

# Quantifying and Mitigating Unimodal Biases in Multimodal Large Language Models: A Causal Perspective

Meiqi Chen\*

Peking University

meiqichen@stu.pku.edu.cn

Yan Zhang†

Peking University

zhyzhy001@pku.edu.cn

Yixin Cao

Singapore Management University

caoyixin2011@gmail.com

Chaochao Lu†

Shanghai AI Laboratory

luchaochao@pjlab.org.cn

## Abstract

Recent advancements in Large Language Models (LLMs) have facilitated the development of Multimodal LLMs (MLLMs). Despite their impressive capabilities, MLLMs often suffer from an over-reliance on unimodal biases (e.g., language bias and vision bias), leading to incorrect answers in complex multimodal tasks. To investigate this issue, we propose a causal framework to interpret the biases in Visual Question Answering (VQA) problems. Within our framework, we devise a causal graph to elucidate the predictions of MLLMs on VQA problems, and assess the causal effect of biases through an in-depth causal analysis. Motivated by the causal graph, we introduce a novel MORE dataset, consisting of 12,000 VQA instances. This dataset is designed to challenge MLLMs’ abilities, necessitating multi-hop reasoning and the surmounting of unimodal biases. Furthermore, we propose two strategies to mitigate unimodal biases and enhance MLLMs’ reasoning capabilities, including a Decompose-Verify-Answer (DeVA) framework for limited-access MLLMs and the refinement of open-source MLLMs through fine-tuning. Extensive quantitative and qualitative experiments offer valuable insights for future research. Our project page is at <https://opencausalab.github.io/MORE>.

## 1 Introduction

Following the success of Large Language Models (LLMs) (Ouyang et al., 2022; Touvron et al., 2023b), Multimodal LLMs (MLLMs) (OpenAI, 2023; Team et al., 2023) have been proposed for various vision-language tasks (Fu et al., 2023; Liu et al., 2023b). Despite their promising results, it remains unclear if they truly understand images and text in the context of multi-modal reasoning.

\*This work was done during her internship at Shanghai AI Laboratory.

†Corresponding author.



Figure 1: Examples of over-reliance on unimodal biases. MLLMs (e.g., LLaVA) erroneously generate answers due to language bias (indicated by the underlined text below the left image) and vision bias (the right image).

As shown in the knowledge-based Visual Question Answering (VQA) problems of Figure 1, when prompted with “Which country is hosting the next World Cup after this venue?” MLLMs such as GPT-4V (OpenAI, 2023) and LLaVA (Liu et al., 2023a) may capture the language bias of “the next World Cup” and think that the next World Cup will be “the 2022 FIFA World Cup held in Qatar” (which is also outdated knowledge), while ignoring the exact venue presented in the image. Similarly, when presented with an image of “The Shard” in London, MLLM directly identifies “The representative building is The Shard” influenced by vision bias, overlooking the specific constraint “in Berlin” mentioned in the question. These inherent issues pose significant challenges to the reasoning capabilities of MLLMs, particularly when faced with more complex questions.

In order to investigate the issue of MLLMs’ Over-REliance (MORE) on such unimodal biases, we propose a causal framework to interpret and quantify language and vision biases. Specifically, we first define the causal graph of MLLM’s prediction on VQA problems. The causal graph is constructed based on a variety of causal factors

Datasets	Knowledge-based	Multi-hop Reasoning	Answer Type	Unimodal Biases Evaluation	Rationale	# Size
Visual7W (Zhu et al., 2016)	✗	✗	Open-ended	✗	✗	327.9K
VQA (v2) (Goyal et al., 2017)	✗	✗	Open-ended	✗	✗	1.1M
FVQA (Wang et al., 2017)	✓	✗	Open-ended	✗	✓	5.8K
OKVQA (Marino et al., 2019)	✓	✗	Open-ended	✗	✗	14K
S3VQA (Jain et al., 2021)	✓	✗	Open-ended	✗	✗	7.5K
A-OKVQA (Schwenk et al., 2022)	✓	✗	Multi-choice	✗	✓	23.7K
INFOSEEK (Chen et al., 2023)	✓	✗	Open-ended	✗	✗	1.4M
<b>MORE (Ours)</b>	✓	✓	Multi-choice	✓	✓	12K

Table 1: Comparison of MORE with other VQA datasets, highlighting its incorporation of external knowledge, multi-hop reasoning, unimodal bias evaluation, and rationale for interpretability.

that are integral to the prediction process, such as image and question text. Then, we identify a set of interventions in the context of VQA problems, thereby ascertaining the causal effect of unimodal biases on the prediction capabilities of MLLMs via *do*-calculus (Pearl, 1995). By quantifying such causal effects, we can evaluate MLLM’s sensitivity and robustness against unimodal biases.

Based on the above causal analysis, we curate a novel dataset termed MORE, comprising 12,000 VQA instances. This dataset advances existing VQA datasets by introducing a dedicated evaluation of unimodal biases. To facilitate the evaluation, we adopt a Multiple Choice Question (MCQ) format, where each instance consists of an image, a question, and four candidate options. The image is sourced from an existing VQA dataset (Chen et al., 2023). For question and option curation, we incorporate a knowledge graph (KG) (Wang et al., 2021), allowing us to better simulate MLLMs to navigate the corresponding spurious paths within the causal graph. Specifically, the options consist of one correct answer, and three distractors targeting language bias, vision bias, and multi-hop reasoning, respectively. We also provide the reasoning path, designated as *causal rationale*, in the KG for each instance, offering interpretability for evaluation. As summarized in Table 1, compared to existing VQA datasets, MORE features external knowledge, multi-hop reasoning, unimodal bias evaluation, and rationale, showing better comprehensiveness. Experimental results on six leading MLLMs reveal that: 1) most MLLMs perform much poorly on MORE, showing an obvious tendency to rely on unimodal biases. 2) MLLMs still struggle to achieve precise semantic understanding when dealing with multimodal reasoning.

Furthermore, to mitigate the unimodal biases and enhance MLLM’s reasoning abilities, we also

propose two strategies. The basic idea is to leverage step-by-step multimodal reasoning to mitigate spurious paths. The first strategy involves prompt engineering, featuring limited-access MLLMs (e.g., GPT-4V (OpenAI, 2023)) through a **Decompose-Verify-Answer** (DeVA) framework to explicitly guide MLLMs to extract information from multiple modals comprehensively. The second strategy requires fine-tuning open-source MLLMs such as LLaVA (Liu et al., 2023a) based on our dataset, instructing them with the generated causal rationale.

Overall, our main contributions are as follows:

- We propose a causal framework with a well-defined causal graph to interpret and quantify the biases in VQA problems.
- We construct a new MORE dataset, designed to challenge MLLMs to overcome reliance on unimodal biases via multi-hop causal reasoning, demonstrating superior comprehensiveness compared to existing VQA datasets.
- We propose two strategies to mitigate the unimodal biases and enhance the reasoning ability of current MLLMs, including prompting and fine-tuning ways.

## 2 A Causal Framework

In this section, inspired by Stolfo et al. (2022), we first introduce the causal graph of MLLM’s prediction on VQA problems. Then, we leverage the causal graph to elucidate the biases inherent in VQA, especially vision and language biases. Finally, we assess the causal effects of such biases on MLLMs’ prediction by performing controlled interventions (Pearl, 1995).

### 2.1 Problem Setup

We consider an entity-centric VQA problem, denoted by  $M$ , consisting of a question  $Q$  and an image  $I$ . The image depicts a specific entity and

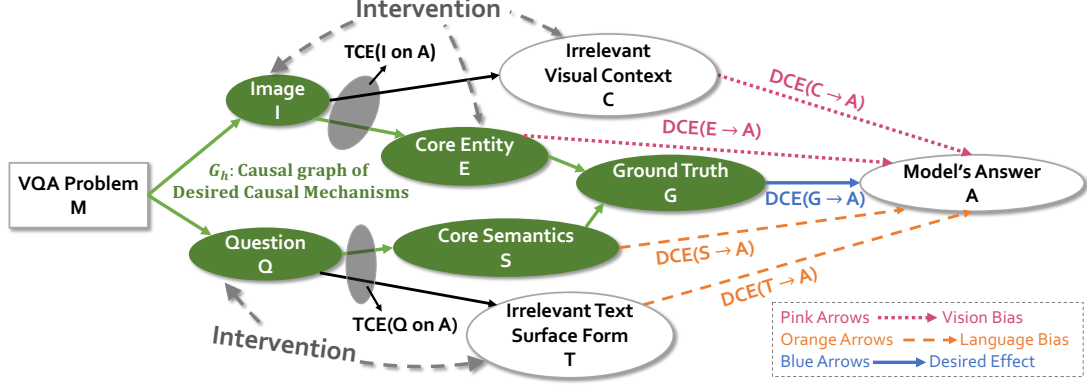


Figure 2: Causal graph of MLLM’s Prediction on VQA problems. We use the green subgraph  $G_h$  to represent the desired causal mechanisms and compare it with the undesired effects of unimodal biases. We quantify the causal effects of each factor by performing controlled interventions of the images ( $I$ ,  $E$ ,  $C$ ) and of the questions ( $Q$ ,  $T$ ).

the question is related to the entity. The question  $Q := (S, T)$  is composed of two distinct elements: the core semantic content  $S$ , which conveys the genuine intent of the question, and the textual surface form  $T$ , which is irrelevant to the question’s core meaning. The image  $I := (E, C)$ , instead, contains the core entity of interest  $E$  and the irrelevant visual context  $C$ . The model’s final answer/prediction is represented by  $A$ . In this paper, we use lowercase letters to represent an instance of its corresponding capital variable.

## 2.2 Causal Graph of MLLMs’ Prediction

Inspired by the intuitive reasoning mechanisms observed in human cognition (Zellers et al., 2019; Fei et al., 2022; Stolfo et al., 2022), we formulate the causal mechanisms underlying human problem-solving in a VQA problem  $m$ :

$$s = f_{c_1}(q), e = f_{c_2}(i), g = f_{c_3}(s, e), \quad (1)$$

where the cognitive process  $f_{c_1}$  is employed to decipher the core semantic meaning  $s$  of the question  $q$ , and  $f_{c_2}$  is employed to extract the core entity  $e$  from the image  $i$ . Then, the function  $f_{c_3}$  is utilized to combine  $s$  and  $e$ , yielding the final result  $g$ . These mechanisms are depicted within the green subgraph  $G_h$  of Figure 2.

In contrast, the process of a model to tackle the same VQA problem  $m$  involves the following causal mechanisms:

$$a = f_b(q, i), \quad (2)$$

where  $a$  represents the model’s prediction. The function  $f_b$  operates as a black box, leaving uncertainty regarding which aspects of  $q$  and  $i$  are considered by the model and how they interact with each other.

To further analyze how a model tackles the VQA problem, we draw all possible causal mechanisms that might take place in the complete causal graph in Figure 2. Notable causal mechanisms include:

**Language Bias** The model may directly process the question  $Q$  in two ways: by focusing on the core semantics  $S$  via the causal path  $Q \rightarrow S \rightarrow A$ , or on the irrelevant part  $T$  via the causal path  $Q \rightarrow T \rightarrow A$ . Both pathways lead to *language bias*, e.g., the focus on “the next World Cup” in Figure 1 (a).

**Vision Bias** The model may attend directly to the entity  $E$  of the image  $I$  via  $I \rightarrow E \rightarrow A$ , or to the irrelevant part  $C$  via  $I \rightarrow C \rightarrow A$ . Both pathways lead to the emergence of *vision bias*, e.g., the focus on the “The Shard” entity in Figure 1 (b).

**Desired Causal Mechanisms** The essence of correct reasoning lies in the model’s grasp of the underlying causal mechanisms necessary for solving a VQA problem. As illustrated in Figure 2), the green subgraph  $G_h$  embodies this comprehension. It implies an understanding of how both the image and the question jointly influence the ground-truth result  $G$  (via  $E \rightarrow G$  and  $S \rightarrow G$ ). Consequently, the model’s predictions should exhibit sensitivity and robustness to changes in the ground truth (Stolfo et al., 2022), namely  $G \rightarrow A$ . No spurious correlations should directly affect  $A$  unless it goes through the mediator  $G$ . Based on the above analyses, we articulate the concept of model sensitivity and robustness on VQA problems:

- *Sensitivity*, which assesses if the model adjusts its predictions properly when the correct answer changes, i.e.,  $A$  responds to changes in  $G$ .
- *Robustness*, which assesses the undesired direct

causal effect of unimodal biases, e.g.,  $C \rightarrow A$  and  $T \rightarrow A$ , where lower effects indicate greater robustness to input changes that do not alter the correct answer.

### 2.3 Causal Analysis of VQA Biases

Once the desired causal mechanism is formulated and the path of unimodal biases is defined, we are able to quantify the causal effects of each causal factor on another by performing controlled interventions (Pearl, 1995).

**Causal Interventions** In the context of VQA, we adopt the following interventions to quantify the causal effects of images and questions on the model’s predictions:

- 1). Interventions on  $Q$ . The question  $Q$  can be modified in two ways: (i) altering both  $S$  and  $T$ , or (ii) altering  $T$  but keeping  $S$  unaffected.
- 2). Interventions on  $I$ . The image  $I$  can be replaced with an alternative image  $I'$  in three ways: (i) altering both  $E$  and  $C$ , or (ii) altering  $C$  but keeping  $E$  unaffected, or (iii) altering  $E$  but keeping  $C$  unaffected.

Note that we do not solely alter  $S$  within  $Q$ , because it is not feasible to intervene on the core semantics  $S$  of a question without affecting the surface text  $T$  of it.

**Formulation of Causal Effects** Next, we explain how we can obtain the causal effect from the interventions. Consider an intervention formulated as  $\text{do}(X : x \rightarrow x')$ , where  $X \in \{I, E, C, Q, T\}$  is a factor in the context of the VQA problem defined by  $M = \{I, Q\}$ . For simplicity, the probability distribution prior to the intervention,  $\mathbb{P}(A | I, Q)$ , is denoted by  $P$ , while the distribution following the intervention is represented as  $P'$ . Following Pearl (1995), we measure the causal effect of each factor in our causal graph by employing the distance metric  $\delta$  to gauge the divergence between distributions  $P$  and  $P'$ :

$$\text{CE} = \delta(P, P'). \quad (3)$$

It is important to note that CE denotes the causal effect, which can refer to either of the two forms: the total causal effect (TCE), signifying the joint effect across all causal paths from one variable to another; and the direct causal effect (DCE), indicating the effect of the directed causal path devoid of intermediary variables (Pearl, 2022).

Following Stolfo et al. (2022), we quantify the causal effect of factor  $X$  on the model’s answer  $A$  by assessing the change in the predicted result:

$$\delta_{\text{cp}}(P, P') := \mathbb{I}(a \neq a') \quad (4)$$

where  $a = \arg \max_x P(x)$ , and  $a' = \arg \max_x P'(x)$ ,  $\mathbb{I}$  denotes an indicator of the “change answer” event.

**Causal Effects of the Questions** For the total causal effect of the question on the answer, we can quantify it by intervening on  $Q$ :

$$\text{TCE}(Q \text{ on } A) := \mathbb{E}_{q' \sim \mathbb{P}(Q)} [\delta(P, P')], \quad (5)$$

where  $P' = \mathbb{P}(A | I, \text{do}(Q = q'))$ .

Note that, this TCE contains two different types of paths that show how  $Q$  affects  $A$ , as illustrated in Figure 2:

- The desired decision-making route we aim for the model to adopt (the  $Q \rightarrow S \rightarrow G \rightarrow A$  path), where the model responds accordingly to alterations in the ground truth answer.
- The spurious correlation that the model might have learned (e.g., the  $Q \rightarrow T \rightarrow A$  path, where the model depends on certain linguistic contexts potentially linked to their prevalence in the training corpus).

Instead of changing the core semantic meaning, maintaining  $S$  constant will enable us to derive the DCE of the textual surface form  $T$  on  $A$ :

$$\text{DCE}(T \rightarrow A) := \mathbb{E}_{q' \sim \mathbb{P}(Q|S)} [\delta(P, P')], \quad (6)$$

where  $P' = \mathbb{P}(A | I, \text{do}(Q = q'))$ .

It should be noted that the  $\text{DCE}(T \rightarrow A)$  equates to the TCE of  $T$  on  $A$ , due to there is no mediator between  $T$  and  $A$ . Besides, as discussed in “Causal Interventions”, solely intervening on  $S$  without influencing the textual surface  $T$  is unattainable. However, comprehending the causal impact of  $S$  on  $A$  becomes feasible through the comparison of two known quantities:  $\text{TCE}(Q \text{ on } A)$  and  $\text{DCE}(T \rightarrow A)$ .

**Causal Effects of the Images** It is evident from Figure 2 that the causal structure of images is exactly the same as that of questions. Therefore, the total causal effect of the image  $I$  on the answer can be directly obtained in a similar manner, i.e.,

$$\text{TCE}(I \text{ on } A) := \mathbb{E}_{i' \sim \mathbb{P}(I)} [\delta(P, P')], \quad (7)$$

where  $P' = \mathbb{P}(A | Q, \text{do}(I = i'))$ .



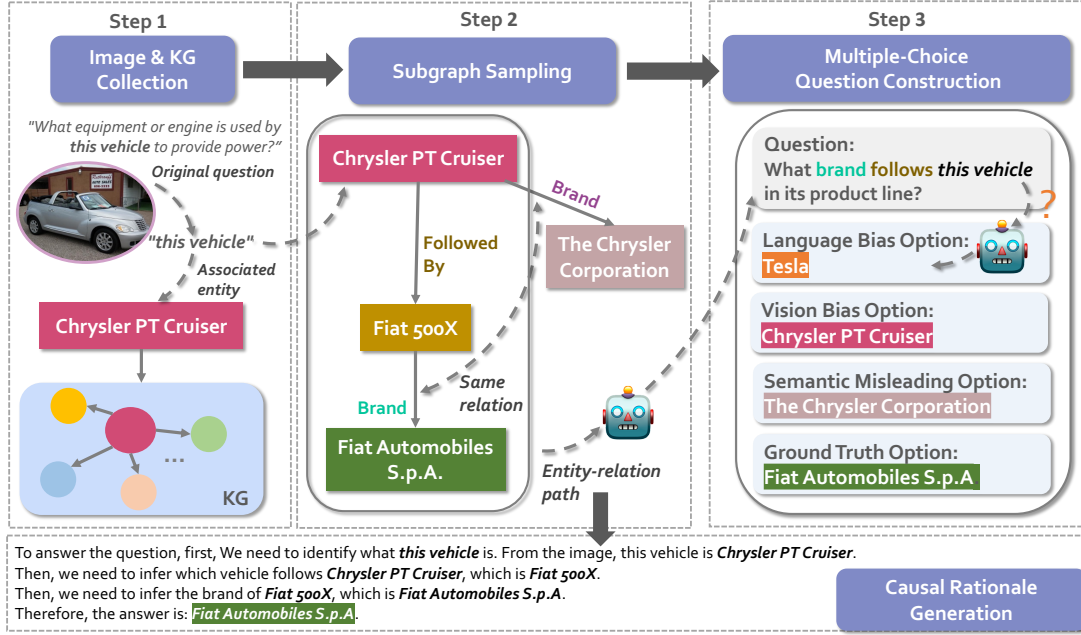


Figure 3: Our framework for generating data of MORE. We first prepare the image source and link the visual entity in a knowledge graph. Then, motivated by the visual and language bias analysis through the causal lens, we construct multiple-choice questions that require MLLMs to overcome unimodal biases and conduct multi-hop reasoning in a sampled subgraph. We also generate the causal (reasoning) rationale for each instance to provide interpretability.

Likewise, maintaining  $E$  constant during each intervention on  $I$  allows us to quantify the DCE of the irrelevant visual context  $C$  on  $A$ , which refers to the strength of the path  $C \rightarrow A$ :

$$\text{DCE}(C \rightarrow A) := \mathbb{E}_{i' \sim \mathbb{P}(I|E)} [\delta(P, P')], \quad (8)$$

where  $P' = \mathbb{P}(A | Q, \text{do}(I = i'))$ .

Maintaining  $C$  and  $G$  constant during each intervention on  $I$  instead allows us to quantify  $\text{DCE}(E \rightarrow A)$ :

$$\text{DCE}(E \rightarrow A) := \mathbb{E}_{i' \sim \mathbb{P}(I|C,G)} [\delta(P, P')], \quad (9)$$

where  $P' = \mathbb{P}(A | Q, \text{do}(I = i'))$ .

Note that in practical situations, it is significantly challenging to intervene on the core entity  $E$  while simultaneously controlling for the irrelevant visual context  $C$ , the ground truth answer  $G$ , and the question  $Q$ . Therefore, this type of intervention will not be considered in subsequent experiments.

Overall, calculating TCE helps us assess a model’s sensitivity (response to changes in ground truth), while DCE evaluates its robustness (stability of predictions against spurious correlations with fixed ground truth).

### 3 Constructing the MORE Dataset

Based on the causal analysis in Section 2, we construct a novel MORE dataset that necessitates

MLLMs to transcend unimodal biases and thoroughly integrate information from both text and images to select the correct answer. To ensure quality, we employ various evaluations to analyze the quality of the generated data. The data generation process is illustrated in Figure 3.

#### 3.1 Images and Knowledge Graph Collection

We begin with an existing VQA dataset, INFOS-EEK (Chen et al., 2023), which links entities depicted in images to information sourced from Wikipedia, requiring a VQA model to answer a question about the associated entity. For instance in Figure 3, the entity associated with the show-cased vehicle is “Chrysler PT Cruiser”, with the question being “What equipment or engine is used by this vehicle to provide power?” In this example, terms such as “Chrysler PT Cruiser”, “this vehicle”, and the relevant region in the image all refer to the same entity.

Based on the images and corresponding entity information, we identify all  $n$ -order neighbors ( $n \in \{1, 2\}$ ) related to the associated entity in a knowledge graph (KG) - Wikidata5M (Wang et al., 2021), which is built upon Wikipedia data.

#### 3.2 Subgraph Sampling

Motivated by the causal analysis discussed in Section 2, we aim to construct multi-hop queries

that require overcoming unimodal biases to be answered correctly. To this end, we first identify a subgraph of an entity and its  $n$ -order neighbors in the KG. We then filter paths in this subgraph that meet two criteria: 1) *Uniqueness of Paths*: the path from the associated entity to the selected neighbor is unique, and 2) *Shared-Type Relations*: they share a same-type relation pointing to a unique entity, and these two pointed entities are not the same. The first criterion guarantees the uniqueness of the correct answer, while the second introduces interference and challenges MLLM’s reasoning ability.

For instance, in the subgraph of Figure 3, “*Fiat 500X*” is a first-order neighbor of “*Chrysler PT Cruiser*”, uniquely linked by the “*followed by*” relation. These two entities are linked to different entities “*The Chrysler Corporation*” and “*Fiat Automobiles S.p.A*” through the same type of relations “*brand*”. Meeting these criteria, we form the multi-hop query: “*Chrysler PT Cruiser*  $\xrightarrow{\text{followed by}}$  *Fiat 500X*  $\xrightarrow{\text{brand}}$  *Fiat Automobiles S.p.A*”, with “*Chrysler PT Cruiser*” as the start and “*Fiat Automobiles S.p.A*” as the answer. In this way, other entities of the filtered subgraph may serve as distractors to challenge MLLMs’ reasoning ability, which will be discussed in the next section.

### 3.3 Multiple-Choice Question Construction

In this subsection, we detail the process of constructing multiple-choice questions with four candidate options.

**Question Generation** After obtaining the subgraph that meets the criteria, we generate questions using the entity-relation path in it. Specifically, to obtain fluent and coherent questions, we feed the path into an LLM to produce the target question text. Our inputs include the names of start entity and all relations within the path. We adopt the in-context learning (ICL) technique (Brown et al., 2020) and provide several examples to LLMs. The prompt template is in Appendix A.1. Comparing various LLMs with adjusted instructions, we find ChatGPT generates the highest-quality multi-hop questions, hence choosing its results for subsequent human evaluation. Finally, to prevent information leakage, entity names in the question are replaced with “this <OBJECT\_NAME>”.

**Language Bias Option** As discussed in Section 2.2, language bias refers to the model overly attending to the information from questions through

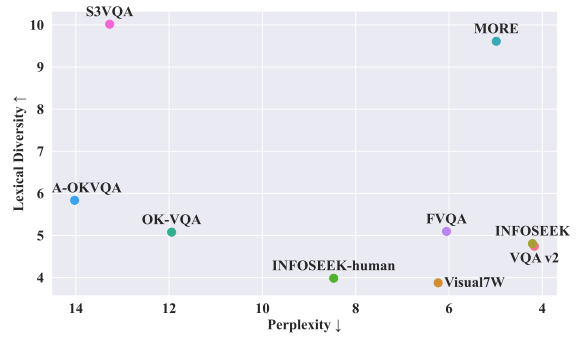


Figure 4: Question quality of MORE compared to other VQA datasets in terms of lexical diversity and fluency.

the undesired path  $Q \rightarrow S \rightarrow A$  or  $Q \rightarrow T \rightarrow A$ . To simulate such a situation, we query an MLLM with the generated question in a *text-only* setup. To ensure that the final options are identical for all MLLMs, we uniformly use the answers obtained from querying GPT-4V. The output for the mentioned example in Figure 3 is “*Tesla*”. The prompt template we use can be seen in Appendix A.2.

**Vision Bias Option** To explore vision bias along the path  $I \rightarrow A$ , we include the visually associated entity name (e.g., “*Chrysler PT Cruiser*”) as an option. This allows us to observe whether the model directly selects it upon encountering an option that aligns with the visual information.

**Semantic Misleading Option** Moreover, we introduce a semantic misleading option, such as “*The Chrysler Corporation*”, to challenge the multi-hop reasoning of MLLMs in KG. This option refers to the entity that is pointed by the relation commonly owned by both the associated entity and its sampled neighbor. For example in Figure 3, upon encountering a question about “*brand*” and “*Chrysler PT Cruiser*”, MLLMs might simply output a direct answer (e.g., “*The Chrysler Corporation*”), ignoring other constraints in the question (e.g., “*followed by*”), hence struggling to choose the correct answer (e.g., “*Fiat Automobiles S.p.A*”).

**Ground Truth Option** Corresponding to the causal path via  $I \rightarrow G$  and  $S \rightarrow G$ , this option is the end entity of the entity-relation path (e.g., “*Fiat Automobiles S.p.A*”). Finally, we check and ensure that each option is distinct from the remaining three to eliminate any potential overlap in samples.

### 3.4 Causal Rationale Generation

Furthermore, the entity-relation paths can assist in generating the reasoning rationale, termed as

*causal rationale*, tailored to the question at hand. In this context, we have employed a heuristic rule-based approach, starting from the associated entity and generating a causal rationale step by step until reaching the ground truth answer. Templates of generation can be found in the Appendix A.3. These generated causal rationales can be utilized to verify whether the reasoning process of MLLMs is correct, thereby providing interpretability. They can also be used to fine-tune specialized MLLMs to enhance their multi-hop reasoning capabilities, with experimental results available in Section 5.3. The causal rationales could further be polished and refined through LLMs such as ChatGPT, and we leave this for future work.

### 3.5 Dataset Statistics

**Statistics of Different Hops** Using the aforementioned construction approach, we automatically generate training data from the train set of INFOS-EEK, and development/test data from the validation set of INFOSEEK. Table 4 in Appendix shows the statistics of our curated dataset MORE. For the purpose of generating questions with depths of 2-hop and 3-hop, we set  $n = 1, 2$ , respectively. We do not include longer-hop queries here due to concerns that their complexity will render the generated questions ambiguous or difficult to comprehend.

**Question Distribution** We categorize the generated questions into distinct types based on their starting  $n$ -grams in Figure 17 in the Appendix. The MORE dataset showcases an extensive lexical diversity in the questions generated.

### 3.6 Quality Analysis

**Question Quality** To ensure the quality of the comprising datasets, we analyze the lexical diversity and the fluency of the generated questions. The considered baselines and metrics are detailed in Appendix C.1. From Figure 4, MORE shows superiority in terms of lexical diversity and fluency, even in comparison to datasets created through human annotator-based question generation.

**Human Evaluation** We also conduct a human evaluation to validate and assess the quality of the generated questions. Results in Appendix C.2 show that 90.6% of the generated questions are classified as valid by the annotators, further demonstrating the quality of our datasets.

## 4 Evaluating MLLMs on MORE

### 4.1 Experimental Setup

**Datasets** We use all the test data of MORE for evaluation. We adopt two different settings: 1) **Open-ended**. Ask the MLLM to generate answers based on the input image and question, or 2) **Multi-choice**. Provide MLLM with four options and let it choose the correct answer from them. The latter setting has a random baseline (accuracy of 25%).

**Baselines** We evaluate various leading MLLMs on our MORE dataset in a zero-shot fashion, including two limited-access MLLMs: GPT-4V (OpenAI, 2023) and Gemini Pro (Team et al., 2023), and four open-source MLLMs: BLIP-2 (6.7B) (Li et al., 2023), InstructBLIP (13B) (Dai et al., 2023), mPLUG-Owl (7B) (Ye et al., 2023), and LLaVA (v1.5, 13B) (Liu et al., 2023a), detailed in Appendix D.1. As for evaluation metrics, we adopt the VQA accuracy (Antol et al., 2015) for all models for a fair comparison. We also quantify the causal effects of images and questions on models’ predictions in Section 4.3.

### 4.2 Evaluation Results

We present the results of MLLMs separately on two-hop, three-hop, and all data of MORE in Table 2. Examples of model prediction can be found in Appendix H. We observe that: 1) The performance of all baselines on MORE is poor (e.g., under the “Multi-choice” setting, only Gemini Pro exceeded the random baseline at 28.9% accuracy), indicating MLLMs’ vulnerability to language and vision biases. 2) There is still a gap between open-source and limited-access models on MORE, especially under the “Open-ended” setting. 3) Most models perform better on two-hop data than on three-hop data (Gemini Pro is particularly outstanding on two-hop data, achieving 33.5% accuracy), suggesting that MLLMs’ reasoning capabilities are challenged when the problems become more complex. 4) GPT-4V performs best under the “Open-ended” setting, but falls short under the “Multi-choice” setting versus Gemini Pro, possibly due to that we use homologous ChatGPT-generated distractors when constructing the language bias option in Section 3.3, which poses a greater challenge to GPT-4V’s judgment. This point is also verified in the following analyses in Section 4.3.

Model	LLM	# Params	MORE (Two-Hop, acc (%))		MORE (Three-Hop), acc (%)		MORE (Overall, acc (%))	
			Open-ended	Multi-choice	Open-ended	Multi-choice	Open-ended	Multi-choice
Random	/	/	/	25.0	/	<b>25.0</b>	/	25.0
BLIP-2	OPT	6.7B	4.0	16.4	1.4	15.4	2.7	15.9
InstructBlip	Vicuna	13B	3.0	17.0	1.6	16.2	2.3	16.6
mPLUG-Owl	Llama	7B	4.0	12.4	8.2	11.4	6.1	11.9
LLaVA	Llama	13B	8.0	20.8	6.8	13.6	7.4	17.5
GPT-4V	-	-	<b>15.8</b>	25.6	<b>15.3</b>	23.2	<b>15.6</b>	24.4
Gemini Pro	-	-	14.2	<b>33.5</b>	10.1	24.4	12.2	<b>28.9</b>

Table 2: MLLMs’ results on the test set of MORE. We report the VQA accuracy (%) under the open-ended generation and the multi-choice settings on two-hop, three-hop, and all data, respectively. “-” denotes not released information.

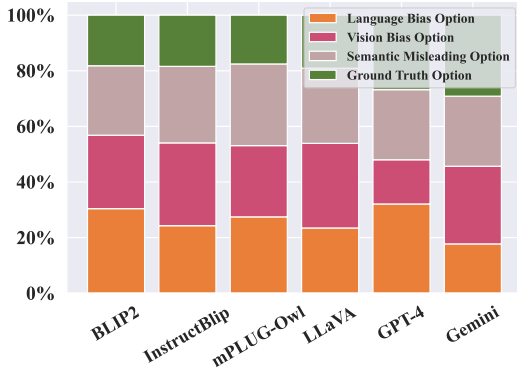


Figure 5: Option distribution of MLLMs.

### 4.3 Causal Analysis of VQA Biases

In this subsection, we analyze the performance of MLLMs through a causal lens.

**Option Distribution** In Figure 5, we show the distribution of options under the “Multi-choice” setting across various MLLMs. We observe that: 1) BLIP2 and GPT-4V often incorrectly choose options indicating language bias, aligning with our prior analysis on GPT-4V. 2) The proportion of either language or vision bias exceeded 40% across all the models, demonstrating the significant impact of unimodal biases on their predictions. 3) To some extent, models’ selection of semantically misleading options indicates some ability to combine visual and textual information, though not fully grasping the problem. This highlights the challenge our MORE dataset poses to current MLLMs. Please note that discrepancies may exist between the proportions of ground truth options presented here and the accuracy values reported in Table 2, as some models’ outputs (e.g., mPLUG-Owl) may not align with the provided options, thus affecting the count of valid answers.

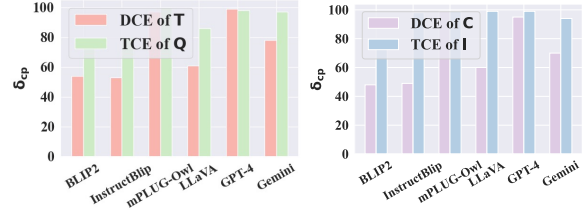


Figure 6: Comparison of direct and total effects of image and question on prediction for MLLMs.

**Causal Effects of Images and Questions** To further analyze the impact of vision bias and language bias on the model’s prediction, we assess the causal effects based on the definitions provided in Section 2.3. Specifically, we randomly select 100 samples for intervention and then measure the average of the effects over all instances to calculate TCE (corresponds to model’s *sensitivity*) and DCE (corresponds to models’ *robustness*). We implement two straightforward and natural intervention ways separately for questions and images. For questions: (i) we change both the core semantics and the textual form, or (ii) we only change the textual form while controlling to keep the core semantics unchanged. For images: (i) we change both the entities of interest and the visual context, or (ii) we only change the visual context while keeping the core entities unchanged. As discussed below Equation (9), we do not conduct sole interventions on the core entity while concurrently controlling for the irrelevant visual context, the ground truth answer, and the question. This is because it is extremely difficult to find such qualified samples to calculate  $DCE(E \rightarrow A)$  in practical scenarios. Overall, a higher TCE is desirable, indicating better sensitivity, while a lower DCE indicates better robustness. Implementation details of the intervention can be found in Appendix E.

From Figure 6: 1) Current MLLMs exhibit high sensitivity (high TCE), a possible reason is that



Model / Acc (%)	OK-VQA	INFOSEEK	MORE	
			Open-ended.	Multi-choice.
GPT-4V	70.4	38.0	15.6	24.4
Gemini Pro	64.9	37.8	12.2	28.9
LLaVA	72.8	16.0	7.4	17.5
GPT-4V (w. the DeVA framework)	<b>74.8</b>	<b>50.6</b>	<b>24.2</b>	30.3
Gemini Pro (w. the DeVA framework)	68.3	50.3	19.5	<b>32.1</b>
LLaVA (trained on INFOSEEK)	72.1	19.7	3.1	13.9
LLaVA (trained on <b>MORE</b> )	73.6	20.2	16.4	28.9

Table 3: Results comparison after employing our strategies on OK-VQA, INFOSEEK, and MORE.

instruction tuning makes models sensitive to variations in input (Stolfo et al., 2022). 2) However, the robustness is relatively low (high DCE), showing predictions change with variations in input even when the ground truth is fixed, suggesting the reliance on spurious paths over genuine causal features. 3) The relationship between MLLMs’ utilization of biases and their performance is intricate, with GPT-4V performing well even amid strong language and vision biases.

## 5 Mitigating VQA Biases for MLLMs

In this section, we explore two strategies to mitigate the unimodal biases of MLLMs.

### 5.1 Experimental Setup

**Dataset** We conduct experiments on three knowledge-based VQA datasets OK-VQA (Marino et al., 2019), INFOSEEK (Chen et al., 2023), and our curated MORE, where relying solely on image content proved insufficient for answering questions. We randomly select 1,000 instances from the validation set of INFOSEEK and OK-VQA, respectively, and report results on the test data of MORE, to ensure a consistent number of samples for evaluation. Since INFOSEEK and OK-VQA do not provide multiple options, we only report the results of them under the “Open-ended” settings.

**Baselines** We focus on three models that have shown superior performance in Table 2: LLaVA, GPT-4V, and Gemini Pro. For limited-access models GPT-4V and Gemini Pro, we propose a prompting-based **Decompose-Verify-Answer** (DeVA) framework. For the open-source model LLaVA, we conduct fine-tuning on various datasets and compare the results. We do not apply the DeVA framework to LLaVA based on findings from our pilot study that LLaVA’s performance in comprehending and responding to our instructions is not on par with that of GPT-4V and Gemini Pro.

### 5.2 A Decompose-Verify-Answer Framework

For a given instance, DeVA begins with a question decomposer to break down the complex question into several easier ones step-by-step, so as to explicitly avoid the model simply taking a spurious path to give the answer. For each subquestion, after providing an answer, DeVA employs a verifier to confirm the accuracy of the answer. The verifier strategizes the utilization of external tools and investigates their outputs, thereby acquiring the necessary context or information needed to provide a precise answer, such as image retrieval and text retrieval. This iterative process of answering and verifying will continue until all subquestions have been resolved. By incorporating external knowledge, the final verified output comprehensively integrated the information of the whole reasoning process to give a correct answer. A detailed illustration of DeVA and the prompt template is in Appendix F.

**Results** From Table 3, we can see that: 1) the DeVA framework significantly enhances the performance of GPT-4V and Gemini Pro across all three datasets. 2) Notably, the improvement on the INFOSEEK dataset is more pronounced than on MORE. This is because INFOSEEK is a dataset focused on information-seeking, and once we provide the models with relevant external knowledge, they can quickly extract the necessary information and output a correct answer. However, the limited results on MORE suggest that even when models are allowed access to additional information, they may still be impeded by language and vision biases. Therefore, how to disentangle and address this issue remains a challenging problem, which we leave for future work.

### 5.3 Fine-tuning LLaVA

We fine-tune LLaVA using the LoRA (Hu et al., 2022) technique on the training sets of INFOSEEK and MORE, respectively. For MORE, we incorporate

the generated causal rationale into the instruction of some samples, so as to enhance the models’ reasoning ability. Other implementation details are in Appendix G. From Table 3, we can see: 1) After fine-tuning, LLaVA shows a significant improvement in performance on OK-VQA, INFOSEEK, and MORE, particularly LLaVA (trained on MORE) outperforms the vanilla GPT-4V on OK-VQA and MORE. 2) While LLaVA (trained on INFOSEEK) does improve its performance on INFOSEEK, it significantly reduces its effectiveness on MORE. In contrast, LLaVA (trained on MORE) demonstrates better generalizability, confirming the effectiveness of incorporating rationales into instructions.

## 6 Related Work

**Multimodal Large Language Models** Recent advancements in LLMs have facilitated the development of powerful MLLMs. Pioneer work has demonstrated their impressive capability on multimodal tasks (OpenAI, 2023; Team et al., 2023; Liu et al., 2023a). Despite the substantial success, the current evaluation of MLLMs mainly focuses on basic visual tasks (Liu et al., 2023b; Fu et al., 2023; Lu et al., 2024), and their reasoning ability has not been investigated in detail.

**Knowledge-based VQA Datasets** Several previous works have studied knowledge-based VQA problems (Wang et al., 2017; Marino et al., 2019; Chen et al., 2023). While these datasets do require knowledge to answer questions, they have the following limitations: 1) They often focus solely on image-related information without requiring the model to perform multi-hop reasoning. 2) The answer types are typically open-ended and do not provide reasoning rationales, making it difficult to evaluate the model’s output. 3) They cannot assess the degree of reliance on language or visual biases by models. As initially introduced in Section 1, Table 1 provides a comparative summary of various datasets alongside our MORE dataset, underscoring its enhanced comprehensiveness.

**Language and Vision Biases in VQA** Studies have found that some VQA models tend to rely on statistical priors present in the training data, rather than truly understanding the content of images (Agrawal et al., 2018). This issue manifests in two primary forms: language bias and vision bias. The former emerges from strong correlations between specific questions and their answers (Ab-

basnejad et al., 2020; Zhu et al., 2020), and the latter stems from the frequent co-occurrence of certain textual and visual elements in the training dataset (Si et al., 2022; Gupta et al., 2022). Recent strategies to mitigate such biases mainly focus on employing data augmentation techniques (Niu et al., 2021). Although the concerns are similar, all the aforementioned methods focus on the statistical bias under the pre-train and fine-tune paradigm, while we focus on the reasoning ability of MLLMs.

## 7 Conclusion

This paper presents a comprehensive approach to quantifying and mitigating the unimodal biases in MLLMs. Through our causal inference framework, we provide an in-depth analysis to assess the causal effects of such biases on the model’s prediction in VQA problems. The introduced MORE dataset challenges MLLMs to engage in multi-hop reasoning and to overcome language and vision biases, thereby pushing the boundaries of their reasoning capabilities. Our proposed solutions, including the DeVA framework and the fine-tuned LLaVA models, demonstrate significant potential in enhancing the reasoning abilities of MLLMs.

## 8 Limitations

Our current generation of rationales is based on heuristic rules. Experiments in Section 5.3 have demonstrated the effectiveness of incorporating rationales into instructions. Therefore, we believe that refining and polishing these rationales with an LLM (e.g., ChatGPT) could be beneficial. Besides, the Wikidata5M dataset we employed was released in 2021, and some information in the knowledge graph may be outdated. Although we have made efforts to manually verify the test set and try to ensure it does not contain incorrect information, it is still inevitable that errors may occur within the extensive training data.

## References

- Ehsan Abbasnejad, Damien Teney, Amin Parvaneh, Javen Shi, and Anton van den Hengel. 2020. [Counterfactual vision and language learning](#). In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10041–10051. IEEE.
- Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. [Don’t just assume; look](#)

- and answer: Overcoming priors for visual question answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4971–4980. IEEE Computer Society.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. [VQA: visual question answering](#). In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2425–2433. IEEE Computer Society.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Samuel Cahyawijaya, Genta Indra Winata, Bryan Wilie, Karissa Vincentio, Xiaohong Li, Adhiguna Kuncoro, Sebastian Ruder, Zhi Yuan Lim, Syafri Bahar, Masayu Khodra, Ayu Purwarianti, and Pascale Fung. 2021. [IndoNLG: Benchmark and resources for evaluating Indonesian natural language generation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8875–8898, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William Cohen. 2022. [MuRAG: Multimodal retrieval-augmented generator for open question answering over images and text](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5558–5570, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- 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.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%\\* chatgpt quality](#).
- Michael A Covington and Joe D McFall. 2010. Cutting the gordian knot: The moving-average type–token ratio (mattr). *Journal of quantitative linguistics*, 17(2):94–100.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. [Instructblip: Towards general-purpose vision-language models with instruction tuning](#).
- Nanyi Fei, Zhiwu Lu, Yizhao Gao, Guoxing Yang, Yuqi Huo, Jingyuan Wen, Haoyu Lu, Ruihua Song, Xin Gao, Tao Xiang, et al. 2022. Towards artificial general intelligence via a multimodal foundation model. *Nature Communications*, 13(1):3094.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiwu Zheng, Ke Li, Xing Sun, et al. 2023. [Mme: A comprehensive evaluation benchmark for multimodal large language models](#). *ArXiv preprint*, abs/2306.13394.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. [Making the V in VQA matter: Elevating the role of image understanding in visual question answering](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 6325–6334. IEEE Computer Society.
- Tianrui Guan, Fuxiao Liu, Xiyang Wu Ruiqi Xian Zongxia Li, Xiaoyu Liu Xijun Wang, Lichang Chen Furong Huang Yaser Yacoob, and Dinesh Manocha Tianyi Zhou. 2023. [Hallusionbench: An advanced diagnostic suite for entangled language hallucination & visual illusion in large vision-language models](#). *arXiv e-prints*, pages arXiv–2310.
- Vipul Gupta, Zhuowan Li, Adam Kortylewski, Chenyu Zhang, Yingwei Li, and Alan L. Yuille. 2022. [Swapmix: Diagnosing and regularizing the over-reliance on visual context in visual question answering](#). In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 5068–5078. IEEE.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Aman Jain, Mayank Kothiyari, Vishwajeet Kumar, Preethi Jyothi, Ganesh Ramakrishnan, and Soumen Chakrabarti. 2021. Select, substitute, search: A new benchmark for knowledge-augmented visual question answering. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2491–2498.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and



- Wen-tau Yih. 2020. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. [Generalization through memorization: Nearest neighbor language models](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). *ArXiv preprint*, abs/2301.12597.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. [Visual instruction tuning](#). *ArXiv preprint*, abs/2304.08485.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023b. [Mmbench: Is your multi-modal model an all-around player?](#) *ArXiv preprint*, abs/2307.06281.
- Chaochao Lu, Chen Qian, Guodong Zheng, Hongxing Fan, Hongzhi Gao, Jie Zhang, Jing Shao, Jingyi Deng, Jinlan Fu, Kexin Huang, et al. 2024. From gpt-4 to gemini and beyond: Assessing the landscape of mlms on generalizability, trustworthiness and causality through four modalities. *arXiv preprint arXiv:2401.15071*.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. [OK-VQA: A visual question answering benchmark requiring external knowledge](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 3195–3204. Computer Vision Foundation / IEEE.
- Philip M McCarthy. 2005. *An assessment of the range and usefulness of lexical diversity measures and the potential of the measure of textual, lexical diversity (MTLD)*. Ph.D. thesis, The University of Memphis.
- Philip M McCarthy and Scott Jarvis. 2010. Mtl-d, vocd-d, and hd-d: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior research methods*, 42(2):381–392.
- Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. 2021. [Counterfactual VQA: A cause-effect look at language bias](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 12700–12710. Computer Vision Foundation / IEEE.
- OpenAI. 2023. Gpt-4 technical report.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Judea Pearl. 1995. Causal diagrams for empirical research. *Biometrika*, 82(4):669–688.
- Judea Pearl. 2022. Direct and indirect effects. In *Probabilistic and causal inference: the works of Judea Pearl*, pages 373–392.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- 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.
- Lucas Shen. 2022. Lexicalrichness: A small module to compute textual lexical richness.
- Qingyi Si, Fandong Meng, Mingyu Zheng, Zheng Lin, Yuanxin Liu, Peng Fu, Yanan Cao, Weiping Wang, and Jie Zhou. 2022. [Language prior is not the only shortcut: A benchmark for shortcut learning in VQA](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3698–3712, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Alessandro Stolfo, Zhijing Jin, Kumar Shridhar, Bernhard Schölkopf, and Mrinmaya Sachan. 2022. [A causal framework to quantify the robustness of mathematical reasoning with language models](#). *ArXiv preprint*, abs/2210.12023.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. 2023. [Gemini: a family of highly capable multimodal models](#). *ArXiv preprint*, abs/2312.11805.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#). *ArXiv preprint*, abs/2302.13971.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023b. [Llama: Open and efficient foundation language models](#). *ArXiv preprint*, abs/2302.13971.



- Peng Wang, Qi Wu, Chunhua Shen, Anthony Dick, and Anton Van Den Hengel. 2017. Fvqa: Fact-based visual question answering. *IEEE transactions on pattern analysis and machine intelligence*, 40(10):2413–2427.
- Qingyun Wang, Lifu Huang, Zhiying Jiang, Kevin Knight, Heng Ji, Mohit Bansal, and Yi Luan. 2019. [PaperRobot: Incremental draft generation of scientific ideas](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1980–1991, Florence, Italy. Association for Computational Linguistics.
- Xiaozhi Wang, Tianyu Gao, Zhaocheng Zhu, Zhengyan Zhang, Zhiyuan Liu, Juanzi Li, and Jian Tang. 2021. [KEPLER: A unified model for knowledge embedding and pre-trained language representation](#). *Transactions of the Association for Computational Linguistics*, 9:176–194.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. 2023. [The dawn of lmms: Preliminary explorations with gpt-4v \(ision\)](#). *ArXiv preprint*, abs/2309.17421.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. 2023. [mplug-owl: Modularization empowers large language models with multimodality](#). *ArXiv preprint*, abs/2304.14178.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. [From recognition to cognition: Visual commonsense reasoning](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6720–6731. Computer Vision Foundation / IEEE.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.
- Xi Zhu, Zhendong Mao, Chunxiao Liu, Peng Zhang, Bin Wang, and Yongdong Zhang. 2020. [Overcoming language priors with self-supervised learning for visual question answering](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pages 1083–1089. ijcai.org.
- Yuke Zhu, Oliver Groth, Michael S. Bernstein, and Li Fei-Fei. 2016. [Visual7w: Grounded question answering in images](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 4995–5004. IEEE Computer Society.

## Question Generation

<b>Task Description:</b> Provide a question according to the starting entity and path (relations split by ‘,’) in a knowledge graph.
<b>Template:</b> Starting entity: <HEAD_ENTITY> Path: <RELATION_PATH_IN_KG>
<b>Examples:</b>
Starting entity: Coca-Cola Path: discoverer or inventor, place of birth <b>Generated Question:</b> Which city is the birthplace of the inventor of Coca-Cola?
Starting entity: James Berkeley Path: place of birth, owned by <b>Generated Question:</b> Who is the owner of the building where James Berkeley was born?
Starting entity: Conwy Castle Path: country, highest point, material used <b>Generated Question:</b> What material is used in the highest point of the country where Conwy Castle is located?

Figure 7: Prompt template of multi-hop question generation.

## Language Bias Option Generation

<b>Task Description:</b> Given a question, provide a specific answer.
<b>Examples:</b>
Question: What is the parent taxon of the main food source of this animal? <b>Answer:</b> Animalia
Question: What is the heritage designation of the burial place of the person who commissioned this building? <b>Answer:</b> UNESCO World Heritage Site.

Figure 8: Prompt template of language bias option generation.

Dataset	#I, Q, A	Len of Q / A	# Ent
MORE-train	10K	14.3 / 2.1	1,261
- 2-hop	4,134	11.6 / 2.0	886
- 3-hop	5,866	16.1 / 2.2	686
MORE-dev	1K	13.8 / 2.3	118
- 2-hop	548	12.2 / 2.2	71
- 3-hop	452	15.8 / 2.5	73
MORE-test	1K	13.9 / 2.4	251
- 2-hop	500	12.3 / 2.2	153
- 3-hop	500	15.6 / 2.6	143

Table 4: Dataset statistics of different hops.

## A Prompt Templates

### A.1 Question Generation

We present the prompt template for generating language bias options in Section 3.3 in Figure 7.

### A.2 Language Bias Option Generation

We present the prompt template for generating language bias options in Section 3.3 in Figure 8.

### A.3 Training Instance with Causal Rationale

An example of the training instance in MORE is shown in Figure 9.

## B Question Distribution

In Figure 17, we categorize the generated questions into distinct types, based on their starting n-grams. The dataset MORE showcases an extensive lexical diversity in the questions generated. This diversity is evidenced by variations in the introductory interrogative words (e.g., “what”, “who”, “where”, etc.), exemplified by phrases like “What is the...”, “In which country...”, and more. Such lexical richness is crucial for mitigating the vulnerability of MLLMs to linguistic variations.

## C Quality Analysis Details

### C.1 Question Quality

To ensure the quality of the comprising datasets, we analyze the lexical diversity and the fluency of the generated questions, which are useful for conducting a robust evaluation using questions that are linguistically diverse and coherent.

## Training Instance with Causal Rationale

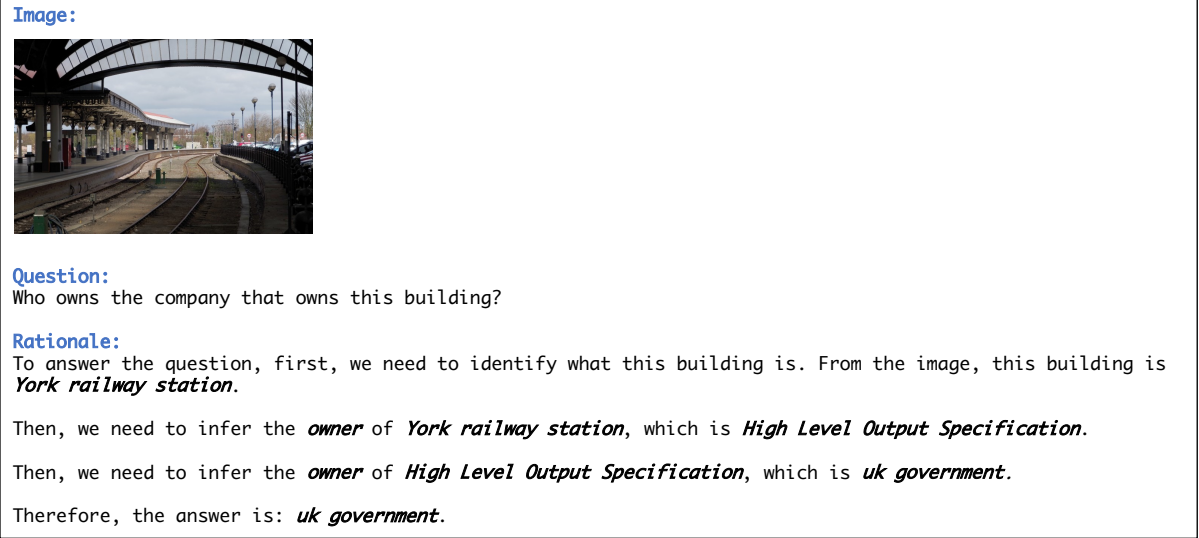


Figure 9: Prompt template of the training instance in MORE.

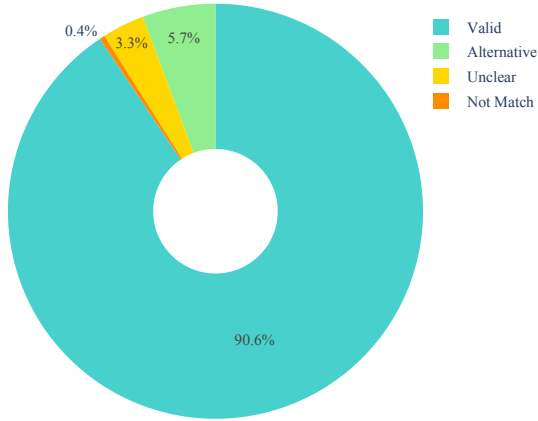


Figure 10: Human evaluation results of MORE.

**Baselines** We select extensive VQA datasets for comparison, including Visual7W (Zhu et al., 2016), VQA (v2) (Goyal et al., 2017), FVQA (Wang et al., 2017), OK-VQA (Marino et al., 2019), S3VQA (Jain et al., 2021), A-OKVQA (Schwenk et al., 2022), and INFOSEEK (Chen et al., 2023) (contains both automated generation version and human-annotated version).

**Evaluation Metrics** For lexical diversity, we utilize three metrics that are not dependent on length: moving average type-token ratio (MATTR) (Covington and McFall, 2010), measure of textual lexical diversity (MTLD) (McCarthy, 2005), and hypergeometric distribution diversity (HDD) (McCarthy and Jarvis, 2010). We average these three metrics for a unified assessment and employ the Lexical-Richness package (Shen, 2022) (version 0.5.03) for

calculation. For fluency, we employ a pre-trained language model GPT2-large (Radford et al., 2019) with 774M parameters to compute the perplexity of the questions, which is often used as a measure by previous work (Wang et al., 2019; Cahyawijaya et al., 2021).

## C.2 Human Evaluation

We conduct a human evaluation of 100 questions randomly chosen from the MORE dataset to validate and assess the quality of the generated questions. This evaluation is carried out by three human annotators, who are provided with detailed guidelines and illustrative examples before starting the evaluation process. For each question, given the visual context and ground truth answer, the annotators are asked to determine whether: 1) the question is valid, 2) the question has a possible answer alternative, 3) the question does not match the answer, or 4) the question is unclear or confusing. For each sample, we ask three human annotators to judge each generated question into one of the four options.

The outcomes of this human evaluation are summarized in Figure 10. The results are encouraging, with 90.6% of the generated questions being classified as valid by the annotators, further demonstrating the quality of our datasets.

## D Baselines and Datasets Details

### D.1 Baselines

For open-source MLLMs, we consider the following baselines:

1) BLIP2 (Li et al., 2023), which utilizes a scalable multimodal pre-training method to enable LLMs to understand images. We employ its BLIP2-OPT (Zhang et al., 2022)-6.7B variant.

2) InstructBLIP (Dai et al., 2023), an instruction-tuned version of BLIP-2 on various tasks including VQA. We employ its InstructBLIP-Vicuna (Chiang et al., 2023)-13B variant.

3) mPLUG-Owl (Ye et al., 2023), which proposes a new two-stage training method for aligning images and text. We employ its mPLUG-Owl-Llama (Touvron et al., 2023a)-7B variant.

4) LLaVA (Liu et al., 2023a), which translates images into texts of captions and bounding boxes, and prompts GPT-4 to generate a multimodal instruct-tuning dataset in the context of seed examples. We employ its LLaVA-Llama (Touvron et al., 2023a)-13B variant.

## D.2 Datasets

**OK-VQA** (Marino et al., 2019) is a VQA dataset that necessitates external knowledge beyond the visible content in images. It is crafted by crowd-sourced participants who form complex questions based on extensive information available on Wikipedia.

**INFOSEEK** (Chen et al., 2023) is another VQA dataset aimed at information-seeking questions that require more than just commonsense knowledge. It combines human-annotated questions with visual entity recognition datasets and Wikidata to generate complex question-answer pairs.

## E Implementation Details of Intervention

In this section, we introduce the implementation details of interventions on questions and images. Given the impracticality of enumerating all possible perturbations of  $T$  and  $R$ , practical results within a specific subset of  $T$  and  $R$  could be attained by intervening on  $Q$  while leaving  $S$  unaffected, and intervening on  $I$  while leaving  $E$  unaffected, respectively (Stolfo et al., 2022).

For questions, in the analysis of TCE, the image stays the same while the question is altered to a different one from the dataset that relates to the same entity, ensuring a ground truth shift post-intervention. In the analysis of DCE, we similarly keep the image constant and let ChatGPT rephrase the question, altering its textual form but preserving its semantic meaning (verified through manual evaluation), thereby maintaining an unchanged ground

truth after the intervention.

For images, in the analysis of TCE, we maintain a constant question while replacing the image with another from the dataset that corresponds to the same question but features different entities, ensuring a change in the ground truth post-intervention. In the analysis of DCE, the question remains unchanged, but the image is replaced with one of the same entity from an alternative source, specifically using Google Lens’s API\*) to keep the ground truth consistent after the intervention.

## F The DeVA Framework

In this section, we first present the overall framework, then we will go over each part of it in detail.

### F.1 Overall Framework

Given a question  $Q$  and an image  $I$ , the VQA task demands the system to return an output  $A$  that concisely answers the question. As shown in Figure 11, we first initialize a question decomposer  $D$  to analyze  $Q$  and break it down into manageable sub-questions. Then, we employ a verifier  $V$  to confirm the accuracy of the original answer to each decomposed subquestion. The generation verifier typically involves active information-seeking and answer-verification, which acquires the necessary context or information needed to investigate and revise the answers. This includes two optional operations: image retrieval to seek images similar to  $I$  and determine their titles, or text retrieval with a specific query to fetch pertinent documents and summarize their content. This iterative process of answering and verifying will continue until we resolve each subquestion. Finally, the verified answer  $A$  is output following the aforementioned reasoning process and retrieved information.

### F.2 Question Decomposer

For a given VQA problem, MLLMs often simply exploit a spurious path to make predictions due to the inherent language and vision biases. In order to alleviate this issue, motivated by Chain-of-Thought reasoning (Wei et al., 2022), we encourage MLLMs to decompose the question  $Q$  before outputting the answer, so as to gradually solve a complex question that requires multi-hop reasoning. As shown in Figure 11, for the question “Which country hosted the next World Cup after this venue?”, our decomposer breaks it down into two subquestions :

\*Web interface available at <https://images.google.com>.



1. *“What this venue is?”*
2. *“Which country hosted the next World Cup after Allianz Arena?”*

Such decomposition will explicitly constrain the model to comprehend and extract the truth semantics of the question, thus avoiding simply exploring a spurious path to give the answer.

### F.3 Verifier

Some works have found that vision illusion and language hallucination may appear in the process of MLLMs’ response generation (Guan et al., 2023; Yang et al., 2023). To alleviate this issue, we adopt the retrieval-augmented generation approaches (Khandelwal et al., 2020; Chen et al., 2022). Specifically, we consider two different retrieval ways for the verifier to choose during each verification step: image retrieval and text retrieval.

**Image Retrieval** Although our framework is applicable to any image retrieval method, in this paper, we mainly utilize Google Image Search to obtain a broad range of information related to the image as provided by Google Lens API. This information encompasses various details, such as knowledge graph entities and captions of analogous or identical images. The availability of these details can vary based on the image input provided to Google Image Search. Then, the verifier gleans relevant information from captions associated with visually similar images, so as to verify the original answer and conduct the next round of reasoning.

**Text Retrieval** Similarly, our framework is applicable to any text retrieval method, we explore a simple, off-the-shelf dense retriever for Wikipedia, GTR (Karpukhin et al., 2020), as our text retriever. First, the verifier constructs a query to perform text retrieval according to the currently generated context, and then the query is input into a GTR model to get related document titles and contents. Finally, the verifier will fetch pertinent documents and summarize their content to verify the immediate answer.

### F.4 Final Verified Answer

Finally, the improved response that takes verification into account is generated. This is executed by a final prompt where the context takes into account all of the previous reasoning steps, the baseline response, and the verification question-answer pairs,

so that the corrections can take place. The prompt template is shown in Figure 12.

## G Implementation Details of Fine-tuning LLaVA

We adopt LLaVA-13B-v1.5 as the base model and employ the LoRA (Hu et al., 2022) technique. During pre-training, only LoRA parameters are optimized. We set the rank of LoRA modules to 128. Our model is optimized with a learning rate of  $2e-4$  and a linear warm-up for the first 3% steps. The batch size is 16 and the number of epochs is 1. All the LoRA parameters are trained on 8 NVIDIA A100 GPUs with 80GB memory. We fine-tune LLaVA on the training sets of INFOSEEK and MORE, respectively. For MORE, we incorporate the generated rationale into the instructions of some samples, so as to enhance the models’ reasoning ability. The template of a training instance can be found in Figure 9.

## H Case Study

We conduct a case study on the development set of MORE in Figure 13~16, including both the “Open-ended” and “Multi-Choice” settings.

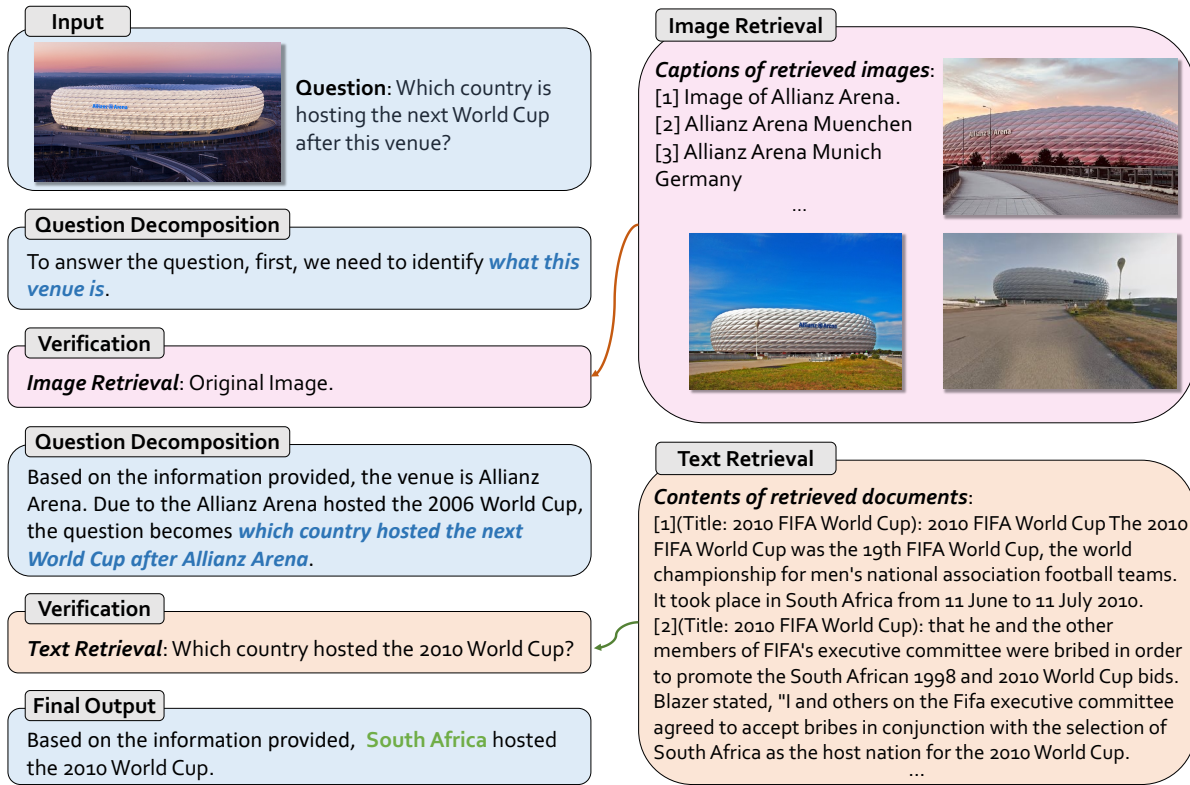


Figure 11: An overview of our proposed Decompose-Verify-Answer (DeVA) Framework.

### Decompose-Verify-Answer Prompting

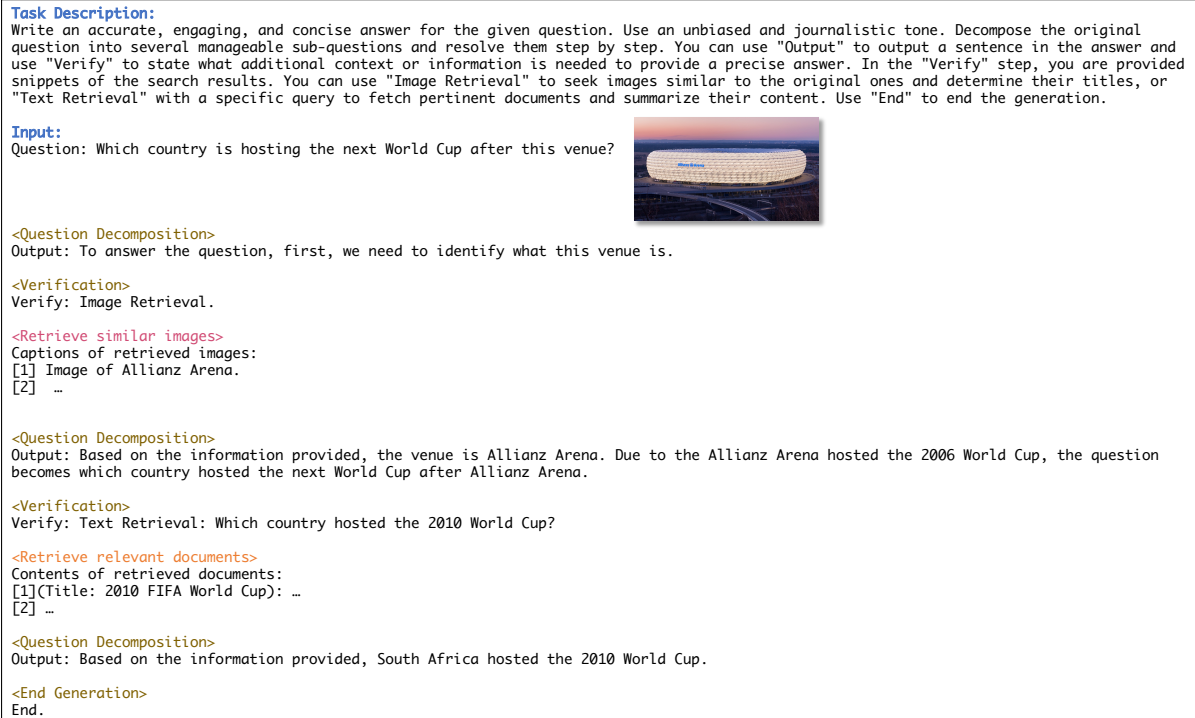
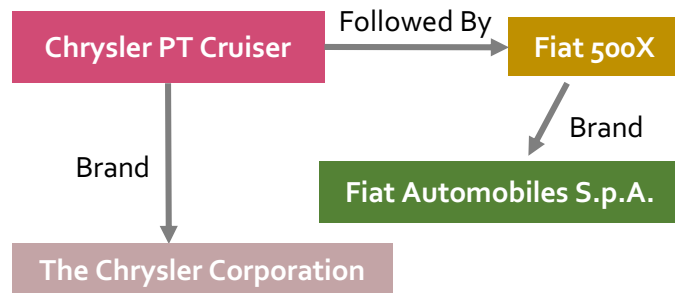


Figure 12: Prompt template of DeVA framework.

## 2-hop case



**Question:** What brand follows this vehicle in its product line?

**BLIP2**

The chrysler pt cruiser

Vision Bias

**InstructBLIP**

chrysler

Vision Bias

**mPLUG-Owl**

The vehicle follows the brand BMW in its product line.

Visual Illusion

**LLaVA**

Fiat

Ground Truth

**GPT-4**

The vehicle in the image is a Chrysler PT Cruiser Convertible. The Chrysler PT Cruiser is a model from the Chrysler brand, which is part of Stellantis, a multinational automotive manufacturing corporation.

Semantic Misleading

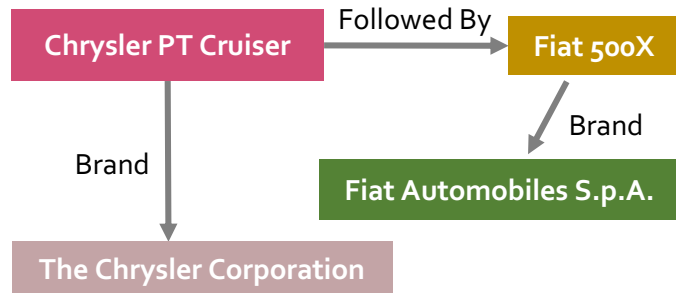
**Gemini Pro**

The Chrysler PT Cruiser is a front-wheel drive, four-passenger convertible that was produced by Chrysler from 2001 to 2010. The PT Cruiser was based on the Chrysler Neon platform and was available in two trim levels: the base model and the GT. The PT Cruiser was powered by a 2.4-liter four-cylinder engine that produced 150 horsepower and 165 lb-ft of torque. The PT Cruiser was discontinued after the 2010 model year.

Semantic Misleading

Figure 13: Two-hop case in the dev set of MORE. The wrong answers are marked in red and the correct is in green.

## 2-hop case (Option)



**Question:** What brand follows this vehicle in its product line?

Choose from the following four options:

- A. Chrysler PT Cruiser
- B. Tesla
- C. The Chrysler Corporation
- D. Fiat Automobiles S.p.A.

### BLIP2

The Chrysler PT Cruiser

Vision Bias

### InstructBLIP

Chrysler PT Cruiser

Vision Bias

### mPLUG-Owl

Chrysler PT Cruiser

The brand of the car that follows the model which this vehicle follows is Chrysler PT Cruiser. This is the correct answer.

Vision Bias

### LLaVA

Chrysler PT Cruiser

Vision Bias

### GPT-4

The vehicle in the image is a Chrysler PT Cruiser, which is a model produced by Chrysler. Therefore, the brand that follows the model which this vehicle follows is Chrysler.

Semantic Misleading

### Gemini Pro

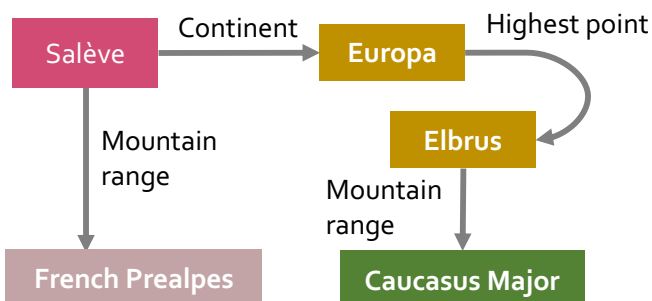
Chrysler PT Cruiser

Vision Bias

Figure 14: Two-hop case (multi-choice setting) in the dev set of MORE. The wrong answers are marked in red and the correct is in green.



### 3-hop case



**Question:** What is the mountain range of the highest point on the continent where this mountain is located?

#### BLIP

The alaska range

Visual Illusion

#### InstructBLIP

rocky

Visual Illusion

#### mPLUG-Owl

The mountain of Mount McKinley, also known as Denali, is located in Alaska, USA. It is the tallest mountain in North America and the tallest mountain in the world when measured from the base of its sea-level summit. The mountain range that Mount McKinley is a part of is the Alaska Range.

Visual Illusion

#### LLaVA

The highest mountain range is the Alps, the highest peak is Mont Blanc.

Language Bias

#### GPT-4

The mountain in the image is the Salève. The Salève is a mountain located in the Prealps, near Geneva, Switzerland, and France.

Semantic Misleading

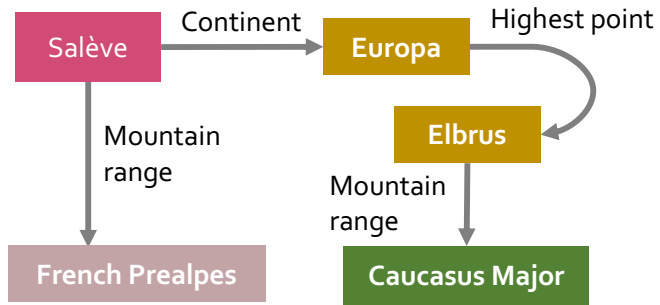
#### Gemini Pro

The mountain range is the Chugach Mountains, which are part of the larger Coast Mountains.

Visual Illusion

Figure 15: Three-hop case in the dev set of MORE. The wrong answers are marked in red.

### 3-hop case (Option)



**Question:** What is the mountain range of the highest point on the continent where this mountain is located?

Choose from the following four options:

- A. Salève B. the Himalayas  
C. French Prealps D. Caucasus Major

#### BLIP2

The Himalayas

Language Bias

#### InstructBLIP

French Prealps

Semantic Misleading

#### mPLUG-Owl

The mountain range of the highest point on the continent where this mountain is located is Salève.

Salève, which is located in the Rhône-Alpes region of France, is a small mountain range that extends from the city of Geneva to the city of Annecy. The summit of this range, Mont Salève, is the highest point on the continent where this mountain is located. Salève is a popular destination for hiking and skiing, and it offers spectacular views of the surrounding Alps, including the Mont Blanc massif and the Swiss Alps. It is also a popular destination for rock climbing, and the steep cliffs and rocky faces of Salève provide a challenging climbing experience.

Vision Bias

#### LLaVA

Caucasus Major

Ground Truth

#### GPT-4

Based on the options provided, the highest point on the continent where the Salève mountain is located would be Mount Everest, as Salève is part of the Prealps, a region in Europe, and Mount Everest is the highest point on the Asian continent. The Greater Caucasus Mountains are also in Asia but are not as high as Mount Everest. Therefore, the correct answer from the options given is the Himalayas.

Language Bias

#### Gemini Pro

The highest point of the mountain range is Mont Blanc, which is located in the French Prealps. The French Prealps are located in Europe.

Semantic Misleading

Figure 16: Three-hop case (multi-choice setting) in the dev set of MORE. The wrong answers are marked in red and the correct is in green.

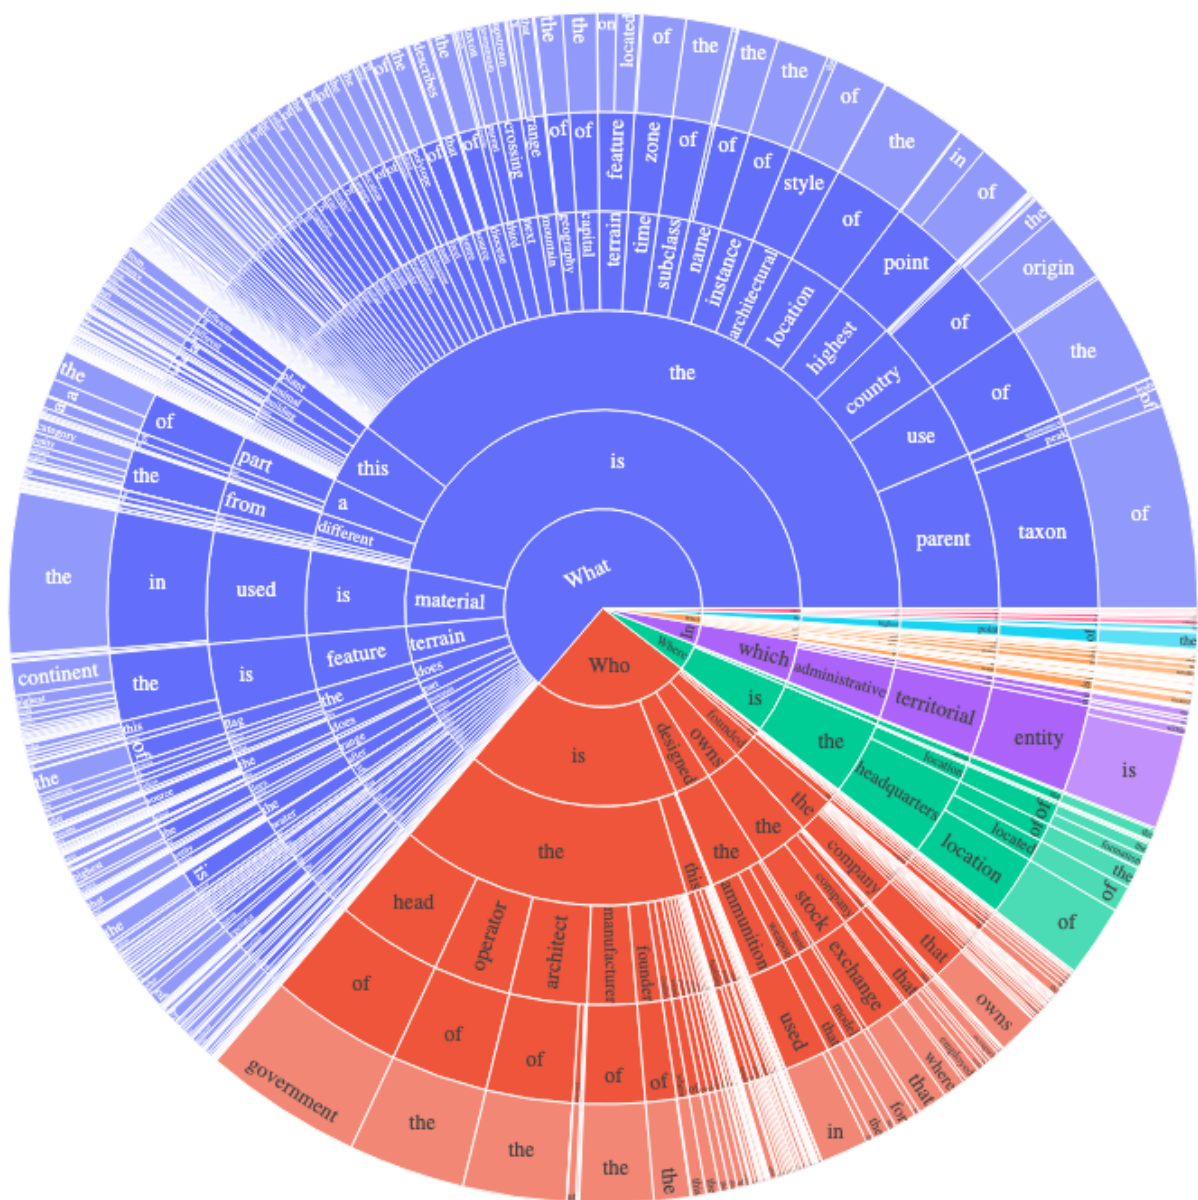


Figure 17: Question prefix distribution in MORE. The arc length is proportional to the number of questions containing the word.