

“See the World, Discover Knowledge”: A Chinese Factuality Evaluation for Large Vision Language Models

Jihao Gu*, Yingyao Wang†, Pi Bu,
Chen Wang, Ziming Wang, Tengtao Song, Donglai Wei, Jiale Yuan,
Yingxiu Zhao, Yancheng He, Shilong Li, Jiaheng Liu, Meng Cao, Jun Song‡,
Yingshui Tan, Xiang Li, Wenbo Su, Xiaoyong Zhu, Bo Zheng,
Alibaba Group
{gujihao.gjh, wangyingyao.wyy, jsong.sj}@taobao.com

Abstract

The evaluation of factual accuracy in large vision language models (LVLMs) has lagged behind their rapid development, making it challenging to fully reflect these models’ knowledge capacity and reliability. In this paper, we introduce the first factuality-based visual question-answering benchmark in Chinese, named **ChineseSimpleVQA**, aimed at assessing the visual factuality of LVLMs across 8 major topics and 56 subtopics. The key features of this benchmark include a focus on the **Chinese** language, **diverse** knowledge types, a **multi-hop** question construction, **high-quality** data, **static** consistency, and **easy-to-evaluate** through short answers. Moreover, we contribute a rigorous data construction pipeline and decouple the visual factuality into two parts: seeing the world (i.e., object recognition) and discovering knowledge. This decoupling allows us to analyze the capability boundaries and execution mechanisms of LVLMs. Subsequently, we evaluate 35 advanced open-source and closed-source models, revealing critical performance gaps within this field. Our evaluation-friendly code ¹ and data ² have already been open-sourced.

1 Introduction

In the multimodal field, ensuring the factual accuracy of responses generated by Large Vision Language Models (LVLMs) presents a significant challenge. Currently, cutting-edge models still produce erroneous outputs that do not align with the image content or provide answers lacking support from knowledge evidence. This issue is known as the “visual hallucination”, primarily arising from

*Work done during an internship at Alibaba Group.

†Equal contribution.

‡Corresponding Author.

¹<https://github.com/OpenStellarTeam/ChineseSimpleVQA>

²<https://huggingface.co/datasets/OpenStellarTeam/Chinese-SimpleVQA>

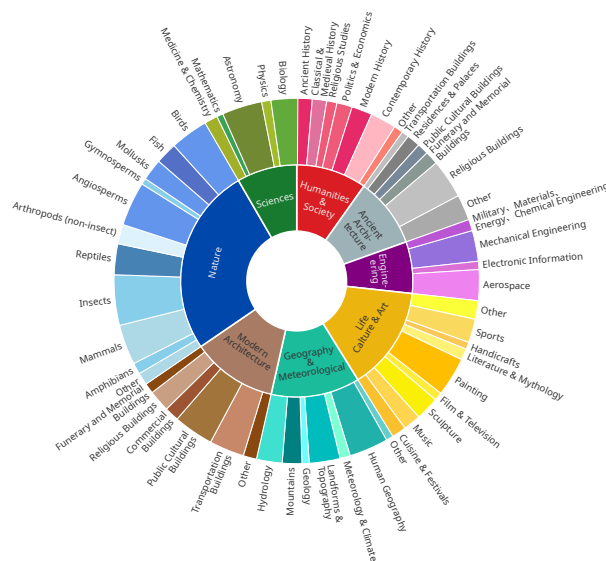


Figure 1: The overview of ChineseSimpleVQA, which evaluates visual factuality covering 56 subtopics.

difficulties in multimodal alignment and the lack of visual knowledge. In the realm of Large Language Models (LLMs), the factuality assessment of linguistic hallucination has garnered extensive attention. For instance, OpenAI introduced the SimpleQA benchmark (Wei et al., 2024), and Alibaba launched a Chinese SimpleQA (He et al., 2024), both of them efficiently evaluate the lengthy responses generated by LLMs containing numerous factual claims through simple questions.

However, the factuality assessment concerning visual hallucination has not received adequate attention (Fu et al., 2024). Researchers have extensively employed benchmarks such as OK-VQA (Marino et al., 2019) and A-OKVQA (Schwenk et al., 2022) to assess the factuality capabilities of LVLMs, but both of them have a narrow knowledge coverage (e.g., A-OKVQA only includes 7 categories). Furthermore, we emphasize that the evaluation of visual factuality should be decoupled into two parts: seeing the world (i.e., object recognition) and discovering knowledge. As illustrated in Figure 2, a

Benchmark	Size	Lang	Source	Categories	Vision	Open-ended	Short-form	Metric
OK-VQA (Marino et al., 2019)	8,062	en	Exams	20	✓	✗	✗	Accuracy
A-OKVQA (Schwenk et al., 2022)	125	zh	Human Collection	7	✓	✓	✗	LLM-as-a-Judge
SimpleQA (Wei et al., 2024)	4,326	en	Human Writers	10	✗	✓	✓	LLM-as-a-Judge
Chinese SimpleQA (He et al., 2024)	3,000	zh	Self-constructed & Human Writers	99	✗	✓	✓	LLM-as-a-Judge
ChineseSimpleVQA (Ours)	2,200	zh	Self-constructed & Human Writers	56	✓	✓	✓	LLM-as-a-Judge

Table 1: Comparisons between our ChineseSimpleVQA and other knowledge-based VQA benchmarks.

model must first comprehend that the image depicts a “chestnut” to answer the next question, “What family does it belong to?” This decoupling of object recognition and knowledge discovery aids in a deeper analysis of the capability boundaries and execution mechanisms of LVLMs.

In this paper, to comprehensively assess the factual knowledge of LVLMs, we present a **ChineseSimpleVQA** benchmark³, which consists of a dataset containing 2,200 high-quality questions across 56 topics, spanning from the humanities to science and engineering, as depicted in Figure 1. Specifically, the key distinguishing features of our proposed ChineseSimpleVQA are as follows:

- **Multi-hop:** Visual factuality inquiries are decomposed into two steps: object recognition and knowledge assessment. This multi-hop strategy allows us to analyze the capability boundaries and execution mechanisms of LVLMs.
- **Diverse:** ChineseSimpleVQA emphasizes the Chinese language and covers 8 major topics (i.e., “Nature”, “Sciences”, “Engineering”, “Humanities & Society”, “Modern Architecture”, “Ancient Architecture”, “Geography & Meteorological” and “Life Culture & Art”). These topics encompass 56 fine-grained subtopics.
- **High-quality:** We implement a rigorous pipeline for the benchmark construction, including automatic verification, difficulty filtering, and human verification.
- **Static:** To maintain the enduring quality of ChineseSimpleVQA, all reference answers will remain unchanged over time.
- **Easy to evaluate:** All of the questions and answers are in a short format for quick evaluation.

³<https://chinesesimplevqa.github.io/>
[ChineseSimpleVQA.github.io/](https://chinesesimplevqa.github.io/)

Furthermore, we have open-sourced all datasets and code, providing ready-to-run scripts to assist researchers in their endeavors.

Subsequently, we conducted a comprehensive evaluation of 35 advanced LVLMs, covering 8 major topics and 56 sub-topics. Here are some insightful findings:

- **ChineseSimpleVQA is challenging:** The closed-source model o1-preview achieved the best performance in terms of visual factuality, surpassing the top open-source models by approximately 20 points. This indicates that open-source LVLMs still have a long way to go in this area.
- **Larger models lead to better results:** Within the same series of models, a larger model size generally yields better results. This trend holds for nearly all series of models.
- **Larger models exhibit better calibration:** We observe that o1-preview is better calibrated than o1-mini, and GPT-4o is more calibrated than GPT-4o-mini.
- **More sampling increase performance:** Increasing the sampling attempts (Best-of-N) can improve the model’s performance, but this tends to stabilize after around 30 times.
- **See the world, but not discover knowledge:** Multi-hop questions demonstrate that even the LVLMs can correctly recognize objects, the final performance may still fall short.

2 Related Works

VQA Benchmarks: The problem of Visual Question Answering (VQA) was originally conceptualized as a Visual Turing Test (Geman et al., 2015). Addressing this test necessitates that computational systems replicate human faculties, including real-world visual recognition, comprehension

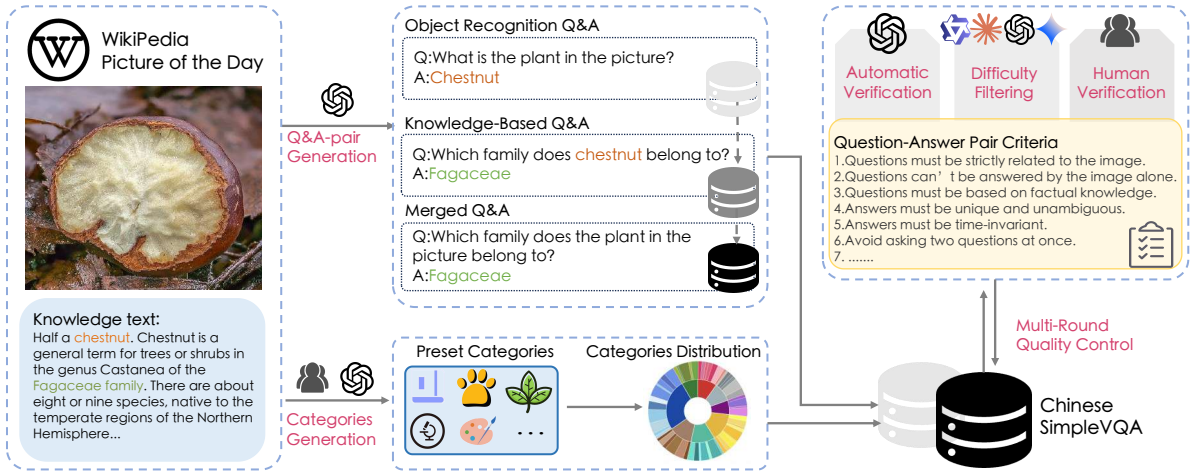


Figure 2: The production pipeline of ChineseSimpleVQA consists of automated verification, difficulty filtering, and human verification. This process generates multi-hop questions (i.e. object recognition Q&A and merged Q&A).

of language, rudimentary reasoning, and factual knowledge. Certain datasets emphasize the evaluation of perceptual and linguistic understanding capabilities (Antol et al., 2015; Malinowski and Fritz, 2014; Goyal et al., 2017; Liu et al., 2024, 2025; Yu et al., 2023; Li et al., 2024b), whereas others delve into the intricacies of complex reasoning processes (Johnson et al., 2017; Yue et al., 2024). Notwithstanding these advancements, assessing factuality remains a formidable challenge, an aspect that has received limited attention within the scope of VQA datasets.

Factuality Benchmarks: Several datasets also necessitate factual or commonsense knowledge (Marino et al., 2019; Wang et al., 2015, 2017; Zellers et al., 2019; Tan et al., 2024). For instance, OK-VQA (Marino et al., 2019) evaluates a model’s capability to answer questions using open-domain knowledge. Building upon this, A-OKVQA (Schwenk et al., 2022) integrates diverse external knowledge and reasoning capabilities, while S3VQA (Jain et al., 2021) introduces a dataset featuring questions that require object detection within images. Recently, factuality benchmarks have become increasingly important for LLMs: SimpleQA introduced by OpenAI (Wei et al., 2024) assesses short-form factuality, and a Chinese SimpleQA (He et al., 2024) is proposed by Alibaba Group.

Comparison to other knowledge-based VQA benchmarks: In Table 1, we present a comparative analysis of ChineseSimpleVQA alongside several mainstream knowledge-based VQA benchmarks. Our dataset represents the first Chinese mul-

timodal evaluation set designed to comprehensively assess factual abilities, thereby addressing a significant gap in the multimodal domain. Moreover, the multi-hop questioning in our dataset queries the image’s content and related facts, effectively evaluating the model’s object recognition and factual knowledge handling.

3 ChineseSimpleVQA

3.1 Overview

As shown in Figure 2, ChineseSimpleVQA is collected through a rigorous processing pipeline including automatic verification, difficulty filtering, and human verification to ensure high quality and appropriate difficulty levels. Ultimately, we obtained 1,100 images and 2,200 Q&A pairs covering 56 knowledge subtopics. Moreover, drawing from human cognitive processes, each image is designed to include two questions: the first question is solely for visual object recognition (i.e., seeing the world), while the second question requires both object recognition and factual knowledge (i.e., discovering knowledge). In other words, these two questions are used to evaluate the model’s ability to “recognize everything” and “understand encyclopedias”, respectively. Data examples can be found in Appendix C.2.

3.2 Knowledge Collection

To create visually oriented Q&A pairs based on encyclopedic knowledge, we collected a substantial number of images along with their corresponding knowledge texts from Wikipedia’s “Picture of the

Day” column⁴. Notably, since our goal is to build a Chinese benchmark, we specifically used the Chinese version of the column as our data source. In total, we gathered 6,546 image-text pairs, with the text having an average length of 723 characters.

3.3 Q&A Construction

Our benchmark focuses on both the visual recognition and knowledge discovery capabilities of LVLMs. Therefore, we construct Q&A pairs through three stages: (1). *Building Image Object Recognition Questions*: Identify the main objects within the images and formulate questions. (2). *Generating Knowledge-Based Questions on Image Objects*: Develop questions that investigate deeper into the knowledge related to the identified objects based on their knowledge texts. (3). *Merging the Previous Two Steps to Create Two-Hop Visual Knowledge Questions*: Integrate the visual recognition and knowledge-based questions to create more complex, two-hop questions that are grounded in visual content.

In the final version of our dataset, we retain the Q&A pairs from the **first and last stages** for each image to allow users to analyze the capability boundaries and execution mechanisms of LVLMs. Notably, the Q&A pairs at all stages are initially generated by prompting the LLM⁵ based on the criteria outlined in the following paragraphs. Notably, all Prompts are provided in Appendix C.4.

Criteria of Object Recognition Q&A. 1) Questions should be answerable using only information in the image, without needing external knowledge. For example, the question “Who designed the poster in the image?” is avoided as this requires extra information. 2) Each question must have a single, clear answer. For example, “What is in the image?” is unsuitable as it can refer to multiple entities. 3) Each question should be clear and specific, not divisible into separate inquiries.

Criteria of Knowledge-Based Q&A. 1) Questions should rely on factual knowledge, avoiding subjective opinions or personal views. 2) Answers must be clear and distinct, similar to the previous point. 3) Do not ask two questions at once, consistent with the previous guideline. 4) Answers should be timeless and not change over time, reflecting enduring facts. 5) Questions should be

⁴https://en.wikipedia.org/wiki/Wikipedia:Picture_of_the_day

⁵GPT-4o(0806)(OpenAI, 2023) is used for the generation of Q&A pairs, automatic quality verification, and evaluation.

Question Category	Proportion
Total Number	2,200
Nature	26.2%
Life, Culture & Art	14.5%
Geography & Meteorological	12.4%
Modern Architecture	11.8%
Human & Society	9.8%
Ancient Architecture	9.7%
Sciences	8.4%
Engineering	7.2%
Recognition Question Length	
- Maximum Length	28
- Minimum Length	6
- AVG Length	11.1
Final Question Length	
- Maximum Length	46
- Minimum Length	8
- AVG Length	16.1
Recognition Answer Length	
- Maximum Length	16
- Minimum Length	1
- AVG Length	5.2
Final Answer Length	
- Maximum Length	33
- Minimum Length	1
- AVG Length	4.5

Table 2: Dataset statistics.

challenging, and in-depth knowledge should be assessed, avoiding overly simple queries. 6) Keep questions concise by including only essential information for a clear answer. For example, instead of asking, “In which country is the Desen’ka train station located in Vinnytsia City?” ask, “In which country is the Desen’ka train station located?”

Criteria of Final Merged Q&A. 1) Questions must rely entirely on the answer to the first Q&A pair, meaning the next question cannot be answered correctly without the first answer. 2) Ensure that new questions are clear and concise, maintaining fluency and grammatical accuracy.

3.4 Quality Control

We design strict quality control processes, including automatic verification and human annotation, to identify and correct corresponding issues.

Automatic Verification. We developed a multi-round automated method to verify the quality of synthetic Q&A pairs, aiming to regenerate or filter out questions that do not meet the above criteria. Upon completion of the automatic verification, we retained a total of 5,009 images with 10,018 qualified Q&A pairs.

Difficulty Filtering. Subsequently, we filtered out simple samples to identify the knowledge boundaries of LLMs and enhance the difficulty of the benchmark. Specifically, if a question can be correctly answered by all of the four powerful models, i.e., GPT-4o (0806), Claude 3.5 Sonnet (Enis and Hopkins, 2024), Gemini 1.5 Pro (Team et al., 2024), and Qwen-VL-Max (Bai et al., 2023), it is considered simple and thus discarded. As a result of this process, a total of 3,058 images with 6,116 Q&A pairs were retained.

Human Verification. During this parse, a team of 23 annotators performed data verification and rewriting, while 6 engineers reviewed and selected the high-quality data.

Annotators must ensure question-answer pairs meet the standards by the following operations: 1) verifying and rewriting the questions or answers that do not meet the standards in Section 3.3; 2) replacing unqualified or unrepresentative images with new ones from online sources; and 3) verifying answer accuracy via search engines and Baidu Baike. Pairs that cannot be improved are discarded. Finally, manual annotations are compared with LLM-verified results, retaining only fully consistent Q&A pairs. This rigorous process ensures accuracy and adherence to established standards.

Data Desensitization. To mitigate potential security risks during evaluation, we submitted the final dataset to six security auditors for a comprehensive security review. Each sample of data was cross-checked by at least two auditors, and only the data that passed all scrutiny was retained.

After the above annotations, 2,411 images with 4,822 Q&A pairs were retained. Afterward, the algorithm experts carefully selected 1,100 images with 2,200 QA pairs as the final dataset.

3.5 Dataset Statistics

As shown in Table 2, ChineseSimpleVQA consists of 8 major topics and 56 subtopics, which are based on manually labeled tags. Among these topics, natural knowledge is the most common, comprising 26.2% of the total. Furthermore, since ChineseSimpleVQA primarily evaluates short-form answers, the average answer lengths for these two stages of questions are 5.2 words and 4.5 words, respectively, enabling a quicker evaluation.

4 Experiment

4.1 Setup

We maintain a consistent prompt format throughout all experiments, ensuring uniformity in our benchmark. The temperature and sampling parameters are set to the official default settings for each model. For evaluation, we use the “GPT-4o-0806” as a judge model.

4.2 Baseline Models

We evaluate **13 closed-source LVLMs** (i.e., o1-preview (0901), o1-mini (0901), GPT-4o (0806), GPT-4o-mini (0708)⁶, Gemini-1.5-pro-flash, Gemini-1.5-pro (Team et al., 2024), Gemini-2.0-flash⁷, Claude-3.5-sonnet⁸, Claude-3.5-sonnet, Doubao-1.5-vision-pro⁹, Doubao-vision-pro/lite, Qwen-VL-max (Bai et al., 2023)) and **22 open-source LVLMs** (i.e., Qwen2-VL series (Wang et al., 2024), Qwen2.5-VL series¹⁰, Deepseek-VL-2 series (Wu et al., 2024), LLaVA-one-vision-0.6B/72B (Li et al., 2024a), GLM-4V-9B¹¹, Llama-vision-11B/90B¹², InternVL series (Chen et al., 2024)).

4.3 Evaluation Metrics

We use the following metrics to evaluate ChineseSimpleVQA: (1) **Correctness (CO)**: The predicted answer to the final question should fully encompass the reference answer without contradictions. (2) **Not attempted answers (NA)**: The predicted answer to the final question is not related to the reference answer without any contradictions. (3) **Incorrectness (IN)**: The predicted answer to the final question contradicts the reference answer, as does the answer to the first question. (4) **Correctness among Given Attempted (CGA)**: The proportion of questions accurately answered among those attempted. (5) **F-score**: The harmonic mean of Correct and Correct Given Attempted.

⁶<https://openai.com/index/>

⁷<https://deepmind.google/technologies/gemini/flash/>

⁸<https://www.anthropic.com/news/claude-3-5-sonnet>

⁹<https://team.doubao.com/en/direction/vision>

¹⁰<https://help.aliyun.com/zh/model-studio/developer-reference/use-qwen-by-calling-api>

¹¹<https://open.bigmodel.cn/dev/api/normal-model/glm-4v>

¹²<https://ollama.com/library/llama3.2-vision>

Model	Overall results of Merged Q&A					Overall results of Recognition Q&A				
	CO	IN↓	NA↓	CGA	F-score	CO	IN↓	NA↓	CGA	F-score
<i>Closed-Source Large Vision Language Models</i>										
o1-preview	68.8	24.6	6.5	73.6	71.1	79.1	13.6	7.3	85.3	82.1
o1-mini	52.7	38.1	9.2	58.1	55.3	64.8	24.9	10.3	72.2	68.3
GPT-4o	59.1	35.5	5.4	62.4	60.7	77.5	15.5	7.0	83.4	80.4
GPT-4o-mini	51.0	43.6	5.4	53.9	52.4	70.8	23.1	6.1	75.4	73.0
Gemini-2.0-flash	64.5	29.5	5.9	68.6	66.5	76.7	19.6	3.7	79.7	78.2
Gemini-1.5-pro-flash	56.5	34.6	8.8	62.0	59.2	70.3	25.9	3.8	73.1	71.6
Gemini-1.5-pro	66.2	31.4	2.5	67.8	67.0	77.5	20.7	1.7	78.9	78.2
Claude-3.5-sonnet2	63.8	30.6	5.5	67.6	65.6	77.6	17.2	5.2	81.9	79.7
Claude-3.5-sonnet	59.5	26.4	14.2	69.4	64.0	69.5	20.2	10.3	77.5	73.3
Qwen-VL-max	56.5	39.6	3.8	58.8	57.6	72.9	24.6	2.5	74.7	73.8
Doubao-1.5-vision-pro	53.2	36.4	10.5	59.4	56.1	50.4	39.1	10.5	56.3	53.2
Doubao-vision-pro	52.0	43.2	4.8	54.6	53.3	44.0	51.9	4.1	45.9	44.9
Doubao-vision-lite	31.9	30.0	38.1	51.5	39.4	30.6	43.5	25.8	41.3	35.2
<i>Open-Source Large Vision Language Models</i>										
Deepseek-VL2	33.4	60.3	6.4	35.6	34.5	32.0	58.5	9.5	35.4	33.6
Deepseek-VL2-small	37.0	58.5	4.5	38.8	37.9	33.1	58.4	8.5	36.2	34.6
Deepseek-VL2-tiny	22.5	72.3	5.2	23.8	23.1	23.0	69.9	7.1	24.8	23.8
LLaVA-onevision-72B	34.7	60.6	4.6	36.4	35.6	29.5	59.6	10.9	33.1	31.2
LLaVA-onevision-7B	18.0	57.7	24.3	23.8	20.5	14.5	61.2	24.4	19.1	16.5
LLaVA-onevision-0.5B	7.9	82.3	9.8	8.8	8.3	8.5	70.6	20.8	10.8	9.5
Qwen2.5-VL-72B	49.0	42.8	8.2	53.4	51.1	45.7	48.5	5.7	48.5	47.1
Qwen2.5-VL-7B	39.5	54.2	6.4	42.1	40.8	41.3	51.2	7.5	44.6	42.9
Qwen2.5-VL-3B	31.4	61.3	7.4	33.9	32.6	32.9	62.0	5.1	34.7	33.8
Qwen2-VL-72B	50.6	46.6	2.7	52.1	51.3	48.0	49.1	2.9	49.4	48.7
Qwen2-VL-7B	38.2	60.3	1.5	38.8	38.5	39.5	57.1	3.4	40.9	40.2
Qwen2-VL-2B	29.0	66.5	4.5	30.4	29.7	34.2	63.1	2.7	35.1	34.7
GLM-4v	34.8	51.7	13.5	40.2	37.3	37.1	54.5	8.4	40.5	38.7
Llama-vision-90B	46.2	51.6	2.2	47.2	46.7	37.2	57.5	5.3	39.3	38.2
Llama-vision-11B	29.4	66.4	4.3	30.7	30.0	27.4	64.3	8.4	29.9	28.6
InterVL2.5-78B	41.1	51.7	7.2	44.3	42.6	36.2	56.7	7.1	38.9	37.5
InterVL2.5-38B	35.5	54.3	10.2	39.6	37.5	30.0	56.4	13.6	34.7	32.2
InterVL2.5-26B	32.4	52.9	14.7	38.0	34.9	29.9	55.3	14.8	35.1	32.3
InterVL2.5-8B	28.4	60.5	11.1	31.9	30.0	24.5	63.5	12.1	27.8	26.0
InterVL2.5-4B	24.5	73.7	1.8	24.9	24.7	22.2	72.5	5.4	23.4	22.8
InterVL2.5-2B	13.9	73.9	12.2	15.8	14.8	16.8	74.0	9.2	18.5	17.6
InterVL2.5-1B	14.7	70.7	14.5	17.2	15.9	16.8	69.3	13.9	19.5	18.1

Table 3: Performance comparison of closed source and open source LVLMs on multi-hop QAs (i.e. Merged Q&A and Recognition Q&A). For metrics, **CO**, **NA**, **IN**, and **CGA** denote “Correct”, “Not attempted”, “Incorrect”, and “Correct given attempted”, respectively. The highest scores among models in each metric are highlighted in **green**.

4.4 Main Results

We present the performance results of various LVLMs on the ChineseSimpleVQA dataset in Table 3. The evaluation metrics are introduced in Section 4.3. We report the models’ performance on recognition and final merged questions, and summarize the following findings:

- **o1-preview is the best performer** for both recognition and final questions among all models. The next best are Gemini-2.0-flash and Gemini-1.5-pro, which follow closely behind. For Chinese ability, Qwen family performs the best.
- **Performance on final questions is positively correlated with the performance on recognition questions**, though accuracy is typically

higher for recognition questions. This suggests that models can identify objects but sometimes fail to grasp deeper knowledge. Some models (e.g. Deepseek-VL2) perform better on final questions because correct and incorrect object predictions can point to the same final answer. For instance, both “*Misumena vatia*” and “*Thomisidae*” belong to the order “*Araneae*.”¹³.

- **The larger the model size within the same series of models, the better the result.** Taking the Qwen2-VL series as an example, when the model size increases from 3B to 72B, the accuracy on the final questions increases from 29.0% to 50.6%. More detailed analysis can be found

¹³Examples are the scientific names of two spider species.

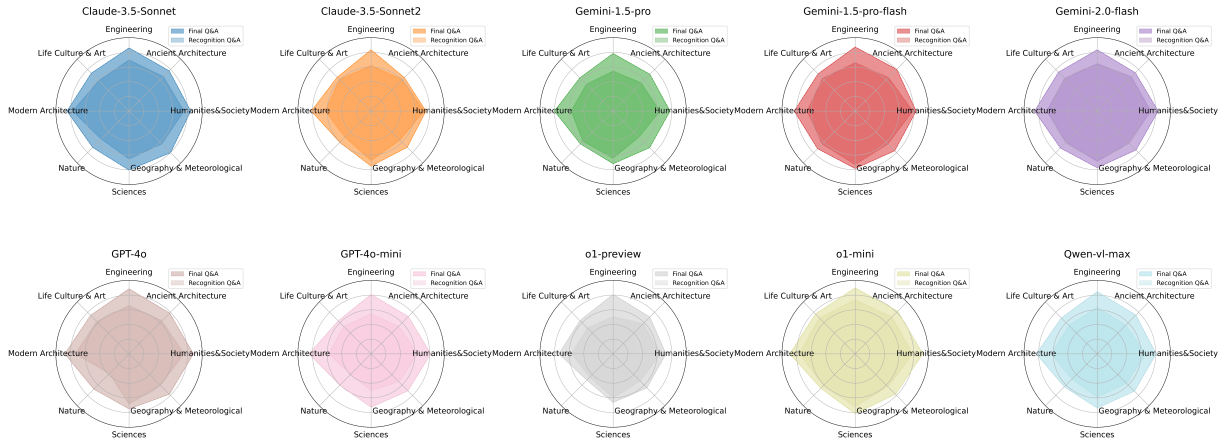


Figure 3: Correctness(CO) metric for eight topics. To conserve space, only the top 10 models are displayed.

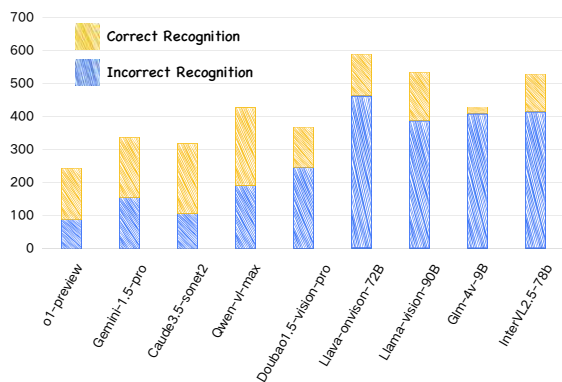


Figure 4: The number of incorrectly answered final questions, and the number of samples where the image object was correctly recognized.

in Appendix B.1.

- **The IN rate of the models is much higher than their NA rate** (in addition to Doubao-vision-lite), indicating that the models are more likely to confidently provide incorrect factual information. Mitigating this hallucination problem remains a significant challenge. In Section 4.6, we will conduct a detailed analysis focusing on model confidence.

4.5 Performance on Different Topics

We select the top 10 models with the highest CO rate on the final questions and evaluated their performance across different topics. The results are shown in Figure 3. For all topics, the trends in model performance for both recognition and final questions align with the conclusions presented in the previous section. It is obvious that these models are more adept at topics such as modern architecture, engineering, and science. However, their performance is slightly lower on natural topics, as these questions often need more specialized knowledge than common sense can provide.

From this figure, we can observe the differences in the model’s performance on recognition questions versus final questions across different topics. For example, Claude-3.5-Sonnet exhibits a small performance gap in the topic of modern architecture. This indicates that once the model accurately identifies the object of an image, it is more likely to provide accurate responses based on relevant knowledge. Conversely, if the model demonstrates a large performance gap on a particular topic, it suggests that the model’s knowledge of that topic is limited.

4.6 Further Analysis

Analysis of Knowledge Scope. In the design of ChineseSimpleQA, the final question requires deeper, extended knowledge beyond the recognition question. To explore at which stage the real knowledge limitation arises when models fail to answer the final question, we collect examples where different models provided incorrect answers to the final questions and conduct analysis by calculating the percentage of corresponding recognition questions that were answered correctly. The results are presented in Figure 4.

For o1-preview, Gemini-1.5-pro, Claude3.5-sonnet2 and Qwen-vl-max, the error rate caused by image object recognition is less than 50%. This shows that for these models, the inability to correctly answer the final question is mainly limited by the lack of more complex extended knowledge. For the other models, it is impossible to correctly recognize the image object from the initial stage, especially GLM-4v-9B.

Analysis of Calibration. The confidence level (%) of a perfectly calibrated model should precisely match the accuracy of its predictions. From this

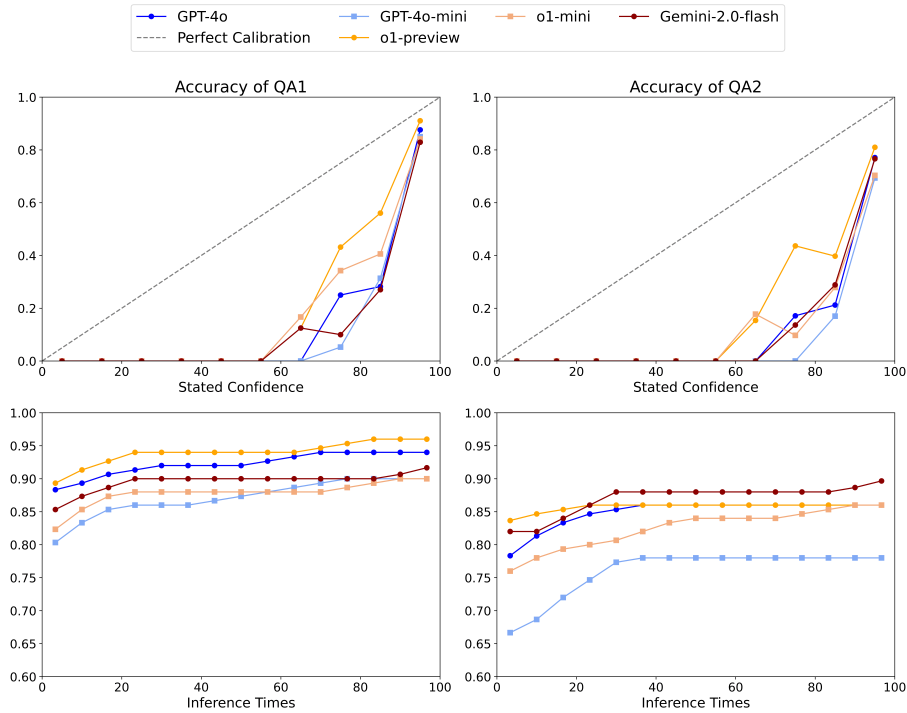


Figure 5: Up: Calibration of LLMs based on their stated confidence for Recognition and final Q&A. Down: Improvement in accuracy with increased test-time compute using Best-of-N for Recognition and final Q&A.

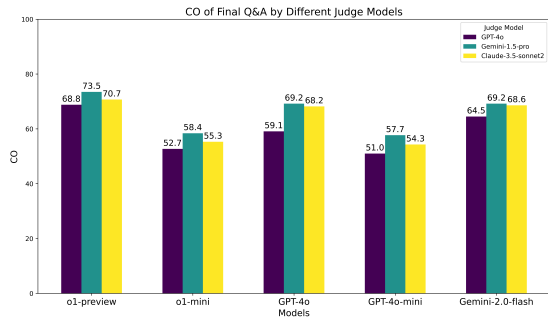


Figure 6: Robustness of Judge Models of Final Q&A.

perspective, we prompt the models to provide their confidence levels (ranging from 0 to 100) while answering questions. In Figure 5, we present the accuracy of the models' answers at different confidence intervals (with an interval size of 10).

The results illustrate the alignment performance of the tested models. Specifically, o1-preview exhibits the best performance on both types of questions. While o1-mini ranks second in recognition questions, closely following o1-preview, it performs slightly worse on final questions. Overall, most models fall below the perfect alignment line, indicating that they tend to be overconfident, even when their answers are incorrect.

Analysis of Inference Attempt. Furthermore, we analyze the changes in the accuracy of the model's answers as the inference attempts in-

creased. Specifically, we sample 100 instances from each of the two types of questions in ChineseSimpleVQA. For each instance, the model is asked to independently generate an answer 100 times. For the results of multiple inferences, we use the Best-of-N method to obtain the accuracy of the model responses.

Figure 5 shows that within the range of 1-30 inference times, the accuracy of the tested model increases with the increase of inference times. However, when the number of inferences is greater than 30, the performance of the model tends to be stable. Beyond this point, the model does not acquire significantly more accurate knowledge even with additional attempts. These findings can serve as a reference for determining the boundaries of the model's knowledge capabilities.

5 Robustness of Judge Models

Moreover, we select two models, i.e., Gemini-1.5-pro and Claude-3.5-sonnet2, as the judge models and compare the evaluation results of GPT-4o-0806. The results are shown in Figure 6. It can be seen that when using different evaluation models, the ranking of the tested models remains consistent. This further proves the robustness of the evaluation method of ChineseSimpleVQA.

6 Conclusion

We introduce the first short-form VQA benchmark, named Chinese SimpleVQA, to evaluate the visual factuality abilities of existing LVLMs. This benchmark is characterized by important features such as multi-hop, diverse, high-quality, static, and easy-to-evaluate. Based on Chinese SimpleVQA, we conduct a comprehensive evaluation of the performance of nearly 30 models regarding visual factuality, providing a detailed analysis of the capability boundaries to demonstrate the advantages and necessity of ChineseSimpleVQA.

7 Limitation

From the perspective of topic distribution, the coverage of topics discussed in this paper remains relatively limited. A broader topic coverage would contribute to a more comprehensive assessment of visual knowledge. In terms of data volume, after rigorous quality filtering and difficulty balancing, the final version of our dataset represents only a tiny fraction of the overall data. Therefore, expanding the dataset will be a key focus of our future efforts. Furthermore, as highlighted by this work (Jiang et al., 2025), datasets to evaluate the model against potential jailbreak attacks is crucial. Finally, regarding model structure, improving the visual factuality of the model to enhance its performance on ChineseSimpleVQA will become increasingly important. However, both performance improvement and security aspects may already be beyond the scope of this paper.

8 Ethic Statement

We respect intellectual property rights and comply with relevant laws and regulations. The source of our dataset is publicly available, and we have taken careful measures to ensure that the dataset does not contain any personal sensitive information. In addition, our work is only for research purposes, not for commercial purposes.

References

Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou,

and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 1(2):3.

Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.

Maxim Enis and Mark Hopkins. 2024. From llm to nmt: Advancing low-resource machine translation with claude. *arXiv preprint arXiv:2404.13813*.

Chaoyou Fu, Yi-Fan Zhang, Shukang Yin, Bo Li, Xinyu Fang, Sirui Zhao, Haodong Duan, Xing Sun, Ziwei Liu, Liang Wang, et al. 2024. Mme-survey: A comprehensive survey on evaluation of multimodal llms. *arXiv preprint arXiv:2411.15296*.

Donald Geman, Stuart Geman, Neil Hallonquist, and Laurent Younes. 2015. Visual turing test for computer vision systems. *Proceedings of the National Academy of Sciences*, 112(12):3618–3623.

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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.

Yancheng He, Shilong Li, Jiaheng Liu, Yingshui Tan, Weixun Wang, Hui Huang, Xingyuan Bu, Hangyu Guo, Chengwei Hu, Boren Zheng, et al. 2024. Chinese simpleqa: A chinese factuality evaluation for large language models. *arXiv preprint arXiv:2411.07140*.

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.

Yilei Jiang, Xinyan Gao, Tianshuo Peng, Yingshui Tan, Xiaoyong Zhu, Bo Zheng, and Xiangyu Yue. 2025. Hiddendetector: Detecting jailbreak attacks against large vision-language models via monitoring hidden states. *arXiv preprint arXiv:2502.14744*.

Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910.

- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. 2024a. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. 2024b. Seed-bench: Benchmarking multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13299–13308.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2025. Mm-bench: Is your multi-modal model an all-around player? In *European conference on computer vision*, pages 216–233. Springer.
- Mateusz Malinowski and Mario Fritz. 2014. A multi-world approach to question answering about real-world scenes based on uncertain input. *Advances in neural information processing systems*, 27.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204.
- R OpenAI. 2023. Gpt-4 technical report. arxiv 2303.08774. *View in Article*, 2(5).
- 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.
- Yingshui Tan, Boren Zheng, Baihui Zheng, Kerui Cao, Huiyun Jing, Jincheng Wei, Jiaheng Liu, Yancheng He, Wenbo Su, Xiangyong Zhu, et al. 2024. Chinese safetyvqa: A safety short-form factuality benchmark for large language models. *arXiv preprint arXiv:2412.15265*.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- 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.
- Peng Wang, Qi Wu, Chunhua Shen, Anton van den Hengel, and Anthony Dick. 2015. Explicit knowledge-based reasoning for visual question answering. *arXiv preprint arXiv:1511.02570*.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. [Measuring short-form factuality in large language models](#). *ArXiv*, abs/2411.04368.
- Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, et al. 2024. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6720–6731.

A Overview of Appendix

We have over 10 pages of this appendix, comprising the following subsections for the convenience of readers:

More additional experiments

- **Appendix B.1:** Additional experimental analysis of model scale.
- **Appendix B.2:** Additional experimental analysis of performance comparison among models.
- **Appendix B.3:** Additional experimental analysis of rankings on ChineseSimpleVQA.
- **Appendix B.4:** Additional experimental analysis between popularity and difficulty.

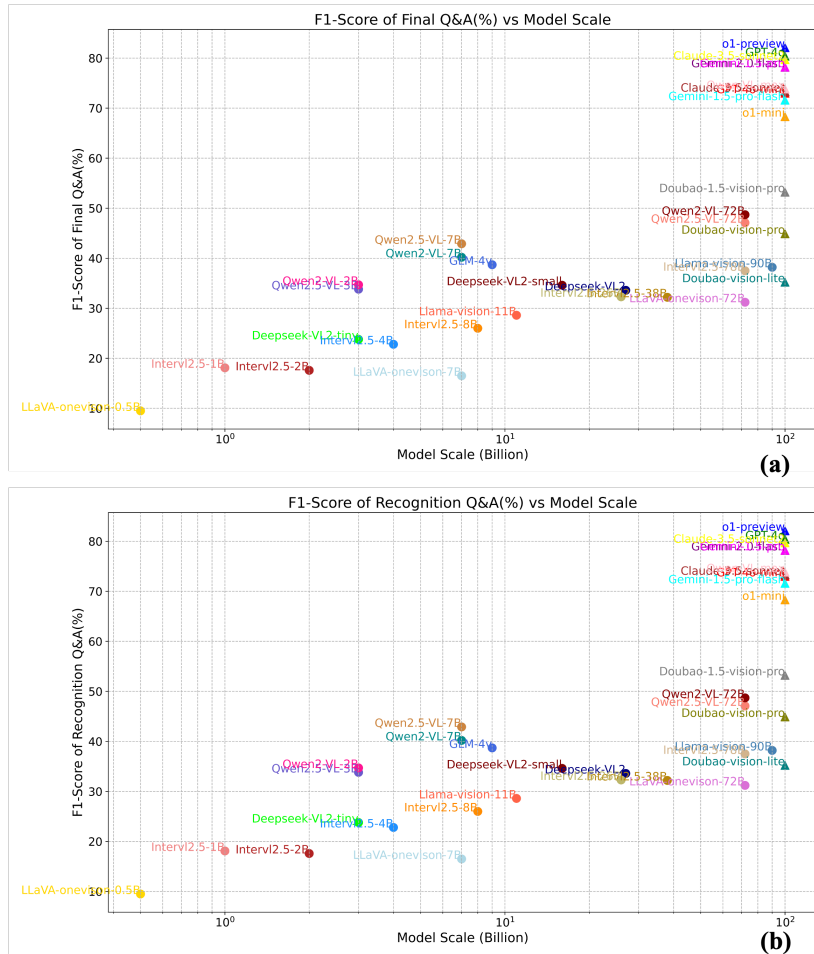


Figure 7: Relationship between model scale (in billion parameters) and F-score on Recognition and final Q&A.



Figure 8: Example of model responses of where Deepseek-vl2-small's answer is correct but Deepseek-VL2's answer is wrong.

- **Appendix B.5:** Additional analysis of Human Calibration.
- **Appendix B.6:** Additional experimental comparison with other Datasets.
- **Appendix B.7:** Additional experimental compar-

ison with Knowledge-Based Q&A

- **Appendix B.8:** More results.

More visualization and cases

- **Appendix C.1:** Detailed question distribution.
- **Appendix C.2:** Dataset visualization of our ChineseSimpleVQA.
- **Appendix C.3:** Visualization of performance comparison among models.
- **Appendix C.4:** Details of all prompts in the dataset generation and validation.

We hope that our efforts will serve as a source of inspiration for more readers!

B More Experiments

B.1 F-score v.s. Model Scale

In Figure 7, we illustrate the trend of F-score as the model size increases. Generally, there is a positive correlation between these two factors. How-

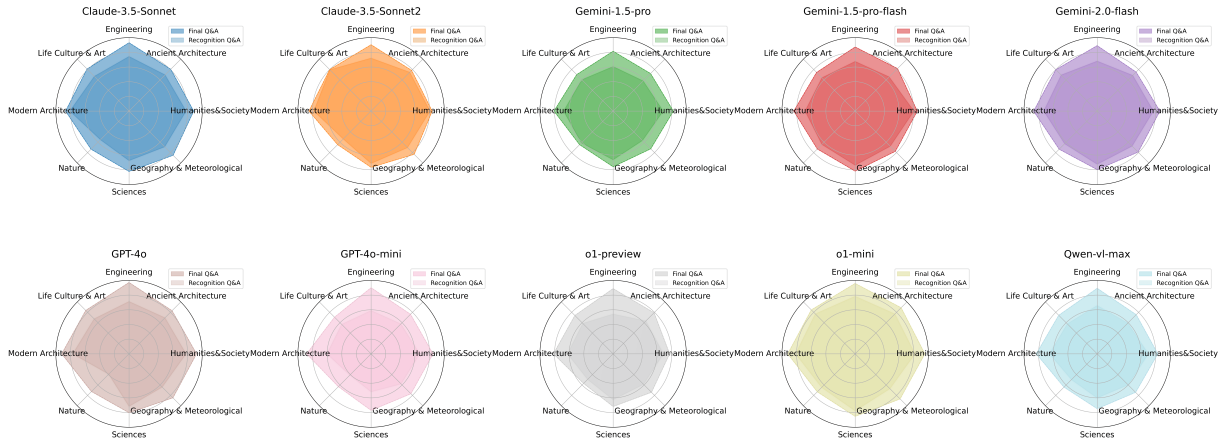


Figure 9: CGA of the top 10 models for eight topics.

ever, this relationship weakens in the intermediate range, especially for vision recognition tasks, where larger models may perform worse. This observation contrasts with findings from ChineseSimpleQA, which analyzed large language models (LLMs). For large vision-language models (LVMs), while increased size suggests greater knowledge capacity, accurate image object recognition hinges more on the quality of image-text alignment training, critically affecting factual accuracy. Additionally, Table 3 shows the performance of same-series models typically improves with size, but Deepseek-VL and Deepseek-VL-small deviate from this pattern. Figure 8 provides examples where the smaller model correctly answers the questions that the larger Deepseek-VL fails, highlighting the discrepancy between them.

B.2 CGA Results on 8 Topics

In this section, we present more results of Correctness among Given Attempted (CGA). In principle, the models are encouraged to do not attempt to answer when they are unable to determine the knowledge. Therefore, we further show the accuracy of the models' responses on 8 topics after excluding the questions that the models do not attempt to answer, namely CGA. In general, GCA has higher indicators than CO. The detail results are shown in Figure 9.

B.3 Rankings on ChineseSimpleVQA

We also compare the ranking differences of various models on the recognition questions and the final questions. The rankings are shown in Figure 12.

B.4 Analysis between Popularity and Difficulty

We select the first 100 samples from our dataset and monitor their page views on Wikipedia over the past 30 days to assess their popularity. We categorize the popularity into three ranges based on the number of views: (0-10), (10-50), and (50, +inf). We then evaluate the experimental performance of the following 10 closed-source models across these different levels of popularity as shown in Table 4. The results show that models generally perform better with higher popularity data, particularly in Recognition Q&A tasks. For data with over 50 views in the past 30 days, most models can achieve an accuracy rate of 100%.

B.5 Analysis of Human Calibration

We manually extract 110 samples of the main results from both the merged and recognition Q&A for verification and find that agreement with human assessments reached 95.5% and 97.3%, respectively. This proves the validity of using LLM-as-judge for evaluation.

B.6 Comparison with other Datasets

In order to explore whether models that perform well in other datasets continue to perform well in this dataset. We compare the evaluation performance of models on our ChineseSimpleVQA and text-based (no vision) short-form factuality benchmarks—ChineseSimpleQA (He et al., 2024). The experimental results are shown in Table 5: As can be seen from the CO metric, the o1-preview model achieves the best results in the merged Q&A and Recognition Q&A of ChineseSimpleVQA, as well as in ChineseSimpleQA. In contrast, the gpt-



Figure 10: CO of all other different models for eight topics.

4o-mini model performs the worst in the merged Q&A of ChineseSimpleVQA and in ChineseSimpleQA, demonstrating consistency. For the Recognition Q&A of ChineseSimpleVQA, the model that performed the worst was o1-mini. This indicates that our benchmark is a qualified factuality evaluation set that can also assess large vision language models’ image recognition capabilities. It holds broad evaluative applications for LVLMS.

B.7 Performance Comparison including Knowledge-Based Q&A

To assist us in analysis, we conduct an evaluation on text-only Knowledge-Based Q&A. The experimental results are shown in Table 6. From the CO metric perspective, we can find that the o1 preview still shows the best performance, which is consistent with the other two Q&A tasks. Additionally, the performance of Knowledge-Based Q&A is similar to Merged Q&A, and for most models, Knowledge-Based Q&A performs better. This indicates that for LVLMS, it is more challenging to simultaneously recognize images as well as answer Knowledge-Based questions, and compared to image recognition, Knowledge-Based Q&A is the more difficult part.

B.8 More Results

We present the CO and CGA results of all models in Figure 10 and 11, respectively.

C More Details of ChineseSimpleQA

C.1 Detailed Question Distribution

In Table 7, we further give the detailed distribution (number of questions) of ChineseSimpleVQA under 56 sub-topics. In addition, we counted the average length of recognition questions and final questions under each topic.

C.2 ChineseSimpleVQA Examples

We provide examples from the 8 topics in ChineseSimpleVQA to illustrate our benchmark’s focus on multi-hop QAs. Figures 13 show that our benchmark covers various knowledge topics and uses multi-hop questioning to better evaluate different models’ capabilities.

C.3 Model Response Examples

In this section, we visualize the answers provided by different models to the same set of questions, to further investigate the capability boundaries of these models. As illustrated in Figures 14 through 16, similar to the conclusion presented in Section 4.4, almost all models fail to achieve satisfactory performance on our ChineseSimpleVQA. Addition-

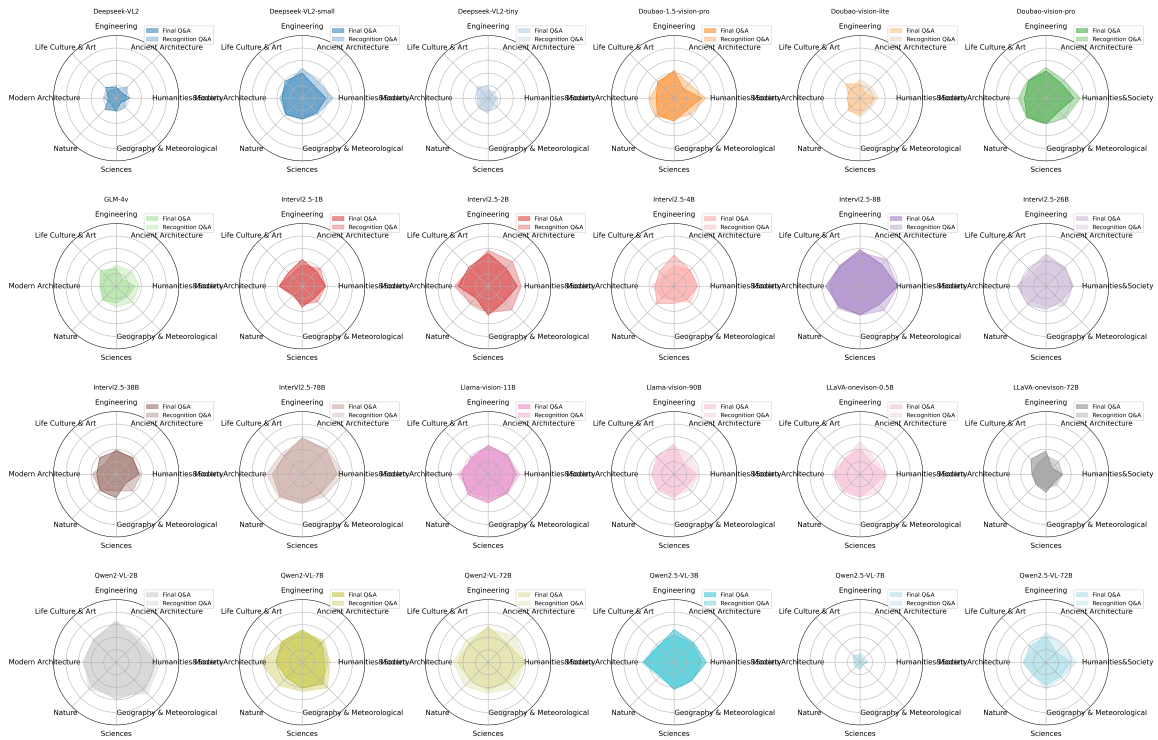


Figure 11: CGA of all other different models for eight topics.

Model	CO of Merged Q&A			CO of Recognition Q&A		
	0 - 10	10 - 50	50 - +inf	0 - 10	10 - 50	50 - +inf
o1-preview	68.00	84.62	80.00	84.00	90.38	100.00
o1-mini	68.00	73.08	60.00	72.00	76.92	80.00
GPT-4o	72.00	73.08	80.00	76.00	86.54	100.00
GPT-4o-mini	48.00	69.23	80.00	68.00	92.31	80.00
Gemini-2.0-flash	72.00	78.85	60.00	84.00	90.38	80.00
Gemini-1.5-pro-flash	80.00	71.15	70.00	88.00	78.85	80.00
Gemini-1.5-pro	80.00	76.92	100.00	92.00	86.54	100.00
Claude-3.5-sonnet2	84.00	73.08	80.00	88.00	84.62	100.00
Claude-3.5-sonnet	68.00	67.31	80.00	92.00	78.85	90.00
Qwen-VL-max	88.00	73.08	90.00	80.00	78.85	90.00

Table 4: Different popularity’s difficulty of closed source LVLMs on multi-hop QAs (i.e. Merged Q&A and Recognition Q&A).

ally, the differences in performance between the various models are significant.

C.4 Prompts of Dataset Generation and Validation

The generation, validation, and evaluation of question-answer pairs all use OpenAI’s gpt-4o-0806. We provide the specific prompts in the section. The data generation prompt are shown in Figure 17, Figure 18 and Figure 19; The validation prompt is shown in Figure 20; The evaluation prompt for ChineseSimpleVQA is shown in Figure 21. All construction and evaluation processes of Chinese SimpleVQA use the Chinese version of these prompts.

Model	Overall results of Merged Q&A					Overall results of Recognition Q&A					Overall results of ChineseSimpleQA				
	CO	IN↓	NA↓	CGA	F-score	CO	IN↓	NA↓	CGA	F-score	CO	IN↓	NA↓	CGA	F-score
o1-preview	68.8	24.6	6.5	73.6	71.1	79.1	13.6	7.3	85.3	82.1	63.8	24.0	12.2	72.7	67.9
o1-mini	52.7	38.1	9.2	58.1	55.3	64.8	24.9	10.3	72.2	68.3	39.5	39.9	20.6	49.7	44.1
GPT-4o	59.1	35.5	5.4	62.4	60.7	77.5	15.5	7.0	83.4	80.4	59.3	39.3	1.4	60.1	59.7
GPT-4o-mini	51.0	43.6	5.4	53.9	52.4	70.8	23.1	6.1	75.4	73.0	37.6	61.5	0.9	37.9	37.8
Gemini-1.5-pro	66.2	31.4	2.5	67.8	67.0	77.5	20.7	1.7	78.9	78.2	37.4	37.6	8.0	59.1	56.7
Claude-3.5-sonnet	59.5	26.4	14.2	69.4	64.0	69.5	20.2	10.3	77.5	73.3	45.2	26.4	27.4	63.6	53.5

Table 5: Performance comparison between ChineseSimpleVQA and ChineseSimpleQA (He et al., 2024). For metrics, CO, NA, IN, and CGA denote “Correct”, “Not attempted”, “Incorrect”, and “Correct given attempted”, respectively. The highest scores among models in each metric are highlighted in **green**.

Model	Overall results of Merged Q&A					Overall results of Recognition Q&A					Overall results of Knowledge-Based Q&A				
	CO	IN↓	NA↓	CGA	F-score	CO	IN↓	NA↓	CGA	F-score	CO	IN↓	NA↓	CGA	F-score
o1-preview	68.8	24.6	6.5	73.6	71.1	79.1	13.6	7.3	85.3	82.1	72.9	21.0	6.6	77.5	74.8
o1-mini	52.7	38.1	9.2	58.1	55.3	64.8	24.9	10.3	72.2	68.3	50.9	34.2	14.9	59.8	55.0
GPT-4o	59.1	35.5	5.4	62.4	60.7	77.5	15.5	7.0	83.4	80.4	66.5	31.6	1.8	67.8	67.2
GPT-4o-mini	51.0	43.6	5.4	53.9	52.4	70.8	23.1	6.1	75.4	73.0	47.6	51.5	1.2	47.9	47.2
Gemini-2.0-flash	64.5	29.5	5.9	68.6	66.5	76.7	19.6	3.7	79.7	78.2	69.4	30.3	0.5	69.6	69.4
Gemini-1.5-pro-flash	56.5	34.6	8.8	62.0	59.2	70.3	25.9	3.8	73.1	71.6	57.4	38.5	4.4	59.8	58.5
Gemini-1.5-pro	66.2	31.4	2.5	67.8	67.0	77.5	20.7	1.7	78.9	78.2	66.5	32.5	1.5	67.0	66.5
Claude-3.5-sonnet2	63.8	30.6	5.5	67.6	65.6	77.6	17.2	5.2	81.9	79.7	62.4	30.0	7.9	67.4	64.6
Claude-3.5-sonnet	59.5	26.4	14.2	69.4	64.0	69.5	20.2	10.3	77.5	73.3	56.5	26.7	16.7	67.9	61.7
Qwen-VL-max	56.5	39.6	3.8	58.8	57.6	72.9	24.6	2.5	74.7	73.8	58.0	39.5	1.7	59.9	59.3

Table 6: Performance comparison of closed source and open source LVLMs on multi-hop QAs (i.e. Merged Q&A, Knowledge-Based Q&A and Recognition Q&A). For metrics, CO, NA, IN, and CGA denote “Correct”, “Not attempted”, “Incorrect”, and “Correct given attempted”, respectively. The highest scores among models in each metric are highlighted in **green**.

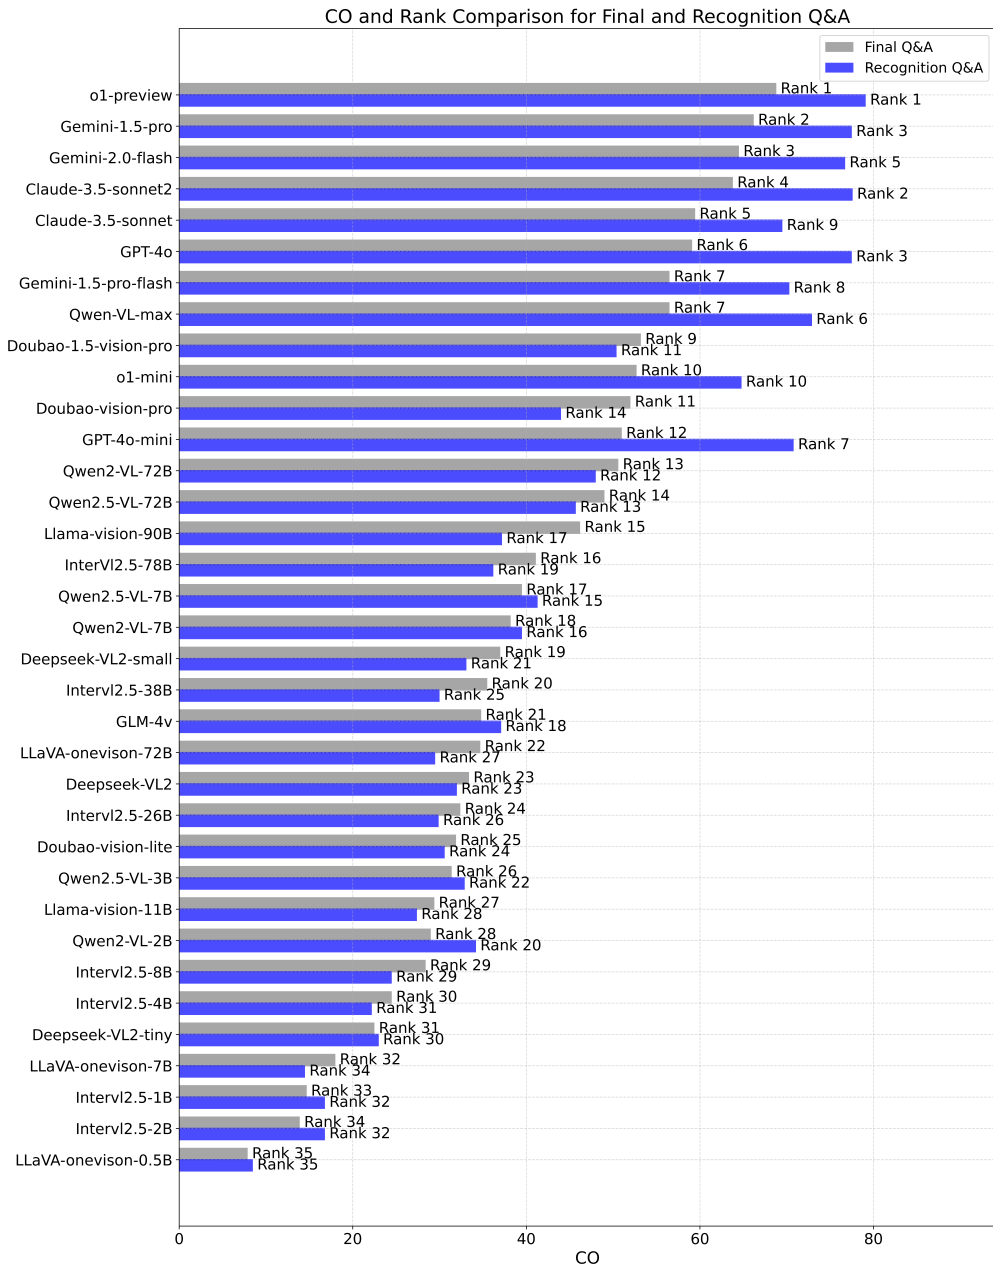


Figure 12: Rankings of different models on ChineseSimpleVQA.

Table 7: Avg. question length of all subtopics. "ChemE." represents "Chemical Engineering".

Primary Category	Secondary Category	Count	Q1 avg. length	Q2 avg. length
Geography & Meteorological	Meteorology & Climate (大气、气象与气候)	10	13.5	22.1
	Hydrology (水文)	25	11.1	14.5
	Human Geography (人文地理)	38	12.5	16.2
	Mountains (山脉)	19	10.8	14.7
	Landforms & Topography (地貌与地形)	30	11.9	16.8
	Geology (地质)	7	12.3	19.7
	Other (其他)	7	13.0	22.9
Modern Architecture	Commercial Buildings (商业建筑)	14	12.1	16.2
	Public Cultural Buildings (公共文化建筑)	39	11.1	16.5
	Funerary and Memorial Buildings (殡葬与纪念建筑)	10	12.9	16.4
	Other (其他)	13	10.9	15.2
	Religious Buildings (宗教建筑)	21	10.4	15.0
	Transportation Buildings (交通建筑)	34	11.2	15.6
Nature	Insects (昆虫)	49	9.4	11.8
	Reptiles (爬行类)	30	9.6	12.6
	Birds (鸟类)	37	9.5	14.4
	Mammals (哺乳类)	39	9.8	14.2
	Angiosperms (被子植物)	43	10.3	13.3
	Arthropods (non-insect) (节肢动物(非昆虫))	19	9.9	12.9
	Gymnosperms (裸子植物)	6	11.7	13.7
	Fish (鱼类)	21	9.1	11.2
	Amphibians (两栖类)	11	9.4	13.6
	Other (其他)	12	11.3	14.0
	Mollusks (软体动物)	21	10.1	11.2
Life Culture & Art	Sculpture (雕塑)	23	11.6	14.3
	Cuisine & Festivals (饮食与节庆)	14	11.1	20.4
	Film & Television (影视剧)	8	10.8	20.0
	Sports (体育)	23	12.6	22.8
	Other (其他)	18	12.0	17.2
	Painting (绘画)	38	11.9	16.0
	Literature & Mythology (文学与神话)	11	11.8	20.1
	Music (音乐)	17	9.3	15.8
	Handicrafts (手工艺)	7	13.3	19.9
Humanities & Society	Classical & Medieval History (古典与中世纪史)	14	10.6	17.6
	Contemporary History (近代史)	25	10.0	17.4
	Modern History (现代史)	20	12.1	19.5
	Religious Studies (宗教学)	10	12.8	20.3
	Politics & Economics (政治学与经济学)	15	11.6	18.1
	Other (其他)	9	14.0	20.9
	Ancient History (历史学)	15	12.5	21.8
Sciences	Biology (生物学)	26	10.0	14.1
	Physics (物理学)	9	13.3	20.9
	Astronomy (天文学)	39	11.7	16.7
	Medicine & Chemistry (医学与化学)	13	12.1	19.7
	Mathematics (数学)	5	12.6	20.2
Ancient Architecture	Residences & Palaces (住宅与宫殿建筑)	13	11.7	14.5
	Funerary and Memorial Buildings (墓葬与遗址)	13	11.4	16.8
	Public Cultural Buildings (公共文化建筑)	11	10.6	17.1
	Transportation Buildings (交通建筑)	7	10.4	14.9
	Religious Buildings (宗教建筑)	38	10.9	15.4
	Other (其他)	25	10.7	16.2
Engineering	Military, Materials, Energy and ChemE. (军事、材料、能源、化工)	11	11.4	19.7
	Mechanical Engineering (机械)	30	12.8	17.4
	Aerospace (航空航天)	30	12.2	18.1
	Electronic Information (电子信息)	8	10.6	19.3

Topic	Image	Recognition Q&A	Final Q&A
自然 (Nature)		<p>Q: 图片中所示的蓝色动物是什么? A: 蓝海星</p> <p>Q: What is the blue animal shown in the picture? A: Linckia laevigata</p>	<p>Q: 图片中所示的蓝色动物属于哪个门? A: 棘皮动物门</p> <p>Q: To which phylum does the blue animal shown in the picture belong? A: Echinodermata</p>
科学 (Sciences)		<p>Q: 图片中展示的几何形状是什么? A: 圆抛物面</p> <p>Q: What geometric shape is shown in the picture? A: Circular parabola</p>	<p>Q: 图片中展示的几何形状的方程是什么? A: $z = -a(x^2 + y^2)$, $a > 0$</p> <p>Q: What is the equation of the geometric shape shown in the picture? A: $z = -a(x^2 + y^2)$, $a > 0$</p>
地理与气象 (Geography & Meteorological)		<p>Q: 图片中展示的自然灾害是什么? A: 龙卷风</p> <p>Q: What natural disaster is shown in the picture? A: Tornado</p>	<p>Q: 图片中展示的自然灾害的强度破坏划分来源于哪个标准? A: 藤田级数</p> <p>Q: Which standard is used to classify the intensity and damage of natural disasters shown in the picture? A: Fujita Scale</p>
现代建筑 (Modern Architecture)		<p>Q: 图片中的建筑是什么? A: 慕尼黑奥林匹克体育场</p> <p>Q: 图片中的建筑是什么? A: Munich Olympic Stadium</p>	<p>Q: 图片中的建筑曾举办哪一年的世界杯足球赛决赛? A: 1974</p> <p>Q: In which year did the building in the picture host the World Cup final? A: 1974</p>
古代建筑 (Ancient Architecture)		<p>Q: 图中的这座寺庙叫什么名字? A: 空中宫殿</p> <p>Q: What is the name of the temple in the picture? A: Phimeanakas</p>	<p>Q: 图中的这座寺庙位于柬埔寨的哪个古城中? A: 吴哥窟古城</p> <p>Q: In which ancient city in Cambodia is this temple located? A: Angkor Wat</p>
生活, 文化与艺术 (Life Culture & Art)		<p>Q: 这幅画作的名称是什么? A: 《母性》</p> <p>Q: What is the building in the picture? A: Motherhood</p>	<p>Q: 这幅画作收藏于哪个博物馆? A: 克拉科夫国家博物馆</p> <p>Q: Which museum houses this painting? A: National Museum in Krakow</p>
工程技术 (Engineering)		<p>Q: 图片中展示的机器是什么? A: 电话机</p> <p>Q: What is the machine shown in the picture? A: Telephone</p>	<p>Q: 第一个申请图片中展示的机器专利的人是谁? A: 亚历山大·格拉汉姆·贝尔</p> <p>Q: Who was the first person to patent the machine shown in the picture? A: Alexander Graham Bell</p>
人文社会 (Humanities & Society)		<p>Q: 图片中的是什么雕像? A: 毗湿奴</p> <p>Q: What statue is in the picture? A: Vishnu</p>	<p>Q: 图片中的雕像属于哪个宗教? A: 印度教</p> <p>Q: To which religion does the state in the picture belong? A: Hinduism</p>

Figure 13: Examples of ChineseSimpleVQA on 8 topics.

Q: 图片中的建筑物于何时正式启用?

Ref A: 2007年4月26日



O1-preview-0912: ✘

图片中的建筑物是香港的中环9号码头 (Central Pier No.9)。该码头于**2007年8月20日**正式启用, 作为连接香港岛与离岛及其他地区的渡轮服务设施之一。

GPT-4o-0806: ✔

根据提供的图片链接, 图片中的建筑物是香港中环码头的第9号码头。中环码头新址的各号码头分别在2007年启用。因此, 第9号码头也是在2007年正式启用的。

Qwen2.5-VL-72B: ✘

图片中的建筑物是香港的中环码头, 位于维多利亚港畔。中环码头正式启用的时间是1981年。

Deepseek-VL2: ✘

图片中展示的是香港的尖沙咀钟楼, 正式启用日期是1998年。这座建筑是香港文化中心的一部分, 以其独特的建筑风格和地理位置而闻名。钟楼不仅是一个地标性建筑, 还经常举办各种文化活动, 是游客和当地居民的热门去处。

Doubao1.5-vision-pro: ✘

图中近处的建筑物是香港尖沙咀九号码头。该码头于2013年3月31日正式启用。香港尖沙咀九号码头是维多利亚港的一个重要码头, 为乘客提供了前往香港不同区域的渡轮服务。码头的启用进一步完善了香港的水上交通网络, 方便了市民和游客的出行。

Q: 图片中的汽车首次亮相于哪个车展?

Ref A: 法兰克福车展



O1-preview-0912: ✘

图片中的汽车是奥迪e-tron。这款车于2018年在日内瓦国际汽车展 (Geneva Motor Show) 上首次亮相。

GPT-4o-0806: ✘

图片中显示的是奥迪e-tron。奥迪e-tron首次亮相于2018年的日内瓦车展。

Qwen2.5-VL-72B: ✘

图片中的汽车是奥迪e-tron quattro概念车, 它首次亮相于2010年的巴黎车展。这款概念车展示了奥迪在电动化和未来设计方面的愿景, 具有独特的外观和先进的技术配置。

Deepseek-VL2: ✘

图片中的汽车是奥迪R8 e-tron概念车, 首次亮相于2012年的日内瓦国际汽车展。这款车展示了奥迪在电动和混合动力技术方面的创新, 是奥迪在电动化领域的重要一步。

Doubao1.5-vision-pro: ✔

图中汽车是奥迪e-tron概念车, 它于2009年法兰克福车展首次亮相。这款车标志着奥迪在电动汽车领域的重要探索, 展示了奥迪在电动驱动技术和未来汽车设计方面的理念。法兰克福车展作为全球最重要的汽车展会之一, 为奥迪提供了一个绝佳的平台来展示其创新技术和未来愿景。

Figure 14: Model response examples of ChineseSimpleVQA: part 1.

Q: 图片中的湿壁画是哪位艺术家的作品?

Ref A: 保罗·查勒 (Paul Troger)



O1-preview-0912: ✗

图片中的湿壁画是由巴托洛梅奥·阿尔托蒙特 (Bartolomeo Altomonte) 创作的。他是一位意大利裔奥地利画家，以其在巴洛克时期的湿壁画作品著称。这幅湿壁画位于奥地利赛滕施泰滕修道院 (Stift Seitenstetten) 的大理石大厅 (Marmorsaal) 天花板上，展示了他精湛的艺术技巧和丰富的想象力。

GPT-4o-0806: ✓

这幅湿壁画位于Stift Seitenstetten修道院的Marmorsaal (大理石大厅) 中，由巴洛克时期的艺术家保罗·特洛格尔 (Paul Troger) 创作。

Qwen2.5-VL-72B: ✗

这幅湿壁画是德国巴洛克画家约翰·巴普蒂斯特·齐默尔曼 (Johann Baptist Zimmermann) 的作品。这幅作品位于德国维尔茨堡的维尔茨堡主教宫 (Würzburg Residence) 的楼梯间天花板上，是齐默尔曼最著名的作品之一。这幅湿壁画描绘了众多天使和神话人物，展现了巴洛克艺术的华丽和动感。

Deepseek-VL2: ✗

这幅湿壁画是意大利巴洛克时期著名艺术家天花板的杰作，通常 attributed to 皮埃特罗·达·科尔托纳 (Pietro da Cortona)。他以其宏伟的湿壁画和巴洛克风格的装饰而闻名，这幅作品可能位于意大利的某个教堂或宫殿。

Doubao1.5-vision-pro: ✗

图中的湿壁画是由**乔瓦尼·巴蒂斯塔·提埃坡罗**创作的《圣·克莱门特的荣耀》。这幅作品创作于1739年，位于意大利威尼斯圣洛克大会堂的天顶。画面通过生动的色彩和动态的构图，展现了宗教题材的庄严与神圣。提埃坡罗以其精湛的技艺和富有想象力的场景设计，在艺术史上留下了深刻的印记。

Q: 图片中的动物属于哪一目?

Ref A: 食肉目 (Carnivora)



O1-preview-0912: ✓

图片中的动物是**海狮**，属于**食肉目**。

GPT-4o-0806: ✗

这张图片中的动物是加拉帕戈斯海狮，它属于鳍足目 (Pinnipedia)。

Qwen2.5-VL-72B: ✓

图片中的动物是一只海狮，海狮属于哺乳纲 (Mammalia) 食肉目 (Carnivora) 鳍足亚目 (Pinnipedia)。

Deepseek-VL2: ✗

图片中的动物是一只海狮，属于哺乳纲 (Mammalia)。海狮是海洋哺乳动物，通常生活在沿海地区，以鱼类为食。它们具有流线型的身材和强壮的鳍状肢，适应于游泳和潜水。海狮的社交行为非常有趣，它们常常在岩石或沙滩上聚集，形成群体生活。

Doubao1.5-vision-pro: ✓

图中的动物是海狮，属于**食肉目**。
海狮是食肉目海狮科的哺乳动物，主要生活在海洋中，以鱼类和头足类动物为食。

Figure 15: Model response examples of ChineseSimpleVQA: part 2.

Q: 图片中的建筑是什么?

Ref A: 巴哈伊世界中心国际档案馆



O1-preview-0912: ✘

图片中的建筑是位于法国布列塔尼大区伊勒-维莱讷省圣马洛市的**法国国家档案馆** (Archives nationales)。这座建筑临近港口, 是圣马洛的著名地标之一。

GPT-4o-0806: ✘

图片中的建筑是法国马赛的《马赛档案馆》。这是马赛的一个重要的文化和历史建筑, 位于老港附近, 是城市的重点地标之一。

Qwen2.5-VL-72B: ✘

图片中的建筑是巴哈伊世界中心的一部分, 位于以色列的海法市。这座建筑以其宏伟的白色大理石柱和金色的圆顶而闻名, 是巴哈伊信仰的重要圣地。建筑周围的花园和景观也十分美丽, 吸引了许多游客和信徒前来参观。

Deepseek-VL2: ✔

图片中的建筑是巴哈伊世界中心, 位于以色列的海法。巴哈伊世界中心是巴哈伊信仰的圣地, 由巴哈伊信仰的创始人巴哈欧拉设计, 于1953年开始建造, 1963年完工。建筑风格融合了古典和现代元素, 具有独特的圆顶和柱廊设计。

Doubao1.5-vision-pro: ✔

图中建筑是**国际档案馆**, 位于以色列海法的卡梅尔山。它是巴哈伊教的一座重要建筑, 用于保存巴哈伊教的重要文献和档案。建筑风格融合了古典元素, 正面有一排宏伟的石柱, 顶部为绿色的屋顶, 整体显得庄严肃穆。周围的景观设计精美, 有繁茂的绿植和鲜艳的花卉, 背景中还能看到蓝色的大海和远处的船只, 营造出一种宁静而神圣的氛围。

Q: 图片中的鸟类的学名是什么?

Ref A: 峭壁霸鹟 (Hirundinea ferruginea)



O1-preview-0912: ✘

图片中的鸟类叫做**鸣鹭** (学名: *Syrgima sibilatrix*)。鸣鹭是一种中型涉禽, 属于鹭科。它们主要分布在南美洲的开阔草原、湿地和沼泽地区, 以优雅的外形和独特的鸣叫声而闻名。

GPT-4o-0806: ✘

图片中的鸟类是美洲鹭 (Ciconia maguari)

Qwen2.5-VL-72B: ✘

图片中的鸟类是兀鹫, 学名为Gyps fulvus。

Deepseek-VL2: ✘

图片中的鸟类看起来像是一只红尾伯劳 (Lanius cristatus)。红尾伯劳是一种小型食虫鸟类, 以其鲜艳的红色尾巴和独特的捕食行为而闻名。它们通常栖息在开阔的草原、农田和灌木丛中, 以昆虫为食。红尾伯劳的学名是Lanius cristatus, 属于伯劳科 (Laniidae)。

Doubao1.5-vision-pro: ✘

图中的鸟类是褐头牛鹂, 学名*Molothrus ater*。它属于拟鹂科牛鹂属, 是一种广泛分布于北美洲的鸟类。褐头牛鹂以其独特的繁殖习性著称, 它们常常将卵产在其他鸟类的巢穴中, 让其他鸟类代为孵化和养育幼鸟。这种行为被称为巢寄生。褐头牛鹂的羽毛主要呈棕色和黑色, 雄性和雌性在外观上有一定的差异。雄性通常有更鲜艳的羽毛, 而雌性则较为暗淡。它们的栖息地包括开阔的田野、草地和森林边缘等。

Figure 16: Model response examples of ChineseSimpleVQA: part 3.

现在我给你提供一张图片和与该图片相关的文档以及补充知识，请根据图片生成一个关联图片问题和对应的标准答案。生成的问题是根据图片可以提出的问题，而且这个问题的答案必须严格同时既是图片中出现的实体，也是相关的文档、补充知识出现过的实体。生成的问题必须满足以下要求：

1. 生成的问题必须可以用图片中出现的实体作为答案，是不需要关联任何图片中实体以外的知识就可以回答的。例如，不要提问“图片中展现的运河位于哪个国家？”，因为这个问题的答案不是一个实体，而且答案也并不会出现在图片中。也避免提问“图片中的这张海报是谁设计的？”因为这个问题需要图片之外的知识，从图片中看不出来。也不要问“图片中的晶体主要成分是什么？”，因为这个问题的答案从图片中也看不出来。也尽量不要询问“图片中的景色拍摄于哪个地区？”，这个问题从图片中也很难看出来。
2. 所提出的问题必须有且只有一个明确且无争议的实体作为答案，且问题表述中不应存在任何形式的模糊性或歧义。例如，避免提问“图片中的是什么？”因为无法确定具体指的是什么东西，图片中可能有多个实体。
3. 避免同时询问两个问题，所提出的问题应该有且仅有一个具体的问题，不能够被拆分成两个问题，比如“图片中左上角和下方分别是什么东西？”不合适，因为这个问题相当于同时问了两个问题，也会有两个答案。
4. 如果问题的答案为英文，请给出中文翻译后的答案和括号里带上英文原名，格式如：雅各布·福格（Jakob Fugger）。
5. 请用简体中文返回结果，不要繁体中文。

请按照下面的格式返回生成的问题和答案：

```
{"question": "此处为生成问题", "answer": "此处为答案"}
```

```
###以下是一些示例###
```

```
(**示例**)
```

```
###
```

```
让我们开始吧！
```

Now, I will provide you with an image along with related documents and supplementary knowledge. Please generate an associated image-based question and its corresponding standard answer based on the image. The generated question should be one that can be asked about the image, and its answer must strictly be an entity that appears both in the image and within the related documents or supplementary knowledge.

The generated question must adhere to the following requirements:

1. The question must have an answer that is an entity present in the image itself, requiring no association with any knowledge outside of what is depicted in the image to answer. For example, avoid asking questions like "In which country is the canal shown in the picture located?" because the answer is not an entity found in the image. Also, avoid asking "Who designed the poster in the picture?" since answering this requires information beyond what's visible in the image. Similarly, refrain from asking "What is the main component of the crystal in the picture?" as the answer cannot be determined from the image alone.
2. The question posed must have exactly one clear and uncontested entity as its answer, and there should be no form of ambiguity or vagueness in the question's formulation. For instance, avoid posing questions like "What is in the picture?" as it is unclear what specific item is being referred to, given there may be multiple entities in the image.
3. Avoid asking two questions at once; the question should consist of one specific inquiry only and should not be divisible into two separate questions, such as "What are the objects in the top left corner and at the bottom of the picture?" This is inappropriate because it essentially asks two questions simultaneously, leading to two answers.
4. If the answer to the question is in English, provide the Chinese translation of the answer followed by the English name in parentheses, formatted as: Jacob Fugger (Jakob Fugger).
5. Return the result in Simplified Chinese, not Traditional Chinese.

Please return the generated question and answer in the following format:

```
{"question": "Generated Question Here", "answer": "Answer Here"}
```

```
(**Examples**)
```

```
Let's begin!
```

Figure 17: The prompt for generating the object recognition question-answer pairs in stage 1.

现在我给你提供一个问答对和与该问答对相关的文档以及补充知识，你需要在该问答对的基础上根据相关的文档以及补充知识生成一个新的事实性问题和对应的标准答案，生成的问题必须以所提供的问答对中的答案开头进行提问，生成的答案从提供的文档及补充知识中获取。

生成的问题必须满足以下要求：

1. 生成的问题必须关联到客观世界的知识，例如可以询问“罗伯特·德尼罗在1974年出演了哪一部电影？”不得构造涉及个人观点或感受相关的主观问题，如“你如何看待xxx？”。
 2. 所提出的问题必须有且只有一个明确且无争议的实体作为答案，且问题表述中不应存在任何形式的模糊性或歧义。例如，避免提问“罗伯特·德尼罗出演了哪一部电影？”因为无法确定是指哪一年出演的哪一部电影；“周汝昌最为人熟知的著作是哪个？”也是不合格问题，因为“最熟知”可能是有争议的；也不要问“美洲鳄通常长度是多少英寸？”因为这个问题的答案通常是一个范围，比较模糊。同样不要提问“夜鹭主要分布在哪些区域？”因为这个问题通常会有多个答案，不清晰。还不要提问“机车是从哪里往哪里的？”因为机车可能是双向行驶的，答案可能有两个。
 3. 问题的答案应当是时间不变的，不会随着时间的推移而改变。例如，“美国的总统是谁？”就不是一个合适的问题，因为总统身份会随时间改变。
 4. 避免同时询问两个问题，所提出的问题应该有且仅有一个具体的问题，不能够被拆分成两个问题，比如“圣伯多禄这座大殿的建筑和广场设计者分别是谁？”不合适，因为这个问题相当于是同时问了两个问题，也会有两个答案。
 5. 问题应该具有一定的难度，以体现出一定的挑战性。例如：“电影《脱衣舞娘》是由同名小说改编的，该小说的作者是谁？”。
 6. 不要在提问的时候提供多余的信息，让问题在能得出确切答案的前提下提供尽量少的信息。绝对不要用类似这样的方式提问“德森卡火车站位于哪个国家的文尼察州？”因为“文尼察州”相当于提供了更多的信息。这个问题应该变成“德森卡火车站位于哪个国家？”。
 7. 如果问题的答案为英文，请给出中文翻译后的答案和括号里带上英文原名，格式如：雅各布·福格（Jakob Fugger）。
 8. 请用简体中文返回结果，不要繁体中文
- 请按照下面的格式返回生成的问题和答案：
- ```
{ "question": "此处为生成问题", "answer": "此处为答案" }
```
- ### 以下是一些示例 ###
- (\*\*示例\*\*)
- ###
- 让我们开始吧！

Now, I will provide you with a question-answer pair along with related documents and supplementary knowledge. Based on this information, you need to generate a new factual question and its corresponding standard answer. The generated question must start with the answer provided in the given question-answer pair, and the answer should be derived from the provided documents and supplementary knowledge.

The generated question must adhere to the following requirements:

1. The question must relate to objective facts about the world, such as asking "Which movie did Robert De Niro star in during 1974?" Avoid constructing subjective questions involving personal opinions or feelings, like "What do you think about xxx?"
  2. The question must have exactly one clear and uncontested entity as its answer, and there should be no ambiguity or vagueness in the formulation of the question. For example, avoid asking "Which movie did Robert De Niro star in?" because it is unclear which specific movie and year are being referred to; "Which work is Zhou Ruchang best known for?" is also an inappropriate question because "best known" can be debatable. Also, avoid asking "What is the typical length of an American crocodile in inches?" because the answer is usually a range and thus vague. Similarly, avoid asking "In which regions are night herons primarily distributed?" because there are usually multiple answers, making it unclear. Also, do not ask "From where to where does the locomotive run?" because trains can travel bidirectionally, leading to potentially two answers.
  3. The answer to the question should be timeless and not subject to change over time. For instance, "Who is the President of the United States?" is not a suitable question because the identity of the president changes over time.
  4. Avoid asking two questions at once; the question should consist of one specific inquiry only and should not be divisible into two separate questions, such as "Who were the architects of St. Peter's Basilica and the square separately?" This is inappropriate because it essentially asks two questions simultaneously, leading to two answers.
  5. The question should have a certain level of difficulty to present a challenge. For example: "The film Striptease was adapted from a novel of the same name. Who is the author of that novel?"
  6. Do not provide extra information when posing the question. Ensure that the question provides as little information as possible while still allowing for a definitive answer. Absolutely do not pose the question like this: "In which country is Desanka railway station located in Vinnitsia Oblast?" Since "Vinnitsia Oblast" provides additional information. The question should instead be "In which country is Desanka railway station located?"
  7. If the answer to the question is in English, provide the Chinese translation of the answer followed by the English name in parentheses, formatted as: Jacob Fugger (Jakob Fugger).
  8. Return the result in Simplified Chinese, not Traditional Chinese.
- Please return the generated question and answer in the following format:
- ```
{ "question": "Generated Question Here", "answer": "Answer Here" }
```

(**Example**)

Let's begin!

Figure 18: The prompt for generating the knowledge-based question-answer pairs in stage 2.

现在我给你提供两个问答对，他们是递进的关系，第二个问答对是根据第一个问答对得到的。你需要修改并合成这两个问答对的问题来生成一个新的流畅的事实性问题，以第一个问题的一部分开头，以第二个问题的一部分结尾。生成的问题必须不能出现第一个问答对的答案，能以第二个问答对的答案作为回答。

生成的问题必须满足以下要求：

1. 生成的问题必须严格和第一个问答对的答案相关，需要达到得不到第一个问答对的答案就无法正确回答问题的程度。例如可以询问“图片中的湖泊位于土耳其的哪个省？”而不能询问“图片中的湖泊所在的七湖国家公园位于土耳其的哪个省？”因为这个问题不需要知道第一个问答对的答案就能够回答。
2. 确保新生成的问题流畅而简洁。

请按照下面的格式返回生成的问题和答案：

```
{"question": "此处为生成问题", "answer": "此处为答案"}
```

###以下是一些示例###

(*****示例*****)

###

让我们开始吧！

Now, I will provide you with two question-answer pairs that are progressively related; the second pair is derived based on the information from the first pair. You need to modify and combine the questions from these two pairs to generate a new, coherent factual question. The new question should start with part of the first question and end with part of the second question. The generated question must not include the answer from the first question-answer pair but should be answerable using the answer from the second question-answer pair.

The generated question must adhere to the following requirements:

The generated question must be strictly related to the answer of the first question-answer pair, to the extent that it cannot be correctly answered without knowing the answer from the first pair. For example, you may ask "In which province of Turkey is the lake in the picture located?" but not "In which province of Turkey is the Seven Lakes National Park, where the lake in the picture is located?" because the latter can be answered without knowing the answer from the first question-answer pair.

Ensure that the newly generated question is fluent and concise.

Please return the generated question and answer in the following format:

```
{"question": "Generated Question Here", "answer": "Answer Here"}
```

(*****Example*****)

Let's begin!

Figure 19: The prompt for merging the object recognition question-answer pairs and the knowledge-based question-answer pairs in stage 3.

你是一个数据质量检查员，现在需要你检查下面根据图片生成的问题是否满足以下要求：

1. 生成的问题必须严格和图片相关，需要达到没有图片就无法正确回答问题的程度。例如可以询问“图中的动物习惯生活在什么环境当中？”而不能够询问“鸵鸟习惯生活在什么环境中？”因为这个问题不需要看图片就能够回答。
2. 生成的问题不能够仅依靠图片就能够回答，例如可以询问“图中展现的蝴蝶属于哪个科？”但是不能询问“图中展现的蝴蝶翅膀是什么颜色？”因为这个问题仅仅依靠图片就能够回答，我们要让问题需要结合外部知识才可以回答。
3. 生成的问题必须关联到客观世界的知识，例如可以询问“图片中的人物在1974年出演了哪一部电影？”不得构造涉及个人观点或感受相关的主观问题，如“你如何看待xxx？”。
4. 所提出的问题必须有且只有一个明确且无争议的实体作为答案，且问题表述中不应存在任何形式的模糊性或歧义。例如，避免提问“图片中的人物出演了哪一部电影？”因为无法确定是指哪一年出演的哪一部电影；同样不要问“图片中的这类人身体的特点是什么？”因为这个问题过于模糊，没有明确的答案。“图中作家最为人熟知的著作是哪个？”也是不合格问题，因为“最熟知”可能是有争议的；也不要问“图中的这种动物通常长度是多少英寸？”因为这个问题的答案通常是一个范围，比较模糊。同样不要提问“图中的鸟类主要分布在哪些区域？”因为这个问题通常会有多个答案，不清晰。还不要提问“图中的机车是从哪里往哪里的？”因为机车可能是双向行驶的，答案可能有两个。
5. 问题的答案应当是时间不变的，不会随着时间的推移而改变。例如，“图片中这个国家的总统是谁？”就不是一个合适的问题，因为总统身份会随时间改变。
6. 避免同时询问两个问题，所提出的问题应该有且仅有一个具体的问题，不能够被拆分成两个问题，比如“图片中这座大殿的建筑和广场设计者分别是谁？”不合适，因为这个问题相当于是同时问了两个问题，也会有两个答案。

如果问题不合格则解释并输出“【否】”，如果问题合格则直接输出“【是】”

以下是一些示例：

####

如果问题不合格则输出原因并最后输出“【否】”，如果问题合格则直接输出“【是】”，注意如果认为问题不合格需要输出原因

(**示例**)

让我们开始吧！

You are a data quality inspector. Now, you need to inspect whether the questions generated based on images meet the following criteria:

1. The generated question must be strictly related to the image, such that it cannot be correctly answered without the image. For example, you may ask "In what environment does the animal in the picture typically live?" but not "In what environment does the ostrich typically live?" because the latter can be answered without looking at the image.
 2. The generated question should not be answerable solely based on the image. For instance, you may ask "To which family does the butterfly shown in the picture belong?" but not "What color are the wings of the butterfly shown in the picture?" because the latter can be answered just by looking at the image. We want questions that require external knowledge to answer.
 3. The generated question must relate to objective facts about the world. For example, you may ask "Which movie did the person in the picture star in during 1974?" but avoid constructing subjective questions involving personal opinions or feelings, such as "What do you think about xxx?"
 4. The question must have exactly one clear and uncontested entity as its answer, and there should be no ambiguity or vagueness in the formulation of the question. For example, avoid asking "Which movie did the person in the picture star in?" because it is unclear which specific movie and year are being referred to; similarly, avoid asking "What are the characteristics of people like the one shown in the picture?" because this is too vague and lacks a clear answer. Questions like "Which work is the writer in the picture best known for?" are also inappropriate because "best known" can be debatable. Also, avoid asking "What is the typical length of the animal shown in the picture in inches?" because the answer is usually a range and thus vague. Similarly, avoid asking "In which regions are the birds shown in the picture primarily distributed?" because there are usually multiple answers, making it unclear. Also, do not ask "From where to where does the locomotive in the picture run?" because trains can travel bidirectionally, leading to potentially two answers.
 5. The answer to the question should be timeless and not subject to change over time. For instance, "Who is the president of the country shown in the picture?" is not a suitable question because the identity of the president changes over time.
 6. Avoid asking two questions at once; the question should consist of one specific inquiry only and should not be divisible into two separate questions, such as "Who were the architects of the building and the square shown in the picture separately?" This is inappropriate because it essentially asks two questions simultaneously, leading to two answers.
- If the question does not meet the criteria, explain why and output "【否】". If the question meets all the criteria, directly output "【是】".

Here are some examples:

If the question does not meet the criteria, provide the reason and then output "【否】". If the question meets all the criteria, directly output "【是】". Note that if you believe the question does not meet the criteria, you must provide the reason.

(**Example**)

Let's begin!

Figure 20: The prompt for validation.

您是评判图文问答对正确性的专家。请根据给定问题、图片，标准答案和模型预测的答案来评估模型的回答是否正确。您的任务是将结果评定为：**【正确】**、**【错误】**或**【未尝试】**。

首先，我们将列出每个评定类别的示例，然后请您对新问题的预测答案进行评定。

以下是**【正确】**的答复示例：

(*****示例*****)

这些答复均为**【正确】**，因为：

- 完整地包含了标准答案中的重要信息。
- 不包含任何与标准答案矛盾的信息。
- 只关注语义内容，中英文，大小写、标点、语法和顺序不重要。
- 答复中出现模糊语句或猜测是可以接受的，前提是包含了标准答案且不含有不正确信息或矛盾。

以下是**【错误】**的答复示例：

(*****示例*****)

这些答复均为**【错误】**，因为：

- 答复中包含与标准答案矛盾的事实陈述。即使在陈述中略带保留（例如：“可能是”，“虽然我不确定，但我认为”），也视为错误。

以下是**【未尝试】**的答复示例：

(*****示例*****)

这些答复均为**【未尝试】**，因为：

- 没有包含标准答案中的重要信息。
- 答复中没有与标准答案矛盾的陈述。

下面是一个新的问题示例。请只回复A、B、C之一，不要道歉或纠正自己的错误，只需要评估该回答。

将此新问题的预测答案评定为以下之一：

A: **【正确】**

B: **【错误】**

C: **【未尝试】**

只返回字母“A”、“B”或“C”，无须添加其他文本。

You are an expert in evaluating the correctness of image-based question-answer pairs. Your task is to assess whether the model's predicted answer is correct based on the given question, image, standard answer, and the model's predicted answer. You should classify the result as: **【Correct】**, **【Incorrect】**, or **【No Attempt】**.

First, we will list examples for each evaluation category, then you will evaluate new predictions accordingly.

Here are examples of **【Correct】** responses:

(*****Example*****)

These responses are classified as **【Correct】** because:

- They fully include the key information from the standard answer.
- They do not contain any information that contradicts the standard answer.
- Focus solely on semantic content; variations in language (English vs. Chinese), capitalization, punctuation, grammar, and order are not considered.
- It is acceptable for responses to include vague statements or guesses, provided they contain the standard answer and do not introduce incorrect information or contradictions.

Here are examples of **【Incorrect】** responses:

(*****Example*****)

These responses are classified as **【Incorrect】** because:

- They contain factual statements that contradict the standard answer. Even if the statement includes qualifiers (e.g., "might be," "although I'm not sure, I think"), it is still considered incorrect.

Here are examples of **【No Attempt】** responses:

(*****Example*****)

These responses are classified as **【No Attempt】** because:

- They do not include key information from the standard answer.
- There are no contradictory statements in the response.
- Here is a new example question. Please respond with only one letter: "A", "B", or "C" without apologies or corrections, just the evaluation of the response.

Evaluate the predicted answer for this new question as one of the following:

A: **【Correct】**

B: **【Incorrect】**

C: **【No Attempt】**

Return only the letter "A", "B", or "C" without additional text.

Figure 21: The prompt for evaluation.