.0	33.9
Ours†	Qwen2-7B	Wiki	20.1	42.3	43.3	42.9
Fine-tuned Models						
Wiki-LLaVA	Vicuna-7B	Wiki	17.7	30.1	27.8	28.9
RoRA-VLM	Vicuna-7B	Wiki+Web	20.3	27.3	25.1	26.9
EchoSight	LLaMA3-8B	Wiki	19.4	-	-	27.7
LLaVA-mR ² AG	Vicuna-7B	Wiki	55.1*	39.1	39.7	39.4
Qwen2-VL-1-Stage†	Qwen2-7B	Wiki	24.3	42.3	43.4	43.0
Qwen2-VL-2-Stage†	Qwen2-7B	Wiki	23.1	36.8	37.6	37.3
Qwen2-VL-MMSTaR†	Qwen2-7B	Wiki	20.9	34.8	35.2	35.1
Ours†	Qwen2-7B	Wiki	27.2	45.2	46.9	46.5

Table 3: Main results (%) on Enc-VQA and InfoSeek with external knowledge. † denotes our method and variants. Note: on Enc-VQA LLaVA-mR²AG* uses Google Lens for retrieval, achieving 62.5% Recall@5, while our Eva-CLIP-based retrieval achieves 31.3% Recall@5.

proposed CoRe-MMRAG achieves the best results among all methods, reaching 42.9% accuracy on InfoSeek and 20.1% on Enc-VQA, surpassing the second-best method, Qwen2-VL-1-Stage, by margins of 2.0% and 2.2%, respectively. Qwen2-VL-2-Stage yields suboptimal results due to potential information loss in unimodal reranking. Qwen2-VL-MMSTaR shows only slight improvements, indicating that current Chain-of-Thought reasoning remains limited for knowledge-intensive VQA. Nonetheless, it still outperforms the parametric-only baseline, highlighting the benefit of retrieved multimodal knowledge. Notably, in zero-shot settings, the performance gain on InfoSeek is more substantial than on Enc-VQA. This discrepancy is likely due to the larger and more complex knowledge base of Enc-VQA, which introduces greater retrieval noise and increases the difficulty of both PRKI and VTKI.

Fine-tuned Setting. Our method maintains superior performance, achieving 46.5% on InfoSeek and 27.2% on Enc-VQA. Furthermore, it exhibits the largest improvements comparing zero-shot setting, with gains of 3.6% on InfoSeek and 7.1% on Enc-VQA. Qwen2-VL-2-Stage, trained with $\mathcal{L}_{\text{VTKI}}$ and \mathcal{L}_{SFT} , shows an improvement of 2.3% improvement over its zero-shot performance on InfoSeek and 6.1% on Enc-VQA. Qwen2-VL-1-Stage, fine-tuned with \mathcal{L}_{SFT} , achieves gains of 2.1% on InfoSeek and 6.4% on Enc-VQA. Qwen2-VL-MMSTaR demonstrates limited improvement,

primarily due to the inadequate quality of its self-generated training instances.




Figure 3: Effectiveness of Our Method in Mitigating PRKI and VTKI. *Top*: Evaluation of Parametric-Retrieved Knowledge Reconciliation, comparing our proposed method with Qwen-2-VL-Param (Param) and Qwen2-VL-1-Stage (1-Stage) under zero-shot and fine-tuned settings. *Bottom*: Evaluation of Visual-Textual Knowledge Reconciliation, showing that our method improves ground-truth entry recognition through both textual and visual modalities across both settings.

Model	Recall@1			Recall@2			Recall@5			Recall@k (K>5)			All
	U-Q	U-E	Acc	U-Q	U-E	Acc	U-Q	U-E	Acc	U-Q	U-E	Acc	Acc
Param	25.0	25.3	25.2	24.4	23.9	24.1	20.7	22.3	21.8	19.3	15.7	16.4	22.1
1-Stage [†]													
Top-1	<u>59.1</u>	<u>58.6</u>	<u>58.8</u>	30.1	27.3	28.0	22.0	22.4	22.1	15.4	15.6	15.6	39.6
Top-2	58.4	57.9	58.2	<u>46.5</u>	<u>45.7</u>	<u>46.0</u>	22.6	25.9	25.1	15.5	15.9	15.8	40.4
Top-5	56.7	56.2	56.4	45.6	44.8	44.7	<u>32.3</u>	<u>34.4</u>	<u>33.8</u>	16.8	16.5	16.6	40.9
Ours													
Top-1	61.5	61.8	61.7	31.7	27.7	28.6	22.8	22.8	22.8	16.6	15.0	15.4	40.7
Top-2	61.0	61.2	61.2	48.4	49.0	48.9	23.8	27.3	26.4	16.9	15.7	16.0	42.9
Top-5	59.7	59.3	59.4	46.2	47.8	47.6	33.2	35.8	35.2	17.5	16.3	16.7	42.9

Table 4: Performance (%) with different numbers of retrieved entries. Recall@ k indicates the presence of ground-truth entry in top- k retrieved results, with accuracy reported across Unseen-Questions (U-Q) and Unseen-Entities (U-E) settings.

Knowledge Reconciliation. Figure 3 illustrates the effectiveness of our method in addressing both VTKI and PRKI under zero-shot and fine-tuned settings. In the zero-shot setting, the model identifies 46.7% of the ground truth Wikipedia entries using the textual modality and 45.6% using the visual modality, with over 40% of the entities correctly recognized by both. Our method improves this recognition rate to 50.1% demonstrating the effectiveness of our method in handling VTKI. For PRKI, our method achieves better retention of parametric knowledge compared to the 1-Stage baseline, indicating more robust reconciliation of parametric-retrieved knowledge inconsistency. In the fine-tuned setting, the model’s ability to leverage multimodal knowledge for identifying ground-truth Wikipedia entries further improves, increasing from 50.1% to 57.3%, demonstrating the effectiveness of the objective $\mathcal{L}_{\text{VTKI}}$. However, we observe that the 1-Stage baseline, trained solely with \mathcal{L}_{SFT} , tends to compromise the retention of correct parametric knowledge, resulting in a larger performance drop (23.8% to 17.6%) relative to the original parametric outputs. In contrast, our method better preserves parametric knowledge, with a modest drop (23.8% to 19.3%), indicating the effectiveness of the $\mathcal{L}_{\text{PRKI}}$ objective.

4.5 Ablation Studies

To validate the effectiveness of our approach, we conduct ablation studies on the InfoSeek validation set.

Effect of Retrieved Entry Count. Increasing the number of retrieved entries generally leads to improved overall recall accuracy. However, its impact varies across different Recall@ k groups, where Recall@ k indicates whether the ground-

truth entry is included in the top- k retrieved results. As shown in Table 4, we observe a clear relationship between the number of retrieved entries (Top- m) and the Recall@ k performance. The first case is when $m \leq k$, meaning the number of retrieved entries does not exceed the evaluation threshold. In this setting, increasing m directly expands the candidate pool, which generally leads to consistent performance improvements. For example, in the Recall@5 group, increasing m from 2 to 5 improves recall accuracy, with the 1-Stage baseline rising from 25.1% to 33.8% and our method from 26.4% to 35.2%, as more relevant entries are included in the input. In contrast, when $m > k$, the inclusion of additional entries yields marginal or even negative returns, likely due to the introduction of irrelevant or noisy information. This effect is particularly evident in lower Recall@ k groups, where precision is more sensitive to input quality. For instance, in the Recall@2 group, increasing m from 2 to 5 leads to a decrease in accuracy, from 46.0% to 44.7% for the 1-Stage baseline and from 48.9% to 47.6% for our method.

Moreover, the impact of m varies across evaluation scenarios. While increasing the number of external entries initially benefits both Unseen-Q and Unseen-E, the latter demonstrates more stable and robust improvements. In contrast, performance degradation caused by excessive retrieval is more pronounced in Unseen-Q, highlighting its greater sensitivity to noisy or irrelevant knowledge.

Effect of Different Prompt. Table 4 presents a comprehensive zero-shot performance analysis between Qwen2-VL-1-Stage and our proposed method. For PRKI resolution, CoRe-MMRAG exhibits enhanced robustness to knowledge noise through carefully designed prompts. In Recall@1

Methods	Recall@5 (%)		
	Unseen-Q	Unseen-E	All
Ours	45.2	46.9	46.5
w/o $\mathcal{L}_{\text{PRKI}}$	44.1	46.0	45.5
w/o $\mathcal{L}_{\text{VTKI}}$	44.2	45.6	45.3
w/o \mathcal{L}_{SFT}	43.3	45.0	44.6

Table 5: Ablation study on training objectives.

group, when increasing retrieved entries from 1 to 2, our method maintains stable performance with accuracy shifting from 61.7% to 61.2%, while Qwen2-VL-1-Stage shows larger degradation from 58.8% to 58.2%. This advantage becomes more evident in the setting of Unseen-Questions, where our method preserves accuracy from 61.5% to 61.0% compared to the Qwen2-VL-1-Stage’s significant drop from 59.1% to 58.4%.

Moreover, our approach demonstrates superior multimodal knowledge integration capabilities. In Recall@2 group, increasing retrieved entries from 1 to 2 yields substantially larger gains as our method improves from 28.6% to 48.9% versus Qwen2-VL-1-Stage advancing from 28.0% to 46.0%. The improvement margin widens further in Unseen-Entities scenarios, with our method achieving progress from 27.7% to 49.0% while the baseline moves from 27.3% to 45.7%, confirming the effectiveness of our method on visual-textual fusion, as visualized in Figure 3.

Effect of Training Objectives. Table 5 demonstrates the contribution of each training objective. Removing any objective leads to performance degradation, with \mathcal{L}_{SFT} causing the most significant decline at 1.9% in overall performance. The absence of $\mathcal{L}_{\text{PRKI}}$ and $\mathcal{L}_{\text{VTKI}}$ also leads to notable performance drops at 1.0% and 1.2% respectively, indicating both objectives are essential for effective knowledge conflict and inconsistency resolution. Our full model achieves the best performance across all metrics, validating the complementary nature of these objectives.

5 Conclusion

In this paper, we present CoRe-MMRAG, a novel framework that addresses two critical challenges in MMRAG: parametric-retrieved knowledge inconsistency and visual-textual knowledge inconsistency. CoRe-MMRAG follows a four-stage pipeline that effectively reconciles internal parametric knowledge with externally retrieved information and leverages joint similarity assessment to

integrate complementary visual and textual signals. To further enhance its capabilities, we introduced a specialized training paradigm with three targeted objectives focused on knowledge source selection, multimodal integration, and answer generation. Extensive experiments on InfoSeek and Enc-VQA benchmarks demonstrate the effectiveness of our approach, achieving performance gains of 5.6% and 9.3% over baseline methods.

6 Limitations and Future Works

Despite the promising results, our work has several important limitations. First, our framework’s effectiveness is heavily dependent on the initial retrieval quality using Eva-CLIP-8B. As demonstrated in our experiments, the retrieval performance (Recall@1) remains relatively low at 45.6% for InfoSeek and 13.3% for Enc-VQA. This limited retrieval performance creates a ceiling effect for the overall system performance, suggesting that improvements in the initial retrieval stage could lead to significant gains in the final results. Second, our approach faces substantial computational challenges. The four-stage process and multiple training objectives require significant computational resources, with training taking approximately 10 hours on 8 H100 GPUs. The model requires substantial memory to process both visual and textual knowledge simultaneously, which may limit its practical deployment in resource-constrained environments. Finally, there are limitations in terms of dataset coverage and generalization. While we demonstrate strong performance on two specific KB-VQA benchmarks, the effectiveness of our approach on other types of multimodal tasks or real-world scenarios remains unexplored. The model’s performance on rare entities or edge cases in the knowledge base may be suboptimal, and there exists a potential domain gap between the Wikipedia knowledge base used for training and real-world applications. Future work should address these limitations to improve the practical applicability of our approach.

7 Acknowledgement

This work was supported in part by the National Natural Science Foundation of China, No.:62476071 and No.:62276155, as well as the Shenzhen Science and Technology Program under Grant No.: KQTD20240729102207002.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, pages 1–100.
- Omar Adjali, Olivier Ferret, Sahar Ghannay, and Hervé Le Borgne. 2024. Multi-level information retrieval augmented generation for knowledge-based visual question answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16499–16513.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avi Sil, and Hannaneh Hajishirzi. 2024. Self-rag: Learning to retrieve, generate, and critique through self-reflection. In *International Conference on Learning Representations*, pages 1–30.
- Zechen Bai, Pichao Wang, Tianjun Xiao, Tong He, Zongbo Han, Zheng Zhang, and Mike Zheng Shou. 2024. Hallucination of multimodal large language models: A survey. *arXiv preprint arXiv:2404.18930*, pages 1–40.
- Davide Caffagni, Federico Cocchi, Nicholas Moratelli, Sara Sarto, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. 2024. Wiki-llava: Hierarchical retrieval-augmented generation for multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1818–1826.
- Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, Alan Ritter, and Ming-Wei Chang. 2023. Can pre-trained vision and language models answer visual information-seeking questions? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 14948–14968.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*, pages 1–21.
- Wenbo Hu, Jia-Chen Gu, Zi-Yi Dou, Mohsen Fayyaz, Pan Lu, Kai-Wei Chang, and Nanyun Peng. 2025. Mrag-bench: Vision-centric evaluation for retrieval-augmented multimodal models. In *International Conference on Learning Representations*, pages 1–24.
- Nikhil Kandpal, Haikang Deng, Adam Roberts, Eric Wallace, and Colin Raffel. 2023. Large language models struggle to learn long-tail knowledge. In *International Conference on Machine Learning*, pages 15696–15707. PMLR.
- Paul Lerner, Olivier Ferret, Camille Guinaudeau, Hervé Le Borgne, Romaric Besançon, José G Moreno, and Jesús Lovón Melgarejo. 2022. Viquae, a dataset for knowledge-based visual question answering about named entities. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 3108–3120.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. 2024. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, pages 1–43.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*, pages 19730–19742. PMLR.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305.
- Zixu Li, Zhiwei Chen, Haokun Wen, Zhiheng Fu, Yupeng Hu, and Weili Guan. 2025. Encoder: Entity mining and modification relation binding for composed image retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5101–5109.
- Weizhe Lin, Jinghong Chen, Jingbiao Mei, Alexandru Coca, and Bill Byrne. 2023. Fine-grained late-interaction multi-modal retrieval for retrieval augmented visual question answering. *Advances in Neural Information Processing Systems*, 36:22820–22840.
- Fan Liu, Zhiyong Cheng, Lei Zhu, Zan Gao, and Liqiang Nie. 2021. Interest-aware message-passing gcn for recommendation. In *Proceedings of the Web Conference 2021*, page 1296–1305. ACM.
- Hanchao Liu, Wenyan Xue, Yifei Chen, Dapeng Chen, Xiutian Zhao, Ke Wang, Liping Hou, Rongjun Li, and Wei Peng. 2024a. A survey on hallucination in large vision-language models. *arXiv preprint arXiv:2402.00253*, pages 1–10.

- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. *Advances in Neural Information Processing Systems*, 36:34892–34916.
- Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. 2018. Attentive moment retrieval in videos. In *Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pages 15–24.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3195–3204.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *Findings of the Association for Computational Linguistics*, pages 2263–2279.
- Thomas Mensink, Jasper Uijlings, Lluís Castrejon, Arushi Goel, Felipe Cadar, Howard Zhou, Fei Sha, André Araujo, and Vittorio Ferrari. 2023. Encyclopedic vqa: Visual questions about detailed properties of fine-grained categories. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3113–3124.
- Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. 2020. Plotqa: Reasoning over scientific plots. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1527–1536.
- Jingyuan Qi, Zhiyang Xu, Rulin Shao, Yang Chen, Jin Di, Yu Cheng, Qifan Wang, and Lifu Huang. 2024. Rora-rlm: Robust retrieval-augmented vision language models. *arXiv preprint arXiv:2410.08876*, pages 1–15.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, pages 1–154.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022. A-okvqa: A benchmark for visual question answering using world knowledge. In *European Conference on Computer Vision*, pages 146–162. Springer.
- Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. 2023. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, pages 1–7.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, pages 1–52.
- Xiao Wang, Jianlong Wu, Zijia Lin, Fuzheng Zhang, Di Zhang, and Liqiang Nie. 2025. Video datafly-wheel: Resolving the impossible data trinity in video-language understanding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(1):1–13.
- Yibin Yan and Weidi Xie. 2024. Echosight: Advancing visual-language models with wiki knowledge. In *Findings of the Empirical Methods in Natural Language Processing*, pages 1538–1551.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. 2024. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, pages 1–26.
- Zhenfei Yin, Jiong Wang, Jianjian Cao, Zhelun Shi, Dingning Liu, Mukai Li, Xiaoshui Huang, Zhiyong Wang, Lu Sheng, Lei Bai, et al. 2023. Lamm: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark. *Advances in Neural Information Processing Systems*, 36:26650–26685.
- Eric Zelikman, Yuhuai Wu, Jesse Mu, and Noah Goodman. 2022. Star: Bootstrapping reasoning with reasoning. *Advances in Neural Information Processing Systems*, 35:15476–15488.
- Jingyi Zhang, Jiaxing Huang, Sheng Jin, and Shijian Lu. 2024a. Vision-language models for vision tasks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5625–5644.
- Tao Zhang, Ziqi Zhang, Zongyang Ma, Yuxin Chen, Zhongang Qi, Chunfeng Yuan, Bing Li, Junfu Pu, Yuxuan Zhao, Zehua Xie, et al. 2024b. Mr²ag: Multimodal retrieval-reflection-augmented generation for knowledge-based vqa. *arXiv preprint arXiv:2411.15041*, pages 1–14.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, pages 1–43.

A Appendix

A.1 Prompts

In this section, we present the prompts used in different MMRAG pipelines, including Qwen2-VL-Param, Qwen2-VL-Oracle, Qwen2-VL-1-Stage, Qwen2-VL-2-Stage, Qwen2-VL-MMSTaR, and our proposed CoRe-MMRAG. Each prompt is illustrated with the format, an example, and the corresponding final answer.

Prompt for Qwen2-VL-Param

Prompt format:

Question: <image>
{question}
Please use parametric knowledge answer the
→ question within 5 words

Example:

Question: <image>
Which historic county does this village belong
→ to?
Please use parametric knowledge answer the
→ question within 5 words

Answer:

"Scotland"

The prompt for Qwen2-VL-Oracle is constructed using the ground-truth Wikipedia entry. Specifically, `wiki_title_gt` and `wiki_content_gt` refer to the title and textual content of the ground-truth entry, respectively.

Prompt for Qwen2-VL-Oracle

Prompt format:

Based on the retrieved document, answer the
→ question
<image>
{question} within 5 words
Context:
Reference Image:<image>
Wiki title: {wiki_title_gt}
Wiki content:{wiki_content_gt}

Example:

Based on the retrieved document, answer the
→ question
<image>
What **is** the parent organization of this building
→ ? within 5 words.
Context:
Reference Image:<image#gt>
Wiki title: Colonial Williamsburg
Wiki content:Colonial Williamsburg **is** a living-
→ history museum **and** private foundation
→ presenting a part of the historic
→ disconstructions...

Answer:

"The Colonial Williamsburg Foundation"

The following is the prompt for Qwen2-VL-1-Stage, where `wiki_title_A/B/C/D/E` and `wiki_content_A/B/C/D` → /E represent the titles and textual contents of the top-5 retrieved Wikipedia entries.

Prompt for Qwen2-VL-1-Stage

Prompt format:

Based on the retrieved document, answer the
→ question
<image>
{question} within 5 words
Reference A:
<image#A>
Wiki title: {wiki_title_A}
Wiki content:{wiki_content_A}
Reference B:
<image#B>
Wiki title: {wiki_title_B}
Wiki content:{wiki_content_B}
Reference C:
<image#C>
Wiki title: {wiki_title_C}
Wiki content:{wiki_content_C}
Reference D:
<image#D>
Wiki title: {wiki_title_D}
Wiki content:{wiki_content_D}
Reference E:
<image#E>
Wiki title: {wiki_title_E}
Wiki content:{wiki_content_E}

Example:

Based on the retrieved document, answer the
→ question
<image>
What **is** the width (in kilometre) of this lake?
→ within 5 words.

Reference A
<image#A>
Wiki title: Wadi Numeira
Wiki content: Wadi Numeira **is** a Wadi in Jordan
→ that **is** known **for** its deep gorge cut through
→ the sandstone...

Reference B
<image#B>
Wiki title: Kalya
Wiki content: Kalya **is** an Israeli settlement
→ organized as a kibbutz **in** the West Bank..

Reference C
<image#C>
Wiki title: Ein Gedi
Wiki content: Ein Gedi, also spelled En Gedi,
→ meaning \"spring of the kid\", is an oasis and
→ a nature reserve in Israel...

Reference D
<image#D>
Wiki title: Dead Sea
Wiki content: The Dead Sea, also known by other
→ names, is a salt lake bordered by Jordan to
→ the east and Israel...

Reference E
<image#E>
Wiki title: Jewish National Fund
Wiki content: Jewish National Fund was founded
→ in 1901 to buy and develop land in Ottoman
→ Syria..."

Answer:

"10"

The following is the prompt for Qwen2-VL-2-Stage, where `wiki_title_A/B/C/D/E` and `wiki_content_A/B/C/` → D/E denote the titles and textual contents of the top-5 retrieved Wikipedia entries. `wiki_title_select` and `wiki_content_select` refer to the title and content of the entry selected during the first-stage reranking.

Prompt for Qwen2-VL-2-Stage

Prompt format for rerank:

```
Identify the most similar wiki context to the
→ question
<image>
{question}
Context A:
Wiki title: {wiki_title_A}
Wiki content:{wiki_content_A}
Context B:
Wiki title: {wiki_title_B}
Wiki content:{wiki_content_B}
Context C:
Wiki title: {wiki_title_C}
Wiki content:{wiki_content_C}
Context D:
Wiki title: {wiki_title_D}
Wiki content:{wiki_content_D}
Context E:
Wiki title: {wiki_title_E}
Wiki content:{wiki_content_E}
The answer should be provided in format: [
→ Reference X] where X is the most similar
→ reference (A/B/C/D/E)
```

Prompt format for generation:

```
Based on the retrieved document, answer the
→ question
<image>
{question} within 5 words.
Context:
Wiki title: {wiki_title_select}
Wiki content: {wiki_content_select}
```

Example:

```
# Rerank:
...Context A: ...
# See Prompt for Qwen2-VL-1-Stage

The answer should be provided in format: [
→ Reference X] where X is the most similar
→ reference (A/B/C/D/E)

# Generation:
Based on the retrieved document, answer the
→ question
<image>
What is the width (in kilometre) of this lake?
→ within 5 words

Context:
Wiki title: Jewish National Fund
Wiki content: Jewish National Fund was founded
→ in 1901 to buy and develop land in Ottoman
→ Syria...
```

Answer:

```
"Reference E"
"Not given"
```

The following is the prompt for Qwen2-VL-MMSTaR, where `wiki_title_A/B/C/D/E` and `wiki_content_A/B/C/D/E` represent the titles and textual contents of the top-5 retrieved Wikipedia entries.

Prompt for Qwen2-VL-MMSTaR (Part I)

Prompt format:

```
Based on the retrieved document, answer the
→ question
<image>
{question} within 5 words
Reference A:
<image#A>
Wiki title: {wiki_title_A}
Wiki content:{wiki_content_A}
```

Prompt for Qwen2-VL-MMSTaR (Part II)

Prompt format:

```
Reference B:
<image#B>
Wiki title: {wiki_title_B}
Wiki content:{wiki_content_B}
Reference C:
<image#C>
Wiki title: {wiki_title_C}
Wiki content:{wiki_content_C}
Reference D:
<image#D>
Wiki title: {wiki_title_D}
Wiki content:{wiki_content_D}
Reference E:
<image#E>
Wiki title: {wiki_title_E}
Wiki content:{wiki_content_E}
Note: Analyze step-by-step and provide final
→ answer in format: []
```

Example:

```
Based on the retrieved document, answer the
→ question
<image>
What place inflows lake? within 5 words.
Reference A: ...
# See Prompt for Qwen2-VL-1-Stage
Note: Analyze step-by-step and provide final
→ answer in format: []
```

Answer:

```
"Step 1: Key Information Analysis
Query Image Analysis:
The query image shows a scenic view of a lake
→ surrounded by mountains and a town. The lake
→ is prominent in the center, with a town on
→ its shores.

Question Analysis:
The question asks about the place that inflows
→ into the lake.

Step 2: Reference Comparison:
Reference A: Describes Gsteigwiler, a village
→ near a lake.

Reference B: Describes Walensee, a lake in
→ Switzerland.

Reference C: Describes Schnepfau, a village near
→ a lake.

Reference D: Describes Lake Lucerne, a lake in
→ Switzerland.

Reference E: Describes Lake Brienz, a lake in
→ Switzerland.

Relevant Reference Selection:
Reference B (Walensee) and Reference D (Lake
→ Lucerne) both describe lakes in Switzerland,
→ but Reference B is more likely to have a
→ river inflowing into it, as it is a larger
→ lake.

Step 3: Answer Generation
Answer: The place that inflows into Lake Walensee
→ is the Seez river. According to the
→ description in Reference B, the Seez,
→ Murgbach, and Linth rivers lead to the lake.
→ The Seez river is the most likely candidate
→ for inflowing into the lake, so the final
→ answer is [Seez river]"
```

The following is the prompt for CoRe-MMRAG, where `wiki_title_A/B/C/D/E` and `wiki_content_A/B/C/D/E` represent the titles and textual contents of the top-5 retrieved Wikipedia entries.

Prompt for CoRe-MMRAG

Prompt format:

```
Question:
<image>
{question}
Step 1:
If the question with image is directly
→ answerable, provide the answer in [] less
→ than 5 words.
If not, proceed to Step 2.

Step 2:
Identify the most similar Wikipedia reference to
→ the {question}
Reference A:
<image#A>
Wiki title: {wiki_title_A}
Wiki content:{wiki_content_A}
Reference B:
<image#B>
Wiki title: {wiki_title_B}
Wiki content:{wiki_content_B}
Reference C:
<image#C>
Wiki title: {wiki_title_C}
Wiki content:{wiki_content_C}
Reference D:
<image#D>
Wiki title: {wiki_title_D}
Wiki content:{wiki_content_D}
Reference E:
<image#E>
Wiki title: {wiki_title_E}
Wiki content:{wiki_content_E}

Step 3:
Based on the most similar reference, answer the
→ question
<image>
{question}
within 5 words

Step 4:
If the answers from Step 1 and Step 3 differ,
→ determine which one is the final answer,
→ output the final answer in less than 5 words.
```

Example:

```
Question:
<image>
What is the width (in kilometre) of this lake?
Step 1:
If the question with image is directly
→ answerable, provide the answer in [] less
→ than 5 words.
If not, proceed to Step 2.

Step 2:
Identify the most similar Wikipedia reference to
→ What is the width (in kilometre) of this
→ lake?
Reference A...
# See Prompt for Qwen2-VL-1-Stage

Step 3:
Based on the most similar reference, answer the
→ question
<image>
{question}
within 5 words

Step 4:
If the answers from Step 1 and Step 3 differ,
→ determine which one is the final answer,
→ output the final answer in less than 5 words.
```

Answer:

```
"2.8 km"
```

A.2 Case Study

A.2.1 VTKI Examples

The cases for VTKI are illustrated in Figure 4. For each input query, we present the top-5 retrieved references (A–E) ranked by embedding similarity between the reference images and the query image, computed using EVA-CLIP-8B. The examples in the figure demonstrate that the Qwen-2-VL model produces different outputs when ranking is based solely on image modality versus text modality. When relying only on image similarity, the model may incorrectly select visually similar but semantically irrelevant references. For instance, in the second row of Figure 4, the model selects reference C as the most similar, but its associated text discusses *Bursaphelenchus xylophilus*, which is unrelated to the question "How many offspring can this bird produce at the same time?" Conversely, when using only textual information, as shown in the fourth row of Figure 4, the model mistakenly selects reference B over reference C. Both references describe food items resembling pudding, but without visual cues, the model cannot determine which is more appropriate. These examples highlight the VTKI problem, where unimodal approaches may lead to incorrect reference selection. In contrast, our proposed CoRe-MMRAG method effectively addresses this issue by leveraging both modalities. When unimodal outputs diverge but include the correct answer, CoRe-MMRAG enables the model to identify and choose the correct reference.

A.2.2 PRKI Examples

The cases for PRKI are illustrated in Figure 5, which follows the same ranking settings as Figure 4. These examples demonstrate that the presence of noisy information in external data can negatively impact the model's ability to express accurate parametric knowledge, leading to incorrect outputs. In such PRKI scenarios, our proposed model effectively mitigates this issue. Notably, the first and second rows in Figure 5 show that our method successfully outputs the correct parametric knowledge, explicitly marked with square brackets ("[]").































 <p>What is the parent organization of this building?</p>	 <p>A</p> <p>Colonial Williamsburg Colonial Williamsburg is a living-history museum and private foundation presenting a part of the historic ...</p>	 <p>B</p> <p>List of the oldest buildings in South Carolina This article attempts to list the oldest extant buildings surviving in the state of South Carolina...</p>	 <p>C</p> <p>Colonial Williamsburg Colonial Williamsburg is a living-history museum and private foundation presenting a part of the historic ...</p>	 <p>D</p> <p>National Register of Historic Places listings in Virginia Beach, Virginia This is a list of the National Register of Historic Places listings in Virginia Beach ...</p>	 <p>E</p> <p>Cherokee Phoenix Cherokee Phoenix is the first newspaper published by Native Americans in the United States and the first published ...</p>
	Ground-Truth Reference: A	By Image Relevance: B	By Text Relevance: A	CoRe-MMRAG: A	
 <p>How many offspring can this bird produce at the same time?</p>	 <p>A</p> <p>Great spotted woodpecker The great spotted woodpecker is a medium-sized woodpecker with pied black and white plumage and a red patch ...</p>	 <p>B</p> <p>Great spotted woodpecker The great spotted woodpecker is a medium-sized woodpecker with pied black and white plumage and a red patch ...</p>	 <p>C</p> <p>Bursaphelenchus xylophilus Bursaphelenchus xylophilus, commonly known as pine wilt nematode or pine wilt nematode (PWN), is a species of nematode ...</p>	 <p>D</p> <p>Woodpecker Woodpecker are part of the family Picidae, which also includes the piculets, wrynecks, and sapsuckers. Members of this family ...</p>	 <p>E</p> <p>Great spotted woodpecker The great spotted woodpecker is a medium-sized woodpecker with pied black and white plumage and a red patch ...</p>
	Ground-Truth Reference: A	By Image Relevance: C	By Text Relevance: A	CoRe-MMRAG: A	
 <p>What is the width (in kilometer) of this lake?</p>	 <p>A</p> <p>Wadi Numeira Wadi Numeira is a Wadi in Jordan that is known for its deep gorge cut through the sandstone. It gives its name to the Bronze ...</p>	 <p>B</p> <p>Kalya Kalya is an Israeli settlement organized as a kibbutz in the West Bank. It was originally established in 1929 ...</p>	 <p>C</p> <p>Ein Gedi Ein Gedi, also spelled En Gedi, meaning "spring of the kid", is an oasis and a nature reserve in Israel, located west of ...</p>	 <p>D</p> <p>Dead Sea The Dead Sea, also known by other names, is a salt lake bordered by Jordan to the east and Israel and the West Bank to the west..</p>	 <p>E</p> <p>Jewish National Fund It has also built 180 dams and reservoirs, developed 250,000 acres (1,000 km²) of land and established more than 1,000 parks ...</p>
	Ground-Truth Reference: D	By Image Relevance: D	By Text Relevance: E	CoRe-MMRAG: D	
 <p>Who is the owner of this food?</p>	 <p>A</p> <p>Coconut The coconut tree (Cocos nucifera) is a member of the palm tree family (Arecaceae) and the only living species of the genus Cocos...</p>	 <p>B</p> <p>Jell-O Jell-O is a variety of gelatin desserts (fruit-flavored gels), puddings, and no-bake cream pies. The original Jell-O ...</p>	 <p>C</p> <p>Crème caramel Crème caramel, flan, caramel pudding or caramel custard is a custard dessert with a layer of clear caramel sauce ...</p>	 <p>D</p> <p>Dessert Dessert is a course that concludes a meal. The course consists of sweet foods, such as confections, and possibly ...</p>	 <p>E</p> <p>Jell-O Jell-O is a variety of gelatin desserts (fruit-flavored gels), puddings, and no-bake cream pies. The original Jell-O ...</p>
	Ground-Truth Reference: B	By Image Relevance: B	By Text Relevance: C	CoRe-MMRAG: B	
 <p>What is the date this aircraft took the first flight?</p>	 <p>A</p> <p>349th Squadron (Belgium) 349th Squadron is a fighter squadron in the Air Component of the Belgian Armed Forces. The squadron traces its origins...</p>	 <p>B</p> <p>RAF Leuchars Royal Air Force Leuchars or RAF Leuchars was a Royal Air Force (RAF) station located in Leuchars, Fife ...</p>	 <p>C</p> <p>Supermarine Spitfire The Supermarine Spitfire is a British single-seat fighter aircraft used by the Royal Air Force and other Allied countries before ...</p>	 <p>D</p> <p>Supermarine Spitfire The Supermarine Spitfire is a British single-seat fighter aircraft used by the Royal Air Force and other Allied countries before ...</p>	 <p>E</p> <p>Supermarine Spitfire The Supermarine Spitfire is a British single-seat fighter aircraft used by the Royal Air Force and other Allied countries before ...</p>
	Ground-Truth Reference: C	By Image Relevance: A	By Text Relevance: C	CoRe-MMRAG: C	

Figure 4: Qualitative results on sample image-question pairs from the InfoSeek dataset. The leftmost column displays the input query image. References A–E are ordered by descending embedding similarity (computed using EVA-CLIP-8B) between the reference images and the query image. "By image relevance" presents the Qwen-2-VL output of most relevant references based solely on reference images, while "By text relevance" relies only on reference texts. Our proposed CoRe-MMRAG model leverages multimodal information to select the most relevant references.

























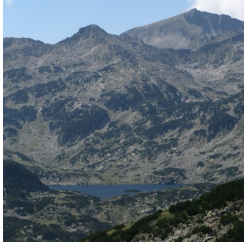





 <p>What country does this island belong to?</p>	 <p>A</p> <p>List of extreme points of the Netherlands This is a list of the extreme points of the Netherlands, the points that are farther... down, north, south, east...</p>	 <p>B</p> <p>Stromboli Stromboli is a sea stack of volcanic origin 2 km (1.2 mi) to the northeast of the island of Stromboli in the Aeolian...</p>	 <p>C ★</p> <p>Panarea Panarea is the smallest of the seven inhabited Aeolian Islands, a volcanic island chain in north of Sicily, southern Italy ...</p>	 <p>D</p> <p>Lehua Lehua Island is a small, crescent-shaped island in the Hawaiian islands, 0.7 miles (1.1 km) north of Ni'ihau ...</p>	 <p>E</p> <p>Volcano Volcano is a rupture in the crust of a planetary-mass object, such as Earth, that allows hot lava, volcanic ash, and gases to ...</p>
	Ground-Truth: Italy	Qwen2-VL-Param: Italy	Qwen2-VL-1-Stage: The Netherlands	CoRe-MMRAG: [Italy]	
 <p>Who is the developer of this aircraft?</p>	 <p>A</p> <p>Water landing In aviation, a water landing is, in the broadest sense, an aircraft landing on a body of water. Seaplanes, such as floatplanes and flying boats.</p>	 <p>B</p> <p>Maldivian Air Taxi Maldivian Air Taxi (MAT) was a domestic carrier in the Maldives and was one of the largest seaplane operators in the world ...</p>	 <p>C ★</p> <p>De Havilland Canada DHC-6 Twin Otter The de Havilland Canada DHC-6 Twin Otter is an aircraft developed by de Havilland Canada</p>	 <p>D</p> <p>De Havilland Canada DHC-6 Twin Otter The de Havilland Canada DHC-6 Twin Otter is an aircraft developed by de Havilland Canada</p>	 <p>E</p> <p>De Havilland Canada DHC-6 Twin Otter The de Havilland Canada DHC-6 Twin Otter is an aircraft developed by de Havilland Canada</p>
	Ground-Truth: de Havilland Canada	Qwen2-VL-Param: de Havilland Canada	Qwen2-VL-1-Stage: Viking Air	CoRe-MMRAG: [de Havilland Canada]	
 <p>What is the total quantity of produced items for this type of aircraft?</p>	 <p>A</p> <p>Czech Airlines Czech Airlines is the flag carrier of the Czech Republic. Its head office is located in the Vokovice area of Prague's 6th district ...</p>	 <p>B</p> <p>List of defunct airlines of Bulgaria This is a list of defunct airlines of Bulgaria.</p>	 <p>C ★</p> <p>Airbus A310 The Airbus A310 is a wide-body aircraft. By the end of production, a total of 255 A310s had been ordered and delivered ...</p>	 <p>D</p> <p>Finnair Finnair is the flag carrier and largest airline of Finland, with its headquarters in Vantaa on the grounds ...</p>	 <p>E</p> <p>Mahan Air Mahan Air, operating under the name Mahan Air, is a privately owned Iranian airline based in Tehran, Iran. It operates scheduled...</p>
	Ground-Truth: 255	Qwen2-VL-Param: 1300	Qwen2-VL-1-Stage: 255	CoRe-MMRAG: 255	
 <p>Which historic county does this village belong to?</p>	 <p>A</p> <p>List of Category A listed buildings in South Lanarkshire This is a list of Category A listed buildings in South Lanarkshire...</p>	 <p>B</p> <p>List of Category A listed buildings in South Lanarkshire This is a list of Category A listed buildings in South Lanarkshire...</p>	 <p>C ★</p> <p>New Lanark New Lanark is a village on the River Clyde, approximately 1.4 miles (2.2 kilometres) from Lanark, in South Lanarkshire</p>	 <p>D</p> <p>List of Category A listed buildings in South Lanarkshire This is a list of Category A listed buildings in South Lanarkshire...</p>	 <p>E</p> <p>List of Category A listed buildings in South Lanarkshire This is a list of Category A listed buildings in South Lanarkshire...</p>
	Ground-Truth: South Lanarkshire	Qwen2-VL-Param: Scotland	Qwen2-VL-1-Stage: South Lanarkshire	CoRe-MMRAG: South Lanarkshire	
 <p>Which city or region does this park locate in?</p>	 <p>A</p> <p>Sivria Sivria is a 2,591 m high peak in the Pirin mountain range, south-western Bulgaria. It is located in the northern part of ...</p>	 <p>B</p> <p>Blagoevgrad Province Blagoevgrad Province also known as Pirin Macedonia or Bulgarian Macedonia is a province (oblast) of southwestern Bulgaria ...</p>	 <p>C</p> <p>Province of Brescia The Province of Brescia is a Province in the Lombardy administrative region of northern Italy. It has a population...</p>	 <p>D</p> <p>Rila National Park Rila National Park is the largest national park in Bulgaria spanning an area of 810.46 km² (312.92 sq mi; 200,270 acres) ...</p>	 <p>E ★</p> <p>Pirin National Park Pirin National Park, originally named Vihren National Park, The park is situated in Blagoevgrad Province...</p>
	Ground-Truth: Blagoevgrad Province	Qwen2-VL-Param: Bulgaria	Qwen2-VL-1-Stage: Blagoevgrad Province	CoRe-MMRAG: Blagoevgrad Province	

Figure 5: Qualitative final answer outputs on sample image-question pairs from the InfoSeek dataset. The leftmost column shows the input query image. References A–E are ordered by descending embedding similarity (computed using EVA-CLIP-8B) between the reference images and the query image. A star symbol indicates the ground-truth reference.