

KG-ViP: Bridging Knowledge Grounding and Visual Perception in Multi-modal LLMs for Visual Question Answering

Zhiyang Li^{1,2,*} Ao Ke^{1,2,*} Yukun Cao^{2,3,†} Xike Xie^{1,2,†}

¹University of Science and Technology of China, China

²Data Darkness Lab, MIRACLE Center, USTC

³School of Computer Science and Technology, Xidian University

{lizhiyang, sa21225249}@mail.ustc.edu.cn

caoyukun@xidian.edu.cn, xkxie@ustc.edu.cn

Abstract

Multi-modal Large Language Models (MLLMs) for Visual Question Answering (VQA) often suffer from dual limitations: knowledge hallucination and insufficient fine-grained visual perception. Crucially, we identify that commonsense graphs and scene graphs provide precisely complementary solutions to these respective deficiencies by providing rich external knowledge and capturing fine-grained visual details. However, prior works typically treat them in isolation, overlooking their synergistic potential. To bridge this gap, we propose **KG-ViP**, a unified framework that empowers MLLMs by fusing scene graphs and commonsense graphs. The core of the KG-ViP framework is a novel retrieval-and-fusion pipeline that utilizes the query as a semantic bridge to progressively integrate both graphs, synthesizing a unified structured context that facilitates reliable multi-modal reasoning. Extensive experiments on FVQA 2.0+ and MVQA benchmarks demonstrate that KG-ViP significantly outperforms existing VQA methods.

1 Introduction

Multi-modal large language models (MLLMs) have demonstrated strong performance on Visual Question Answering (VQA) (Marino et al., 2019; Liu et al., 2024; Xu et al., 2024; Xiong et al., 2025), a task that requires joint reasoning over visual content and textual queries. However, MLLMs still suffer from dual limitations stemming from deficiencies in both knowledge grounding and visual perception.

From the knowledge perspective, when confronted with knowledge-intensive questions, MLLMs are prone to generating hallucinations (Li et al., 2023; Guan et al., 2024; Li et al., 2025b). To mitigate this issue, Retrieval-Augmented

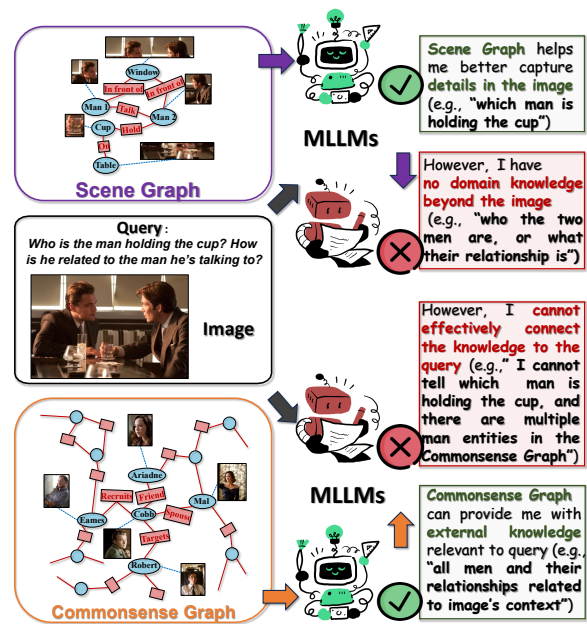


Figure 1: **Complementary roles of Scene Graphs and Commonsense Graphs.** Using a scene from *Inception* as an example: Given the query asking “Who is the man holding the cup?”, the **scene graph** explicitly captures visual relations (e.g., *holding*) yet fails to recognize character identities (*Cobb, Robert*). Conversely, the **commonsense graph** offers semantic facts (e.g., *Cobb targets Robert*) but suffers from grounding ambiguity, as it lacks visual cues to anchor abstract nodes to the specific person in the image. Thus, neither graph alone can support the complete reasoning chain.

Generation (RAG) (Fan et al., 2024) has been widely adopted to incorporate external knowledge. In particular, Graph-based RAG (Edge et al., 2024; Sun et al., 2024; Cao et al., 2025) has recently emerged as a compelling paradigm, where knowledge is organized into structured entities and relations (e.g., *commonsense graphs*), enabling more precise knowledge retrieval and facilitating explicit multi-hop reasoning.

From the perception perspective, prior work (Herzig et al., 2023; Ma et al., 2023; Mitra et al., 2024) has revealed that MLLMs exhibit limited visual understanding, particularly in capturing fine-grained relationships and interactions within complex scenes. To enhance the perceptual

* Equal contribution.

† Corresponding authors.

capability of MLLMs, a promising direction is to introduce *scene graphs* (Mitra et al., 2024; Wang et al., 2024), which provide a structured representation of images by explicitly modeling visual objects and their spatial relations.

Crucially, although both lines of research adopt a similar strategy—utilizing structured representations via commonsense and scene graphs—they are typically treated as independent components. This separation overlooks the potential synergy derived from their structural similarity and semantic complementarity. Concretely, as illustrated in Figure 1, scene graphs excel at visual grounding, capturing explicit entities and their dynamic interactions (e.g., *holding*, *next to*) directly from the image; however, they are limited to visible content and lack the external context required for deep reasoning. In contrast, commonsense graphs encapsulate rich semantic priors and relational knowledge (e.g., *is friend of*, *used for*). Yet, their effectiveness heavily relies on accurate entity alignment; without precise visual anchoring, the retrieval process is susceptible to visual noise, often introducing irrelevant or misleading information.

Motivated by these observations, we propose to bridge visual perception and external knowledge by fusing scene and commonsense graphs into a unified structured representation to empower reasoning. While conceptually intuitive, realizing this synergy poses non-trivial challenges across the entire pipeline, ranging from graph construction and efficient retrieval to effective cross-modal fusion.

First, regarding graph construction, widely used commonsense graphs in VQA are predominantly textual, inherently limiting joint reasoning across modalities (Yuan et al., 2025). Therefore, transforming them into Multi-modal Knowledge Graphs (MMKGs) (Liu et al., 2019) is essential for reliable answer generation. Nevertheless, existing MMKGs (Lee et al., 2024; Zhu et al., 2022) are largely generic, and pipelines for constructing domain-specific MMKGs tailored to complex VQA tasks are not yet well-established.

Second, in retrieval, bridging visual objects with external knowledge remains challenging. Existing VQA methods (Sarwar, 2025; Jian et al., 2024) typically rely on coarse global vision–text alignment, which often overlooks fine-grained visual cues (e.g., interactions) and cause ambiguity when multiple entities share similar textual descriptions. Finally, regarding graph fusion, the inherent heterogeneity between commonsense and scene graphs

poses a significant obstacle. Since the two graphs differ in modality, topology, and node attributes, achieving reliable node alignment and coherent cross-modal fusion remains a challenge.

To address these issues, we propose **KG-ViP**, a unified framework that bridges fine-grained visual perception and external knowledge. First, to ensure knowledge support, we introduce a robust pipeline for constructing domain-specific commonsense graphs in the form of MMKGs. This pipeline accommodates diverse data scenarios, constructing the graph from heterogeneous sources and enriching entities with visual attributes. Building on this, we design a novel Retrieval–Fusion–Generation workflow for KG-ViP. Specifically, we employ a two-stage retrieval strategy that integrates textual and visual modalities. It utilizes expanded query semantics derived from textual retrieval to prune the noisy scene graph, and subsequently employs the refined visual cues to conduct visual retrieval, thereby precisely anchoring external knowledge. This process yields query-aware scenes and commonsense subgraphs. To fuse these heterogeneous graphs, we implement a comprehensive strategy comprising cross-modal entity alignment and agent-based refinement, resulting in a unified graph. Finally, this graph is fed into the MLLMs, providing enhanced support for reliable reasoning.

The main contributions of this paper are summarized as follows: (1) We identify the dual limitations of MLLMs in knowledge grounding and visual perception, advocating for a synergistic integration of scene and commonsense graphs to enhance joint reasoning; (2) To realize this, we propose KG-ViP, a unified framework featuring a pipeline for graph construction, a two-stage retrieval strategy, and a cross-modal fusion mechanism; (3) Extensive empirical studies demonstrate the effectiveness and superiority of our framework across various evaluation tasks.

2 Related Work

2.1 MLLMs for Knowledge-based VQA

Knowledge-Based VQA (KB-VQA) typically requires incorporating external knowledge to answer questions that demand information beyond visual content. Recently, with the rapid advancement of MLLMs, leveraging MLLMs for KB-VQA has garnered significant research attention (Gui et al., 2022; Si et al., 2023; Lin et al., 2023a). In particular, retrieval augmentation has proven effective in

enhancing KB-VQA performance, as demonstrated by representative methods such as RAVQA (Lin and Byrne, 2022), Wiki-LLaVA (Caffagni et al., 2024), EchoSight (Yan and Xie, 2024), FilterRAG (Sarwar, 2025), and LLM-RA (Jian et al., 2024). Specifically, FilterRAG retrieves knowledge by jointly embedding the question and the image split into 2×2 patches, while LLM-RA employs MLLMs to generate image captions from which key entities are extracted for retrieval. However, these methods largely overlook the rich structural scene information and fine-grained details inherent in images. We explicitly model images as scene graphs to guide knowledge retrieval, significantly enhancing answer accuracy.

2.2 Scene Graph for MLLM Reasoning

Scene graphs bridge the visual-semantic gap by explicitly modeling objects and their spatial interactions (Chang et al., 2021; Li et al., 2024). While foundational in early VQA, they have recently become instrumental in enhancing the fine-grained perception of MLLMs (Yang et al., 2025; Kim et al., 2024). For instance, CCOT (Mitra et al., 2024) interprets images as scene graphs to augment CoT prompting for MLLMs. M3COT (Lee et al., 2025) improves VQA robustness by unifying multi-view scene graphs, while MMCD (Li et al., 2025b) employs structural perturbations of scene graphs to suppress model hallucinations. However, these methods primarily treat scene graphs as internal visual parsing artifacts and overlook the potential role as explicit anchors for external knowledge. Consequently, they remain limited in knowledge-intensive reasoning, motivating the integration of scene and commonsense graphs.

3 Methodology

In this work, we focus on Knowledge-based VQA. Given an image I_q and a textual query Q , the goal is to generate an answer A by bridging visual perception with external knowledge. As illustrated in Figure 2, our KG-ViP framework comprises three distinct stages, detailed as follows.

3.1 Multi-modal Graph Construction

An MMKG extends traditional KGs by integrating structured relational facts with heterogeneous data from diverse modalities. Formally, we define an MMKG as $\mathcal{G} = (\mathcal{E}, \mathcal{R})$, where \mathcal{E} denotes the set of entities (nodes), and \mathcal{R} represents the set of relations (edges) connecting these entities.

Each entity $e \in \mathcal{E}$ is associated with various attributes, encompassing multi-modal content (e.g., images) and metadata (e.g., textual labels). In this work, we construct two distinct types of MMKGs to facilitate visual reasoning: **(1) Commonsense Graph \mathcal{G}_c** : An MMKG encapsulating large-scale, domain-specific commonsense facts spanning diverse modalities. For instance, within the cinematic domain, this graph organizes character relationships and plot synopses into an MMKG structure, enriched with visual content such as character portraits. **(2) Scene Graph \mathcal{G}_s** : An MMKG derived from the specific query image provided for MLLM reasoning tasks. It captures concrete objects along with their spatial relationships and interactions to aid query-specific visual understanding. For example, given a single movie still, this graph explicitly models the spatial configurations and actions of the depicted characters and objects.

3.1.1 Commonsense Graph Construction

We prioritize using off-the-shelf commonsense graphs when available for the target domain. Otherwise, we construct the graph from heterogeneous sources (e.g., multi-modal data, existing textual knowledge bases) via a two-step pipeline: *Textual Graph Construction* and *Multi-modal Entity Alignment*. Notably, the pipeline is orthogonal to specific MMKG construction techniques, allowing for the integration of various existing methods (Zhu et al., 2022; Lee et al., 2024; Yuan et al., 2025). We outline the general workflow below.

Textual Graph Construction. The first stage transforms heterogeneous sources—including unstructured texts and existing structured knowledge bases—into a unified graph. Crucially, we leverage LLMs to extract salient entities and explicit relations from unstructured texts (e.g., movie plot synopses) (Edge et al., 2024; Guo et al., 2024), constructing the graph from scratch. Additionally, we can integrate existing structured knowledge bases by consolidating independent triples via standard techniques, such as embedding-based entity alignment (Sun et al., 2020) and semantic relation normalization (Lomaeva and Jain, 2022; Zhang and Soh, 2024). To ensure topological quality, the graph undergoes refinement—such as pruning low-confidence edges and extracting the giant connected component (Lin et al., 2024).

Multi-modal Entity Alignment The second stage enriches the textual entities with corresponding multi-modal objects (e.g., images or video clips).

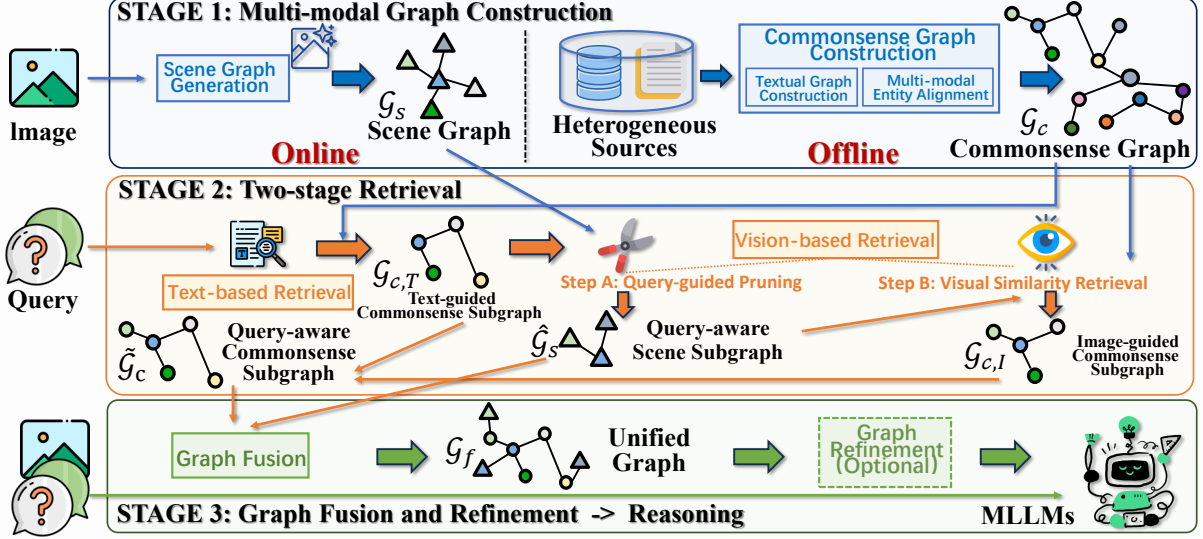


Figure 2: **The Overview of KG-ViP Framework.** KG-ViP operates in three stages. First, we prepare a domain *Commonsense Graph* \mathcal{G}_c and derive a *Scene Graph* \mathcal{G}_s from the input image I_q . Second, the query Q retrieves a text-guided subgraph $\mathcal{G}_{c,T}$ from \mathcal{G}_c , which is used to filter \mathcal{G}_s into a refined $\hat{\mathcal{G}}_s$. Subsequently, $\hat{\mathcal{G}}_s$ guides the retrieval of visual knowledge $\mathcal{G}_{c,I}$ to form the final commonsense subgraph $\tilde{\mathcal{G}}_c$. Finally, $\hat{\mathcal{G}}_s$ and $\tilde{\mathcal{G}}_c$ are fused into a unified graph \mathcal{G}_f , providing the structured context for the MLLM to generate the answer A .

For entities associated with official media repositories or explicit metadata, we employ high-precision matching based on semantic consistency. Conversely, for entities lacking direct associations, we utilize an LLM-driven web search agent (Kangur et al., 2025; Braun et al., 2025) to retrieve and filter candidate objects. For instance, to visually ground the entity “Iron Man”, the LLM agent retrieves a set of movie stills and retains only those where visual attributes (e.g., red and gold armor) statistically align with the entity’s textual definition, selecting the top- k most representative instances.

3.1.2 Scene Graph Generation

For a specific query, we generate a scene graph that transforms the raw image into explicit objects along with their spatial relationships and interactions. Given an input image I_q , we employ an MLLM or a specialized scene graph generation model (Im et al., 2024; Li et al., 2025a; Dutta et al., 2025; Kim et al., 2024) to detect a set of objects \mathcal{E}_s and identify the relations \mathcal{R}_s between pairs of objects. Notably, to prevent the scene graph from merely replicating the MLLM’s inherent coarse-grained perception, we can employ optimization strategies such as region-wise scanning (Chen et al., 2024), multi-turn iterative refinement (Khandelwal and Sigal, 2022), and ensemble with specialized models (Im et al., 2024; Ren et al., 2024) to extract dense visual details. Consequently, this process yields the scene graph $\mathcal{G}_s = (\mathcal{E}_s, \mathcal{R}_s)$. Crucially,

the detected objects \mathcal{E}_s serve as both fine-grained visual grounding and structural anchors for external knowledge retrieval.

3.2 Two-stage Retrieval

Since large-scale commonsense graphs typically contain millions of relational facts, efficiently retrieving information relevant to the query is critical. Existing methods (Khan et al., 2025; Lin et al., 2023a; Wan and Yu, 2025) typically rely on text-text or vision-text embedding matching for retrieval, which overlooks vision-centric retrieval and suffers from two limitations: (1) *Textual Bias*: relevant knowledge may be missed if the key visual entities are not explicitly mentioned in the textual query; (2) *Visual Noise*: vision–text matching may incorrectly retrieve semantically related but contextually irrelevant nodes, such as different characters played by the same actor or different items within the same category. To address these issues, we introduce a two-stage retrieval method that comprehensively extracts query-relevant commonsense information from both text and visual modalities.

Stage 1: Text-based Retrieval. The first stage focuses on acquiring a broad semantic context derived from the textual query. Given a commonsense graph $\mathcal{G}_c = (\mathcal{E}_c, \mathcal{R}_c)$ and a query Q , we identify a set of entities \mathcal{E}_{text} from Q . For each entity $t \in \mathcal{E}_{text}$, we utilize the text embedding similarity to find the corresponding node in \mathcal{G}_c and extract

a k -th order subgraph centered on it. This extraction significantly reduces the search space to ensure scalability. The union of these subgraphs forms an initial candidate subgraph $\mathcal{G}_c^k = (\mathcal{E}_c^k, \mathcal{R}_c^k) \subseteq \mathcal{G}_c$.

To further remove irrelevant or weakly related entities, we estimate the importance of each entity $e \in \mathcal{E}_c^k$ by a generalized ranking mechanism, such as Random Walk with Restart (RWR) (Wu et al., 2014) or Personalized PageRank (PPR) (Yang et al., 2024). We then select the top- n ranked entities $\mathcal{E}_{c,T} \subseteq \mathcal{E}_c^k$, together with the relations among them $\mathcal{R}_{c,T} \subseteq \mathcal{R}_c^k$, to form a compact text-guided commonsense subgraph $\mathcal{G}_{c,T} = (\mathcal{E}_{c,T}, \mathcal{R}_{c,T})$.

Stage 2: Vision-based Retrieval. Visual input I_q provides complementary entity cues for VQA that are often absent from the textual query Q . However, naively aligning all entities from the scene graph \mathcal{G}_s derived from I_q with the commonsense graph is inefficient. Owing to the long-tailed distribution of scene graph entities, many detected objects are weakly related or irrelevant to the query, leading to high retrieval cost and noisy knowledge grounding. To overcome this, we propose a coarse-to-fine strategy involving *Query-guided Pruning* followed by *Visual Similarity Retrieval*.

Step A: Query-guided Pruning. We first utilize the text-guided commonsense subgraph $\mathcal{G}_{c,T}$ obtained in Stage 1 to filter visual noise. Specifically, we prompt an MLLM to prune the original scene graph \mathcal{G}_s by removing entities that are irrelevant to the $\mathcal{G}_{c,T}$. The remaining entities $\hat{\mathcal{E}}_s$ and their relations form a query-aware scene subgraph $\hat{\mathcal{G}}_s = (\hat{\mathcal{E}}_s, \hat{\mathcal{R}}_s)$.

This strategy offers significant advantages over traditional query-based filtering. While existing VQA approaches (Lin and Byrne, 2022; Jian et al., 2024) typically identify salient visual entities by conditioning on the textual query Q , such filtering is limited by the sparse semantic scope of the query and may discard entities that are implicitly relevant through relational dependencies. For example, a visual entity not mentioned in the query may still be related to the queried concepts via commonsense relations (e.g., a tool associated with a queried action). By conditioning on $\mathcal{G}_{c,T}$, which encodes a richer set of entities and local relations, we enable more comprehensive entity recognition from redundant visual information.

Moreover, pruning the scene graph with respect to $\mathcal{G}_{c,T}$ facilitates subsequent graph fusion (See Section 3.3). Restricting $\hat{\mathcal{G}}_s$ to a semantic subspace overlapping with the commonsense graph reduces

structural incompatibility and alignment ambiguity, thereby enabling more reliable cross-modal entity alignment and knowledge integration.

Step B: Visual Similarity Retrieval. With the refined visual entities $\hat{\mathcal{E}}_s \in \hat{\mathcal{G}}_s$, we proceed to retrieve complementary knowledge from the commonsense graph \mathcal{G}_c . Prior knowledge-based methods (Chen et al., 2023; Sarwar, 2025) commonly retrieve relevant commonsense graph entities by computing vision-to-text embedding similarity, treating textual descriptions as proxies for visual objects. However, such alignment may overlook subtle visual cues and cause ambiguity when multiple entities share similar textual descriptions.

In contrast, we perform direct vision-to-vision retrieval by matching the visual regions of scene graph entities $e_s \in \hat{\mathcal{E}}_s$ with the visual attributes (i.e., representative images) of candidate entities $e_c \in \mathcal{E}_c$ in the commonsense graph. Specifically, we extract embeddings for e_s and e_c using vision encoders (e.g., Cambrian-1 (Tong et al., 2024), CLIP (Radford et al., 2021)) and compute similarity scores between each pair. A generalized ranking mechanism is then applied to select the top- m most relevant entities $\mathcal{E}_{c,I} \subseteq \mathcal{E}_c$, forming the vision-guided commonsense subgraph $\mathcal{G}_{c,I} = (\mathcal{E}_{c,I}, \mathcal{R}_{c,I})$.

Finally, we take the union of the text-guided subgraph $\mathcal{G}_{c,T}$ and the vision-guided subgraph $\mathcal{G}_{c,I}$ to form a unified query-aware commonsense subgraph: $\tilde{\mathcal{G}}_c = \mathcal{G}_{c,T} \cup \mathcal{G}_{c,I} = (\tilde{\mathcal{E}}_c, \tilde{\mathcal{R}}_c) \subseteq \mathcal{G}_c$, which aggregates complementary knowledge retrieved from both modalities and provides the necessary context for the VQA task.

3.3 Graph Fusion and Refinement

We have obtained two query-aware heterogeneous subgraphs: the scene subgraph $\hat{\mathcal{G}}_s$ that provides fine-grained visual details, and the commonsense subgraph $\tilde{\mathcal{G}}_c$ that offers rich knowledge grounding. Maintaining them as a separate structure risks a semantic disconnect. For example, MLLMs may fail to anchor a generic scene node like "man holding a cup" in $\hat{\mathcal{G}}_s$ to the specific identity of "Detective Cobb" in $\tilde{\mathcal{G}}_c$. Such fragmented representations hinder the synthesis of visual evidence and external knowledge, and may even introduce spurious or misleading information during reasoning.

To address this issue, we fuse the two subgraphs into a unified representation by explicitly aligning nodes that refer to the same real-world entity. Since textual labels in the scene graph typically

describe generic visual attributes (e.g., appearance or common object categories) rather than precise identities, direct text-to-text matching is unreliable. We therefore prioritize visual grounding for cross-graph alignment. Specifically, for each pair $(e_i, e_j) \in \hat{\mathcal{E}}_s \times \hat{\mathcal{E}}_c$, we compute a similarity S_{ij} :

$$S_{ij} = \alpha \langle \phi(I_i), \psi(T_j) \rangle + (1 - \alpha) \langle \phi(I_i), \phi(I_j) \rangle,$$

where ϕ and ψ denote the visual and text encoders, respectively, I_i is the cropped visual region of the scene node e_i ; T_j and I_j denote the textual content and the associated image of the commonsense node e_j . The first term measures the cross-modal alignment between the visual object and the knowledge concept, while the second term captures the visual similarity between the object and the entity’s visual reference. For aligned entity pairs, we merge their attributes to construct a coherent representation: we prioritize visual properties from $\hat{\mathcal{G}}_s$ to ensure strict grounding in the image context, while augmenting them with fine-grained description and labels from $\hat{\mathcal{G}}_c$ to enhance semantic precision. The resulting unified graph \mathcal{G}_f encapsulates both the dynamic visual details and the external knowledge.

To mitigate potential redundancy or missing links in the consolidated graph \mathcal{G}_f , we employ an optional MLLM-driven agent (Asai et al., 2024; Zhai, 2024; Choi et al., 2025) to iteratively refine the topology by evaluating its alignment with Q . At each step, the agent selects one of three actions: (1) **Expand**, to retrieve missing critical nodes that bridge disconnected evidence; (2) **Prune**, to remove noise or irrelevant entities; and (3) **Terminate**, to conclude the process. This refinement operates for a maximum of t steps, ensuring \mathcal{G}_f becomes structurally concise and semantically aligned with the user’s intent before final reasoning. Finally, we input \mathcal{G}_f into the MLLM, together with the original visual input I_q and the textual query Q , to produce more reliable responses.

4 Evaluation

4.1 Benchmarks

We evaluate our framework on two benchmarks: **FVQA 2.0+**. First, we adopt FVQA 2.0 (Lin et al., 2023b), a widely used benchmark for VQA requiring external knowledge. Since its original knowledge base consists solely of textual triplets, we upgrade it to FVQA 2.0+ to fit the multi-modal setting. Following the *Commonsense Graph Construction* pipeline in Section 3.1, we enrich the original entities with visual representations, resulting

in a commonsense graph with 1,152 nodes, 1,767 edges, and 3,342 images. We utilize the original 2,820 QA samples, where each entry comprises a query image, a question requiring external knowledge, and a ground-truth natural language answer. **MVQA**. Since FVQA 2.0 primarily involves simple one-hop reasoning over the knowledge graph, we introduce a new benchmark, MVQA, designed to evaluate complex multi-hop reasoning. Constructed based on 50 movies from MovieBench (Wu et al., 2024), we built a cinematic commonsense graph using the same construction methodology, comprising 1,271 nodes, 1,468 edges, and 4,611 images. Furthermore, we developed an automated generation pipeline (detailed in Appendix A.2) to create 1,433 complex QA samples. Unlike previous datasets, these queries necessitate deep reasoning that bridges fine-grained visual cues in the query image with multi-hop relational paths in the commonsense graph.

4.2 Experimental Setup

Baselines. We compare **KG-ViP** against three categories of baselines: (i) *vanilla MLLM* methods that rely on the intrinsic reasoning ability of MLLMs, including **Zero-shot**, **Fine-tuning**, and **Chain-of-Thought (CoT)** (Wei et al., 2023); (ii) *scene graph-augmented* methods, represented by **CCoT** (Mitra et al., 2024), which converts images into scene graphs to enhance CoT reasoning; (iii) *retrieval-augmented* methods that incorporate external knowledge, including **FilterRAG** (Sarwar, 2025) and **LLM-RA** (Jian et al., 2024). In addition, we implement a RAG baseline using a naive text-based retrieval strategy, termed **NaiveRAG**. Implementation details are provided in Appendix A.1.

Evaluation Metrics. Following prior work (Liu et al., 2024; Xia et al., 2024; Chang et al., 2025), we adopt three widely used metrics for evaluation. METEOR (Banerjee and Lavie, 2005) calculates the harmonic mean of unigram precision and recall, capturing semantic variations while applying a fragmentation penalty for word order. Semantic Answer Similarity (SAS) (Risch et al., 2021) utilizes a cross-encoder to assess the semantic equivalence between predicted and reference answers. LLM-as-a-Judge (Gu et al., 2025) leverages an LLM to assign a scalar score that evaluates the relevance of the generated response with the ground-truth answer. We denote this metric as LLM-J for short.

Implementation Details. We adopt Qwen2.5-VL-7B as the backbone MLLM for entity extraction,

Table 1: **Comparison of different methods on FVQA 2.0+ and MVQA.** Results are reported in terms of LLM-J (LLM-as-a-Judge), METEOR, and SAS.

Method	FVQA 2.0+				MVQA			
	LLM-J \uparrow		METEOR \uparrow	SAS \uparrow	LLM-J \uparrow		METEOR \uparrow	SAS \uparrow
	Qwen2.5-7B	DeepSeek-V3.2			Qwen2.5-7B	DeepSeek-V3.2		
<i>Vanilla MLLM Reasoning</i>								
Zero-shot	36.79	38.90	15.12	32.70	14.41	15.84	7.70	9.18
CoT	39.96	40.35	18.30	36.65	16.09	17.78	8.20	9.69
Fine-Tuning	43.16	43.09	19.82	39.52	16.23	16.85	8.33	11.22
<i>Scene Graph-Augmented Reasoning</i>								
CCoT(Mitra et al., 2024)	43.34	44.32	20.35	40.56	16.41	18.64	8.69	10.23
<i>Retrieval-Augmented Reasoning</i>								
NaiveRAG	46.25	48.09	21.52	43.93	19.47	20.78	9.11	13.17
FilterRAG(Sarwar, 2025)	50.16	51.73	22.46	45.72	20.11	20.84	12.41	15.49
LLM-RA (Jian et al., 2024)	51.90	52.64	22.96	46.82	22.53	24.10	11.43	17.22
KG-ViP (Ours)	59.50	60.40	23.83	50.75	30.46	35.44	18.44	23.25

scene graph generation, graph refinement, and final answer generation. For retrieval, all-MiniLM-L6-v2 and Cambrian-1-8B are employed as the text and visual encoders, respectively. The iteration of graph refinement is fixed to $t = 1$ to balance refinement quality and efficiency. Qwen2.5-7B and DeepSeek-V3.2 (671B) are used as judge models for the LLM-J metric. All experiments are conducted on four NVIDIA A100 GPUs. More details are provided in Appendix A.1

4.3 Main Results

Table 1 presents the main results on FVQA 2.0+ and MVQA. Overall, our proposed KG-ViP consistently achieves the best performance across all evaluation metrics on both benchmarks, showing its strong multi-modal reasoning capability. Specifically, on FVQA 2.0+, KG-ViP improves the LLM-J (DeepSeek-V3.2) by 7.76% and the SAS by 2.93%, compared to the previous leading baseline LLM-RA. The results underscore the importance of incorporating structured representations for visual inputs in VQA. While FilterRAG and LLM-RA also identify key visual entities and retrieve related external knowledge, they largely ignore the spatial relations among visual entities. In contrast, KG-ViP explicitly models entities and their spatial interactions with the scene graph, thereby improving the visual perception and grounding of MLLMs. MVQA poses a more challenging setting that involves complex multi-hop reasoning over multiple entities and relations, demanding precise alignment between visual grounding and external knowledge. In this scenario, KG-ViP exhibits an even larger performance margin, outperforming LLM-RA by 11.34% in the LLM-J (DeepSeek-V3.2) and 6.03% in SAS. We attribute these gains to the unified and

Table 2: Effectiveness of Graph Fusion and Refinement

Method	FVQA 2.0+		MVQA	
	LLM-J \uparrow	METEOR \uparrow	LLM-J \uparrow	METEOR \uparrow
KG-ViP	60.40	23.83	35.44	18.44
w/o Graph Fusion	58.11	23.43	29.80	15.10
Δ	-2.29	-0.40	-5.64	-3.34
w/o Graph Refinement	60.18	23.52	33.19	18.19
Δ	-0.22	-0.31	-2.25	-0.25

Table 3: Impact of Retrieval Mode

Text Retrieval	Vision Retrieval		FVQA 2.0+		MVQA	
	V \rightarrow T	V \rightarrow V	LLM-J \uparrow	METEOR \uparrow	LLM-J \uparrow	METEOR \uparrow
		\checkmark	58.34	22.66	27.56	14.20
\checkmark	\checkmark		58.66	23.36	32.02	15.50
\checkmark		\checkmark	60.40	23.83	35.44	18.44

compact graph representation produced by KG-ViP, which integrates the scene graph and commonsense graph, providing a low-noise and structured context to enhance multi-modal reasoning.




Figure 3: Impact of vision encoder selection on KG-ViP.

4.4 Ablations and Analysis

We perform a comprehensive ablation and analysis study of KG-ViP. For the LLM-J metric, DeepSeek-V3.2 serves as the default judge model.

Effectiveness of Graph Fusion and Refinement.

To validate the effectiveness of graph fusion and refinement, we evaluate two ablated variants: (i) providing the MLLM with isolated scene and commonsense subgraphs without fusion, and (ii) using the fused graph without refinement. As shown in Table 2, removing graph fusion results in substan-

Table 4: Impact of Scene Graph Pruning.

Dataset	Metric	Method		
		KG-ViP w/o pruning	KG-ViP query pruning	KG-ViP graph pruning
FVQA 2.0+	LLM-J \uparrow	60.04	59.77	60.40
	METEOR \uparrow	22.71	23.47	23.83
MVQA	LLM-J \uparrow	30.99	33.39	35.44
	METEOR \uparrow	15.88	17.06	18.44

Table 5: Robustness across MLLM architectures.

MLLM	Method	LLM-J \uparrow	METEOR \uparrow	SAS \uparrow
Qwen2.5-VL 3B-Instruct	Zero-Shot	7.98	3.14	4.38
	KG-ViP	24.39	5.80	12.51
	Δ	+16.41	+2.66	+8.13
Qwen2.5-VL 32B-Instruct	Zero-Shot	13.16	7.41	8.73
	KG-ViP	37.35	18.83	26.96
	Δ	+24.19	+11.42	+18.23
GLM-4.1V 9B-Thinking	Zero-Shot	9.70	3.84	6.04
	KG-ViP	28.54	9.50	19.92
	Δ	+18.84	+5.66	+13.88
GLM-4.5V 106B	Zero-Shot	19.10	10.67	11.70
	KG-ViP	40.58	22.34	31.68
	Δ	+21.48	+11.67	+19.98

tial performance degradation on both benchmarks, with the LLM-J dropping by 5.64% on MVQA and 2.29% on FVQA 2.0+. This highlights the critical role of a unified graph representation in synergizing visual evidence with commonsense knowledge. Furthermore, disabling the graph refinement step also leads to a noticeable decline in performance, demonstrating that pruning redundant nodes is essential for maintaining a high-quality, task-relevant graph structure.

Impact of Retrieval Mode. KG-ViP leverages both textual and visual modalities to extract query-relevant commonsense information. Regarding the visual branch, retrieval can be implemented via either vision-to-text ($V \rightarrow T$) or vision-to-vision ($V \rightarrow V$) matching. KG-ViP specifically adopts $V \rightarrow V$ matching to mitigate the ambiguity inherent in $V \rightarrow T$ approaches, where distinct visual entities may share similar textual descriptions. We evaluate the performance of these retrieval configurations in Table 3. The results indicate that the absence of text-based retrieval leads to consistent performance degradation regardless of the vision-based mode, confirming that both modalities provide indispensable cues. Furthermore, when text-based retrieval is retained, $V \rightarrow V$ matching consistently outperforms $V \rightarrow T$ matching. These findings suggest that direct visual alignment captures subtle visual cues more effectively, leading to more reliable knowledge retrieval and reasoning.

Impact of Scene Graph Pruning. To mitigate vi-

sual noise stemming from the long-tailed distribution of scene graphs, KG-ViP prunes task-irrelevant entities by conditioning the scene graph on the text-guided commonsense subgraph. To validate this strategy, we compare it against Query Pruning, which relies solely on the raw query text for filtering. As shown in Table 4, our proposed graph pruning consistently outperforms both no pruning and Query Pruning. This demonstrates that while pruning the scene is essential for noise reduction, leveraging the enriched entity set from the commonsense subgraph provides a more robust semantic reference than the sparse query text, effectively preserving critical visual evidence for reasoning.

Selection of Vision Encoders. To analyze the impact of visual representation on KG-ViP, we investigate several vision encoders, including CLIP, BLIP, and Cambrian-1, on the MVQA benchmark. As illustrated in Figure 3, Cambrian-1 significantly outperforms both CLIP and BLIP. We attribute this performance gap to the distinct nature of these models: CLIP and BLIP are primarily designed as text-vision alignment models, which prioritize matching global visual features with textual descriptions. In contrast, Cambrian-1 functions as a specialized vision-centric encoder, optimized for capturing fine-grained visual details and spatial structures. These advantages make Cambrian-1 particularly well-suited for the vision-centric retrieval and reasoning scenarios required by KG-ViP.

Robustness across MLLM Architectures. Table 5 presents a comprehensive comparison across MLLMs of varying scales on MVQA, including Qwen2.5-VL-3B/32B and GLM-4.1V-9B/GLM-4.5V-106B. When integrated with our KG-ViP framework, these models exhibit consistent performance gains, with average improvements of 20.23%, 7.85%, and 15.06% across the LLM-J, METEOR, and SAS metrics, respectively. This confirms KG-ViP’s robustness across diverse MLLM architectures and scales.

5 Conclusion

In this paper, we address the dual limitations of MLLMs in knowledge grounding and visual perception for VQA. We identify that these gaps can be effectively bridged by integrating commonsense graphs, which provide external knowledge, with scene graphs, which capture fine-grained visual details. Unlike prior approaches that utilize these graphs in isolation, we propose KG-ViP, a uni-

fied framework that synergizes them through a novel retrieval-and-fusion pipeline, offering a unified structured context for joint reasoning. Extensive experiments show that our method significantly enhances performance, establishing a robust paradigm for reliable multi-modal reasoning.

Limitations

While KG-ViP demonstrates robust performance in multi-modal reasoning, we identify two aspects for future exploration. First, our framework relies on off-the-shelf models for scene graph generation. Consequently, the quality of visual parsing sets an upper bound on downstream reasoning, particularly in challenging scenarios involving severe occlusion or small-scale objects. However, since KG-ViP is model-agnostic, it can directly benefit from future advancements in scene graph generation techniques without structural changes. Second, to ensure online inference efficiency, we currently employ an offline-constructed commonsense graph. While this design choice minimizes latency, applying the framework to domains with evolving information would necessitate periodic graph updates or re-indexing to maintain knowledge currency. Future work could explore dynamic graph mechanisms to address such temporal shifts.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC) under Grants 62472400 and 62271465, the National Key R&D Program of China under Grant 2025YFC3408300, and the Suzhou Basic Research Program under Grant SYG202338.

References

- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. [Self-RAG: Learning to retrieve, generate, and critique through self-reflection](#). In *The Twelfth International Conference on Learning Representations*.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Tobias Braun, Mark Rothemmel, Marcus Rohrbach, and Anna Rohrbach. 2025. [DEFAME: Dynamic Evidence-based FAct-checking with Multimodal Experts](#). In *Proceedings of the 42nd International Conference on Machine Learning*.
- 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](#). Preprint, arXiv:2404.15406.
- Yukun Cao, Zengyi Gao, Zhiyang Li, Xike Xie, S. Kevin Zhou, and Jianliang Xu. 2025. [Lego-graphrag: Modularizing graph-based retrieval-augmented generation for design space exploration](#). *Proc. VLDB Endow.*, 18(10):3269–3283.
- Eun Chang, Zhuangqun Huang, Yiwei Liao, Sagar Ravi Bhavsar, Amogh Param, Tammy Stark, Adel Ahmadyan, Xiao Yang, Jiaqi Wang, Ahsan Abdullah, Giang Nguyen, Akil Iyer, David Hall, Elissa Li, Shane Moon, Nicolas Scheffer, Kirmani Ahmed, Babak Damavandi, Rakesh Wanga, and 3 others. 2025. [Wearvqa: A visual question answering benchmark for wearables in egocentric authentic real-world scenarios](#). Preprint, arXiv:2511.22154.
- Xiaojun Chang, Pengzhen Ren, Pengfei Xu, Zhihui Li, Xiaojiang Chen, and Alex Hauptmann. 2021. A comprehensive survey of scene graphs: Generation and application. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(1):1–26.
- 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, Singapore. Association for Computational Linguistics.
- Zuyao Chen, Jinlin Wu, Zhen Lei, Zhaoxiang Zhang, and Changwen Chen. 2024. [Gpt4sgg: Synthesizing scene graphs from holistic and region-specific narratives](#). Preprint, arXiv:2312.04314.
- Changin Choi, Wonseok Lee, Jungmin Ko, and Wonjong Rhee. 2025. [Multimodal iterative rag for knowledge-intensive visual question answering](#). Preprint, arXiv:2509.00798.
- Amartya Dutta, Kazi Sajeed Mehrab, Medha Sawhney, Abhilash Neog, Mridul Khurana, Sepideh Fatemi, Anish Pradhan, M. Maruf, Ismini Lourentzou, Arka Daw, and Anuj Karpatne. 2025. [Open world scene graph generation using vision language models](#). Preprint, arXiv:2506.08189.
- Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. 2024. [From local to global: A graph rag approach to query-focused summarization](#). arXiv preprint arXiv:2404.16130.

- Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. A survey on rag meeting llms: Towards retrieval-augmented large language models. In *Proceedings of the 30th ACM SIGKDD conference on knowledge discovery and data mining*, pages 6491–6501.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Yuanzhuo Wang, Wen Gao, Lionel Ni, and Jian Guo. 2025. *A survey on llm-as-a-judge*. Preprint, arXiv:2411.15594.
- Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. 2024. Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14375–14385.
- Liangke Gui, Borui Wang, Qiuyuan Huang, Alex Hauptmann, Yonatan Bisk, and Jianfeng Gao. 2022. *Kat: A knowledge augmented transformer for vision-and-language*. Preprint, arXiv:2112.08614.
- Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, and Chao Huang. 2024. *Lightrag: Simple and fast retrieval-augmented generation*.
- Roei Herzig, Alon Mendelson, Leonid Karlinsky, Asaf Arbelle, Rogerio Feris, Trevor Darrell, and Amir Globerson. 2023. *Incorporating structured representations into pretrained vision & language models using scene graphs*. Preprint, arXiv:2305.06343.
- Jinbae Im, JeongYeon Nam, Nokyoung Park, Hyungmin Lee, and Seunghyun Park. 2024. *Egrt: Extracting graph from transformer for scene graph generation*. Preprint, arXiv:2404.02072.
- Pu Jian, Donglei Yu, and Jiajun Zhang. 2024. Large language models know what is key visual entity: An llm-assisted multimodal retrieval for vqa. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 10939–10956.
- Uku Kangur, Krish Agrawal, Yashashvi Singh, Ahmed Sabir, and Rajesh Sharma. 2025. Multireflect: Multimodal self-reflective rag-based automated fact-checking. In *Proceedings of the 1st Workshop on Multimodal Augmented Generation via Multimodal Retrieval (MAGMaR 2025)*, pages 1–17.
- Muhammad Junaid Khan, Adil Masood Siddiqui, Hamid Saeed Khan, and Jaleed Khan. 2025. Enhancing visual question answering with common sense knowledge: a data-driven neurosymbolic graph routing approach. *International Journal of Data Science and Analytics*, pages 1–16.
- Siddhesh Khandelwal and Leonid Sigal. 2022. *Iterative scene graph generation*. Preprint, arXiv:2207.13440.
- Kibum Kim, Kanghoon Yoon, Jaehyeong Jeon, Yeonjun In, Jinyoung Moon, Donghyun Kim, and Chanyoung Park. 2024. Llm4sgg: Large language models for weakly supervised scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 28306–28316.
- Insu Lee, Wooje Park, Jaeyun Jang, Minyoung Noh, Kyuhong Shim, and Byonghyo Shim. 2025. *Towards comprehensive scene understanding: Integrating first and third-person views for lvlms*. Preprint, arXiv:2505.21955.
- Junlin Lee, Yequan Wang, Jing Li, and Min Zhang. 2024. Multimodal reasoning with multimodal knowledge graph. *arXiv preprint arXiv:2406.02030*.
- Hongsheng Li, Guangming Zhu, Liang Zhang, Youliang Jiang, Yixuan Dang, Haoran Hou, Peiyi Shen, Xia Zhao, Syed Afaq Ali Shah, and Mohammed Benamoun. 2024. Scene graph generation: A comprehensive survey. *Neurocomputing*, 566:127052.
- Lin Li, Chuhan Zhang, Dong Zhang, Chong Sun, Chen Li, and Long Chen. 2025a. *Interaction-centric knowledge infusion and transfer for open-vocabulary scene graph generation*. Preprint, arXiv:2511.05935.
- Xinwei Li, Li Lin, Shuai Wang, and Hanqian Wu. 2025b. *Seeing beyond hallucinations: Llm-based compositional information extraction for multimodal reasoning*. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '25*, page 1000–1010, New York, NY, USA. Association for Computing Machinery.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023. *Evaluating object hallucination in large vision-language models*. Preprint, arXiv:2305.10355.
- Fake Lin, Xi Zhu, Ziwei Zhao, Deqiang Huang, Yu Yu, Xueying Li, Zhi Zheng, Tong Xu, and Enhong Chen. 2024. Knowledge graph pruning for recommendation. *arXiv preprint arXiv:2405.11531*.
- Weizhe Lin and Bill Byrne. 2022. *Retrieval augmented visual question answering with outside knowledge*. Preprint, arXiv:2210.03809.
- Weizhe Lin, Jinghong Chen, Jingbiao Mei, Alexandru Coca, and Bill Byrne. 2023a. *Fine-grained late-interaction multi-modal retrieval for retrieval augmented visual question answering*. Preprint, arXiv:2309.17133.
- Weizhe Lin, Zhilin Wang, and Bill Byrne. 2023b. *FVQA 2.0: Introducing adversarial samples into fact-based visual question answering*. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 149–157, Dubrovnik, Croatia. Association for Computational Linguistics.

- Ye Liu, Hui Li, Alberto Garcia-Duran, Mathias Niepert, Daniel Onoro-Rubio, and David S Rosenblum. 2019. Mmkg: multi-modal knowledge graphs. In *European Semantic Web Conference*, pages 459–474. Springer.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2024. *Mmbench: Is your multi-modal model an all-around player?* *Preprint*, arXiv:2307.06281.
- Maria Lomaeva and Nitisha Jain. 2022. Relation canonicalization in open knowledge graphs: a quantitative analysis. In *European Semantic Web Conference*, pages 21–25. Springer.
- Zixian Ma, Jerry Hong, Mustafa Omer Gul, Mona Gandhi, Irena Gao, and Ranjay Krishna. 2023. Crepe: Can vision-language foundation models reason compositionally? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10910–10921.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Chancharik Mitra, Brandon Huang, Trevor Darrell, and Roei Herzig. 2024. Compositional chain of thought prompting for large multimodal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, and 1 others. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmlR.
- Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kun-chang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, Zhaoyang Zeng, Hao Zhang, Feng Li, Jie Yang, Hongyang Li, Qing Jiang, and Lei Zhang. 2024. *Grounded sam: Assembling open-world models for diverse visual tasks.* *Preprint*, arXiv:2401.14159.
- Julian Risch, Timo Möller, Julian Gutsch, and Malte Pietsch. 2021. Semantic answer similarity for evaluating question answering models. *arXiv preprint arXiv:2108.06130*.
- Nobin Sarwar. 2025. Filterrag: zero-shot informed retrieval-augmented generation to mitigate hallucinations in vqa. *arXiv preprint arXiv:2502.18536*.
- Qingyi Si, Yuchen Mo, Zheng Lin, Huishan Ji, and Weiping Wang. 2023. *Combo of thinking and observing for outside-knowledge VQA.* In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10959–10975, Toronto, Canada. Association for Computational Linguistics.
- Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel M. Ni, Heung-Yeung Shum, and Jian Guo. 2024. *Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph.* *Preprint*, arXiv:2307.07697.
- Zequan Sun, Qingheng Zhang, Wei Hu, Chengming Wang, Muhao Chen, Farahnaz Akrami, and Chengkai Li. 2020. A benchmarking study of embedding-based entity alignment for knowledge graphs. *Proceedings of the VLDB Endowment*, 13(11):2326–2340.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. 2024. *Cambrian-1: A fully open, vision-centric exploration of multimodal llms.*
- Xueyao Wan and Hang Yu. 2025. *Mmgraphrag: Bridging vision and language with interpretable multimodal knowledge graphs.* *Preprint*, arXiv:2507.20804.
- Jingyi Wang, Jianzhong Ju, Jian Luan, and Zhidong Deng. 2024. *Llava-sg: Leveraging scene graphs as visual semantic expression in vision-language models.* *Preprint*, arXiv:2408.16224.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. *Chain-of-thought prompting elicits reasoning in large language models.* *Preprint*, arXiv:2201.11903.
- Weijia Wu, Mingyu Liu, Zeyu Zhu, Xi Xia, Haoen Feng, Wen Wang, Kevin Qinghong Lin, Chunhua Shen, and Mike Zheng Shou. 2024. *Moviebench: A hierarchical movie level dataset for long video generation.* *Preprint*, arXiv:2411.15262.
- Yubao Wu, Ruoming Jin, and Xiang Zhang. 2014. *Fast and unified local search for random walk based k-nearest-neighbor query in large graphs.* In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, page 1139–1150, New York, NY, USA. Association for Computing Machinery.
- Peng Xia, Kangyu Zhu, Haoran Li, Tianze Wang, Weijia Shi, Sheng Wang, Linjun Zhang, James Zou, and Huaxiu Yao. 2024. *Mmed-rag: Versatile multimodal rag system for medical vision language models.* *arXiv preprint arXiv:2410.13085*.
- Siheng Xiong, Ali Payani, Yuan Yang, and Faramarz Fekri. 2025. *Deliberate reasoning in language models as structure-aware planning with an accurate world model.* In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 31900–31931, Vienna, Austria. Association for Computational Linguistics.

Dexuan Xu, Yanyuan Chen, Jieyi Wang, Yue Huang, Hanpin Wang, Zhi Jin, Hongxing Wang, Weihua Yue, Jing He, Hang Li, and Yu Huang. 2024. [MLeVLM: Improve multi-level progressive capabilities based on multimodal large language model for medical visual question answering](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 4977–4997, Bangkok, Thailand. Association for Computational Linguistics.

Yibin Yan and Weidi Xie. 2024. [EchoSight: Advancing visual-language models with Wiki knowledge](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1538–1551, Miami, Florida, USA. Association for Computational Linguistics.

Dongil Yang, Minjin Kim, Sunghwan Kim, Beong woo Kwak, Minjun Park, Jinseok Hong, Woontack Woo, and Jinyoung Yeo. 2025. [Llm meets scene graph: Can large language models understand and generate scene graphs? a benchmark and empirical study](#). *Preprint*, arXiv:2505.19510.

Mingji Yang, Hanzhi Wang, Zhewei Wei, Sibow Wang, and Ji-Rong Wen. 2024. [Efficient algorithms for personalized pagerank computation: A survey](#). *IEEE Transactions on Knowledge and Data Engineering*, 36(9):4582–4602.

Xu Yuan, Liangbo Ning, Wenqi Fan, and Qing Li. 2025. [mkg-rag: Multimodal knowledge graph-enhanced rag for visual question answering](#). *Preprint*, arXiv:2508.05318.

Wenjia Zhai. 2024. [Self-adaptive multimodal retrieval-augmented generation](#). *Preprint*, arXiv:2410.11321.

Bowen Zhang and Harold Soh. 2024. [Extract, define, canonicalize: An llm-based framework for knowledge graph construction](#). *arXiv preprint arXiv:2404.03868*.

Xiangru Zhu, Zhixu Li, Xiaodan Wang, Xueyao Jiang, Penglei Sun, Xuwu Wang, Yanghua Xiao, and Nicholas Jing Yuan. 2022. [Multi-modal knowledge graph construction and application: A survey](#). *IEEE Transactions on Knowledge and Data Engineering*, 36(2):715–735.

A Appendix

A.1 Experiment Implementation Details

Hardware & Environment. All experiments were conducted on a server running Ubuntu 20.04.6 LTS, equipped with an Intel(R) Xeon(R) Platinum 8358 CPU @ 2.60GHz, 400GB of RAM, and four NVIDIA A100-80G GPUs. We implement our framework using PyTorch. To enable efficient similarity search, we utilize the FAISS library for vector indexing and retrieval.

Model Configuration. Table 6 details the architectures of the models employed in this study. For inference, all Large Language Models (LLMs) are served via the vLLM library to ensure high throughput.

Table 6: Main Models Used in This Paper

Model Name	Base Architecture
all-MiniLM-L6-v2	BERT (Transformer Encoder)
Cambrian-1-8b	Llama-3
clip-vit-base-patch32	ViT-B/32
blip-itm-base-coco	ViT-B + BERT
Qwen2.5-7B-Instruct	Transformer Decoder
Qwen2.5-VL-3B-Instruct	Qwen2.5 + ViT
Qwen2.5-VL-7B-Instruct	Qwen2.5 + ViT
Qwen2.5-VL-32B-Instruct	Qwen2.5 + ViT
GLM-4.1V-9B-Thinking	GLM-4 + ViT
GLM-4.5V-106B	GLM-4.5 (MoE) + ViT

Baseline Settings. Detailed configurations for different experimental settings are as follows:

- **Zero-Shot:** We directly input the query and the image into the model to assess its inherent reasoning capabilities without external augmentation.
- **Chain of Thought (CoT):** We append the prompt “Let’s think step by step” to the input to induce stepwise reasoning from the MLLM.
- **Fine-Tuning:** We employ the LLaMA-Factory framework to incorporate domain knowledge from the commonsense graph using the LoRA (Low-Rank Adaptation) method. The configuration includes a rank of 8, a scaling factor (α) of 16, and a dropout rate of 0 applied to all linear layers. The training employs the AdamW optimizer with a learning rate of 5×10^{-5} (cosine decay), gradient clipping at a maximum norm of 1.0, and BF16 mixed precision. To handle long contexts, we set the cutoff length to 2048, utilizing a per-device batch size of 2 with 8 steps of gradient accumulation.
- **Naive RAG:** We segment the reference text into chunks of 1024 tokens and utilize the all-MiniLM-L6-v2 model for embedding generation. During retrieval, the top-3 chunks ($k = 3$) are selected for context augmentation.

A.2 Benchmark Details

FVQA2.0+ This dataset is built upon FVQA 2.0, an extension of the original FVQA dataset designed for the Fact-based Visual Question Answering task. FVQA 2.0 tests a system’s ability

to answer visually related questions by leveraging external knowledge graphs, effectively addressing limitations such as small scale, imbalanced answer distribution, and overfitting. It comprises 2,820 QA pairs. On this basis, we propose **FVQA 2.0+**. We utilized an LLM-driven web search agent to perform visual modality augmentation on the textual knowledge graph of FVQA 2.0, constructing a Multimodal Knowledge Graph (MMKG) containing 1,152 nodes, 1,767 edges, and 3,342 images.

MVQA This benchmark is constructed based on MovieBench, a hierarchical dataset including high-level plot summaries and shot-level visual descriptions of classic movies. We selected 50 movies to construct an MMKG comprising 1,271 nodes, 1,468 edges, and 4,611 images. Furthermore, we employed an automated question generation pipeline to create 1,433 QA pairs centered on movie plots.

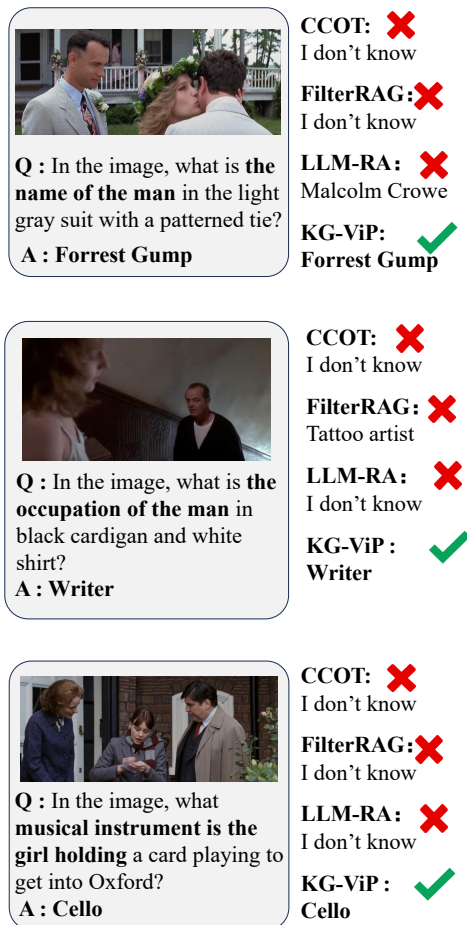


Figure 4: Qualitative comparison of VQA results. The visualization demonstrates the reasoning process and answer generation across different methods.

A.3 Case Study

QA Generation Pipeline To alleviate the scarcity of manually annotated VQA pairs, we employ MLLMs to synthesize QA samples via prompt engineering. Specifically, the generation process is conditioned primarily on scene images and their textual descriptions to identify visible entities, interactions, and event cues. Concurrently, we incorporate an accessible commonsense graph as an external source of background knowledge to constrain the scope of knowledge dependency and enhance factual diversity (e.g., character identities, relationships, and plot-related common sense). In implementation, we retrieve local knowledge fragments associated with candidate entities from the graph and provide them as auxiliary context to the MLLM. This encourages the generation of questions that require reasoning over structured background knowledge beyond mere visually discernible content. Detailed prompt templates are provided in Appendix A.4.

Figure 4 illustrates a qualitative comparison between our method and baseline approaches, including LLM-RA and FilterRAG. It can be observed that the latter are prone to generating “hallucinated” responses that appear plausible but are factually incorrect or ambiguous. In contrast, our proposed method consistently handles complex queries with high precision, particularly in scenarios requiring deep-level reasoning and multi-step knowledge integration.

A.4 Prompt Templates

Here we present the key prompts used in our framework.

Scene Graph Generation

You are a visual grounding + relation extraction system. Give you an Image. Your task is to extract key objects and relationships from this Image. Image size: width={w}px, height={h}px. Return ONLY one valid JSON object.

Rules:

- Use bbox_px: [x1, y1, x2, y2] in pixels.
- 0 ≤ x1 < x2 ≤ {w}, 0 ≤ y1 < y2 ≤ {h}.
- confidence in [0, 1]. Omit uncertain ones.

Schema:

```
{
  "entities": [
    {
      "entity_id": "e1",
      "name": "person",
      "category": "person|object|animal|other",
      "bbox_px": [10, 20, 100, 200],
      "confidence": 0.90
    }
  ]
}
```

```

    }
  ],
  "relations": [
    {
      "subject_id": "e1",
      "predicate": "next_to",
      "object_id": "e2",
      "confidence": 0.70,
      "evidence": "Two persons are adjacent."
    }
  ]
}

```

Query-guided Pruning

You are selecting which image-graph entities/relations should be kept for further commonsense graph retrieval. You can use the information in a text-guided commonsense subgraph.

Text-guided commonsense subgraph: {} **Scene graph entities:** {} **Scene graph relations:** {}

Task: Decide which scene graph entities and scene graph relations are relevant to the text-guided commonsense subgraph. Return ONLY valid JSON with schema:

Schema:

```

{
  "keep_entities": [
    "unique_name1", ...
  ],
  "keep_relations": [
    {
      "s": "unique_name1",
      "p": "predicate",
      "o": "unique_name2"
    }
  ]
}

```

LLM-as-a-Judge

You are an answer relevance evaluator. Given a reference answer and a model answer, your task is to output a single float score between 0 and 1 based on the following criteria:

1.0: Fully correct and complete semantic match.
 0.8~0.9: Mostly correct, minor omissions or extras.
 0.6~0.7: Correct core meaning but brief or incomplete.
 0.3~0.5: Partially correct, captures only some aspects.
 0.0~0.2: Incorrect or irrelevant.

Reference Answer: reference

Model Answer: prediction

Now output a single float score between 0 and 1, noting else!

Graph Refinement

You are a knowledge graph analysis expert. Please analyze whether the current subgraph adequately answers the query question and provide optimization suggestions.

Query: "{}"

Query Image: {}

Current Subgraph Information:

- Total Nodes: {}

- Total Edges: {}

Please analyze and provide optimization suggestions, focusing on:

- Does the subgraph adequately cover the key aspects of the query?
- Which nodes are core nodes that need to be further searched, and which may be noise nodes that need to be deleted?

Please respond strictly in the following JSON format, without any additional content:

Schema:

```

{
  "analysis": "Detailed analysis...",
  "recommendation": "Expand|Prune|Terminate",
  "nodes_to_expand": [],
  "nodes_to_prune": [],
  "reason": "Decision rationale",
  "confidence": 0.8
}

```

VQA Data Generator

You are a Visual Question Answering (VQA) data generator.

Inputs:

- (1) ONE scene image (the question image)
- (2) Scene description (characters, plot)
- (3) Reference facts (entities, attributes)

Task: Generate one question grounded in the scene image and the scene description. The question should not be fully answerable from the Image alone, and is expected to benefit from the provided reference facts (retrieved from a commonsense graph).

Instructions: Look at the scene image first and choose a salient character (prefer one that is clearly visible).

Guidelines (Person-centric questions): Try to involve at least one aspect that is typically not directly visible in the Image, such as name/identity, actor/portrays, occupation/role, condition/suffers from, ownership, or other personal attributes mentioned in the reference facts.

Schema:

```

{
  "qa_pairs": [
    {
      "Question": "...",
      "Answer": "..."
    }
  ]
}

```

Scene Image: {}

Scene description: {}

Reference facts: {}