

SciVQR: A Multidisciplinary Multimodal Benchmark for Advanced Scientific Reasoning Evaluation

Longteng Guo^{1,2,*}, Xuanxu Lin^{1,2,*}, Dongze Hao^{3,*}, Tongtian Yue^{1,2},
Pengkang Huo^{1,2}, Jiatong Ma^{1,2}, Yuchen Liu^{1,2}, Jing Liu^{1,2,†}

¹Institute of Automation, Chinese Academy of Sciences (CASIA)

²School of Artificial Intelligence, University of Chinese Academy of Sciences (UCAS)

³OPPO AI Center

{longteng.guo, jliu}@nlpr.ia.ac.cn

linxuanxu2025@ia.ac.cn, haodongze@oppo.com

Abstract

Scientific reasoning is a key aspect of human intelligence, requiring the integration of multimodal inputs, domain expertise, and multi-step inference across various subjects. Existing benchmarks for multimodal large language models (MLLMs) often fail to capture the complexity and traceability of reasoning processes necessary for rigorous evaluation. To fill this gap, we introduce SciVQR, a multimodal benchmark covering 54 subfields in mathematics, physics, chemistry, geography, astronomy, and biology. SciVQR includes domain-specific visuals, such as equations, charts, and diagrams, and challenges models to combine visual comprehension with reasoning. The tasks range from basic factual recall to complex, multi-step inferences, with 46% including expert-authored solutions. SciVQR not only evaluates final answers but also examines the reasoning process, providing insights into how models reach their conclusions. Our evaluation of leading MLLMs, including both proprietary and open-source models, reveals significant limitations in handling complex multimodal reasoning tasks, underscoring the need for improved multi-step reasoning and better integration of interdisciplinary knowledge in advancing MLLMs toward true scientific intelligence. The dataset and evaluation code are publicly available at <https://github.com/CASIA-IVA-Lab/SciVQR>.

1 Introduction

The rapid advancement of large-scale models, particularly multimodal large language models (MLLMs) (Li et al., 2025; Zhu et al., 2024; Trinh et al., 2024), has brought us closer to achieving human-level intelligence, capable of processing and reasoning across diverse modalities such as

text and images. Among human cognitive abilities, scientific reasoning is particularly crucial: analyzing multimodal cues, performing multi-step inference, integrating domain knowledge, and drawing systematic conclusions across diverse scientific fields. Despite advances in natural language understanding, current models still struggle with tasks requiring deep, structured reasoning grounded in multimodal scientific content. To rigorously assess progress in this direction, it is essential to benchmark MLLMs on their ability to perform scientific reasoning involving diverse subjects, multimodal inputs, and challenging multi-step reasoning.

We highlight three principles necessary for such evaluation:

(1) Disciplinary coverage: Scientific reasoning spans far beyond a few canonical subjects like physics or mathematics. A truly comprehensive benchmark should cover a broad range of scientific domains, reflecting the diversity of real-world knowledge.

(2) Difficulty diversity: A well-designed benchmark should cover a wide spectrum of difficulty levels, ranging from basic factual or conceptual reasoning to college- and graduate-level, multi-hop inference. This enables more nuanced evaluation across different reasoning depths and helps prevent performance saturation at either end of the spectrum.

(3) Solution traceability: Evaluating only the final answer is insufficient for understanding a model’s reasoning capabilities. Benchmarks should include key intermediate solution steps to enable fine-grained analysis of reasoning fidelity and structure.

While existing benchmarks address parts of this challenge, none cover all dimensions. For example, MMMU (Yue et al., 2024a) spans multiple subjects but lacks difficulty diversity and reasoning traces. ScienceQA (Lu et al., 2022) includes

*Equal contribution.

†Corresponding author.

solution steps but is limited to K-12 content with shallow reasoning. OmniMath (Gao et al., 2025) targets advanced mathematical reasoning but is narrow in scope and monolingual.

As a result, there remains a clear gap in benchmarking the emerging scientific capabilities of next-generation MLLMs, which requires disciplinary breadth, reasoning depth, and cross-linguistic generalization.

To bridge this gap, we introduce SciVQR, a comprehensive multimodal benchmark for scientific reasoning in MLLMs. Covering 54 sub-fields across 6 core scientific domains (mathematics, physics, chemistry, geography, astronomy, and biology), SciVQR ensures broad disciplinary representation. Figure 1 illustrates sampled questions from each subject.

The questions are sourced from high school, college, and graduate-level exams, quizzes, textbooks, and competitions, and span a wide range of difficulty levels (easy, medium, and hard). Many require advanced, multi-step reasoning, such as solving differential equations or interpreting complex molecular structures. Notably, 46% of the questions are accompanied by detailed, expert-authored solution traces to support precise reasoning evaluation. SciVQR also incorporates diverse scientific image formats, such as equations, charts, chemical structures, geological diagrams, and biological illustrations. It is developed through a rigorous, multi-stage quality assurance process involving both human experts and LLM assistance. SciVQR challenges models to integrate fine-grained visual understanding, deep subject knowledge, and sophisticated reasoning, thereby setting a new standard for scientific multimodal evaluation in MLLMs.

We evaluate a series of open-source and advanced proprietary MLLMs, including GPT-4o (Hurst et al., 2024) and o1 (Jaech et al., 2024), on SciVQR. Our key findings are summarized as follows:

- **Narrowing gap between proprietary and open-source models:** While there remains a substantial performance gap between proprietary models and open-source models, leading open-source models like Qwen2.5-VL-72B (Bai et al., 2025) (61.8%) and the reasoning-optimized MiMo-VL-7B (Xia et al., 2025) (69.2%) have started to close the gap, even surpassing non-reasoning proprietary models such

as GPT-4o (Hurst et al., 2024) (60.7%).

- **Disciplinary performance differences:** Models perform significantly worse in mathematics and physics due to the greater emphasis on quantitative analysis and reasoning. In contrast, subjects like biology and geography, which rely more on textual understanding and knowledge, see relatively better performance.
- **Reasoning optimization leads to better performance:** Reasoning-optimized models such as o1 (Jaech et al., 2024) (77.2%), o4-mini (OpenAI, 2024) (79.3%), and MiMo-VL (Xia et al., 2025) (69.2%) substantially outperform non-reasoning models like GPT-4o (Hurst et al., 2024) (60.7%), with marked improvements across all subjects. Notably, mathematics and physics scores experience the greatest boost, underscoring the importance of reasoning capabilities.
- **Reasoning optimization improves reasoning traces:** Our CoT evaluations reveal that reasoning-optimized models tend to generate more faithful, complete, and structured reasoning traces. This suggests that optimizing for step-by-step reasoning, rather than relying solely on instruction tuning, is crucial for improving scientific inference performance.

Our findings emphasize that achieving true multimodal scientific reasoning requires deeper multi-step inference and the integration of diverse knowledge across subjects.

2 Related Work

2.1 Multimodal Large Language Models

The development of MLLMs (Liu et al., 2023; Li et al., 2025, 2023a; Wang et al., 2024b; Bai et al., 2025; Liu et al., 2025a, 2024b; Wu et al., 2024) builds on advances in multimodal learning and the success of LLMs (Achiam et al., 2023; Yang et al., 2024; Liu et al., 2024a). Early vision-language frameworks established the foundation for modern systems, which now combine strong vision encoders with LLMs to handle diverse visual inputs. Open-source initiatives and scalable training, including visual instruction tuning (Liu et al., 2023) and synthetic data generation (Zhang et al., 2025; Maaz et al., 2024), have improved adaptability and generalization. While closed-source models like

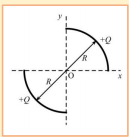
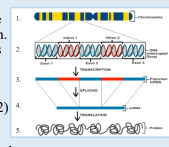

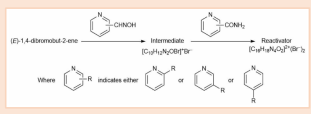





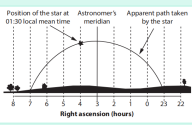
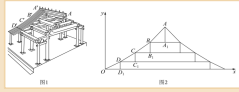
<p>Physics</p> <p>Question: Each of the two curved rods shown in the picture form one quarter of a circle with a radius R. Both rods carry a uniformly distributed electric charge $+Q$. Which of the following choices correctly expresses the net electric field and net electric potential at the origin? Assume $V \rightarrow 0$ as $R \rightarrow \infty$.</p>  <p>A. Electric Field : zero, Electric Potential : zero. B. Electric Field : zero, Electric Potential : $\frac{2kQ}{R}$. C. Electric Field : $\frac{2kQ}{R^2}$, Electric Potential : zero. D. Electric Field : $\frac{\sqrt{2}kQ}{R^2}$, Electric Potential : $\frac{2kQ}{R}$. E. Electric Field : $\frac{2kQ}{R^2}$, Electric Potential : $\frac{2kQ}{R}$.</p> <p>Answer: D</p>	<p>Biology</p> <p>Question: Eukaryotic genes tend to consist of coding regions (exons) and non-coding regions (introns). The figure shows how such a gene leads to the production of a protein. Which of the following statements is true?</p>  <p>A. Thymine content of (1) and (2) is approximately equal. B. The process occurring between (2) and (3) takes place in the cytosol. C. (4) can hybridise with (2). D. The number of amino acid residues in (5) must equal the number of nucleotide residues in (2). E. All processes occurring between (3) and (5) take place in the nuclei.</p> <p>Answer: E</p>	<p>Geography</p> <p>Question: Gabi and Wade remembered they discussed a similar hot spot chain, forming the Hawaiian Islands, in year 9. Going back through their science notes, which of course they kept, they found a map of the Hawaiian-Emperor Seamount chain (Figure).</p>  <p>Using the Hawaiian-Emperor Seamount chain map (Figure 9), what direction did they decide the Pacific Plate has been travelling for the past 45 million years?</p> <p>A. North B. South C. Northwest D. Southwest E. Northeast F. Southeast</p> <p>Answer: C</p>
<p>Chemistry</p> <p>Question: This group of reactants was synthesised via a 2-step pathway, using a series of related reagents:</p>  <p>Where  indicates either  or  or  or </p> <p>How many different reactants could be made with the above set of reagents?</p> <p>Answer: 4</p>	<p>Astronomy</p> <p>Question: The figure shows a Simplified sketch made by an astronomer observing the apparent path taken by a star during the night. The position of the star is shown when the astronomer's local mean time was 01:30.</p>  <p>State the number of hours that the star spends above the astronomer's horizon during this night's observation. The number of hours: ___hours</p> <p>Answer: 8</p>	<p>Mathematics</p> <p>Question:</p>  <p>图1是中国古代建筑中的举架结构, AA', BB', CC', DD' 是桁, 相邻桁的水平距离称为步, 垂直距离称为举, 图2是某古代建筑屋顶截面的示意图. 其中 DD_1, CC_1, BB_1, AA_1 是举, OD_1, DC_1, CB_1, BA_1 是相等的步. 相邻松的举步之比分别为 $\frac{DD_1}{OD_1} = 0.5$, $\frac{CC_1}{DC_1} = k_1$, $\frac{BB_1}{CB_1} = k_2$, $\frac{AA_1}{BA_1} = k_3$. 已知 k_1, k_2, k_3, k_4 成公差为 0.1 的等差数列, 且直线 OA 的斜率为 0.725, 则 $k_3 =$</p> <p>A. 0.75 B. 0.8 C. 0.85 D. 0.9</p> <p>Answer: D</p>

Figure 1: Sampled SciVQR examples from each subject. To solve the questions, the models need strong multimodal reasoning capabilities.

GPT-4o (Hurst et al., 2024) show strong visual reasoning, open-source models such as LLaVA (Liu et al., 2023), MiniGPT-4 (Zhu et al., 2024), and Qwen2-VL (Wang et al., 2024b) advance progress by refining architectures, expanding datasets, and broadening applications.

Reasoning capabilities of MLLMs have evolved alongside text-based LLMs (Wei et al., 2022; Guo et al., 2025). Though once criticized as mere token predictors (Huang and Chang, 2023), LLMs now demonstrate skill in logic (Xin et al., 2025), problem-solving (Rein et al., 2024), and code generation (Hui et al., 2024). This has spurred multimodal reasoning research, with methods like Visual-RFT (Liu et al., 2025b), VLM-R1 (Shen et al., 2025), and Vision-R1 (Huang et al., 2026), emphasizing perception-reasoning integration for complex visual tasks.

2.2 Multimodal Reasoning Evaluation

Evaluation of multimodal systems has progressed from early benchmarks such as VQA (Antol et al., 2015) and MSCOCO (Lin et al., 2014), which tested perception and alignment, to more comprehensive ones. Recent benchmarks like POPE (Li et al., 2023b), LAMM (Yin et al., 2023), MM-Bench (Xu et al., 2023), and SEED (Li et al., 2024)

expand coverage to hallucination detection, cross-modal reasoning, and real-world simulation. They assess capabilities from low-level recognition to high-level contextual understanding, reflecting increasing model complexity.

For reasoning-specific evaluation, targeted benchmarks focus on scientific and mathematical reasoning. MathVista (Lu et al., 2024) and OlympiadBench (He et al., 2024) integrate diagrams, equations, and domain knowledge. MathVerse (Zhang et al., 2024) removes textual redundancy in visual problems to test pure graphical reasoning. MMMU-Pro (Yue et al., 2024b) further introduces filtered protocols, vision-only settings, and adversarial questions to avoid text shortcuts. These works mark a shift toward inseparable modality fusion, where visual analysis and iterative reasoning are essential.

Despite these advances, none provide a comprehensive evaluation of scientific reasoning. Many focus on perception or commonsense reasoning, while domain-specific ones face limitations. Some target only K-12 content with shallow reasoning (e.g., ScienceQA (Lu et al., 2022)), while others (e.g., EMMA (Hao et al., 2025)) restrict themselves to mathematics, physics, and chemistry, missing broader scientific knowledge. Most

benchmarks also provide only final answers, without step-by-step solutions needed to assess reasoning fidelity.

3 The SciVQR Benchmark

3.1 Overview of SciVQR

We present SciVQR, a carefully curated benchmark developed to assess the multimodal scientific reasoning capabilities of foundation models across a diverse set of visually grounded scientific tasks. As shown in Figure 2, SciVQR covers 6 scientific subjects, including mathematics, physics, chemistry, astronomy, geography, and biology, and each subject contains various specialized subfields. The questions in our benchmark were manually collected from different academic competitions, college and graduate-level exam sets, and college textbooks. We categorize all questions into easy, medium, and hard levels with reference to their original source materials.

3.2 Data Curation Process

Data Collection: Our benchmark is built in two stages. First, we collect competitions, graduate exams, college exams, and textbooks across multiple subjects, prioritizing resources rich in images to ensure sufficient visual questions. This yields 15 scientific competitions, 9 college entrance exams, and 6 college or graduate-level exams and textbooks from different countries. In the second stage, we recruit over 20 postgraduates to annotate the questions, following copyright and licensing rules to exclude restricted sources. All annotators receive fair compensation based on the local competitive hourly rate. To mitigate data contamination and language priors, we select questions where visual information is indispensable rather than auxiliary. Annotators capture screenshots of visual questions, then extract text using the Mathpix OCR API. All images are standardized to PNG format, and since each image naturally aligns with one question, we obtain one-to-one image-question pairs. This process produces a diverse pool of 5k questions.

Data Quality Control: To ensure the quality of the collected data, we use a web-based tool to filter the annotated questions. In the first stage, we exclude samples with low-resolution or incorrect images, and verify the correctness of the textual content, ensuring that all examples conform to the required quality standards. In the second stage,

we use lexical overlap and Levenshtein distance to identify and remove duplicate problems. After these steps, we obtain the final 3,254 science-VQA instances for our benchmark. In the third stage, we categorize the problems of different subjects into different subfields. We first provide a predefined list of possible subfields to GPT-4V (OpenAI, 2023), prompting the model to assign an initial subfield category to each question. We then carefully review and correct these preliminary labels. Subsequently, we assign a difficulty level to each question, guided by its original source material. Across all these stages, we perform iterative manual verification and consistency checks to ensure the reliability of the final annotations.

3.3 Data Analysis

Statistic	Number
Total questions	3,254
- Multiple-choice questions	2,545 (78.2%)
- Free-form questions	709 (21.8%)
- Questions with an explanation	1490 (45.8%)
English questions	2661 (81.8%)
Chinese questions	593 (18.2%)
Subjects	6
Subfields	54
Difficulty levels	3
Maximum question length	495
Average question length	77.2
Maximum explanation length	782
Average explanation length	102.4

Table 1: Main statistics in SciVQR.

Key Statistics. Table 1 summarizes the main statistics of SciVQR. The dataset contains 3,254 questions, balancing complexity and practicality with both structured and unstructured formats. Specifically, 78.2% of the questions are multiple-choice, offering a controlled evaluation setting, while the remaining 21.8% are free-form, encouraging open-ended reasoning and expression. Notably, 45.8% of all questions include explanations, offering insights into reasoning processes that can support interpretability and educational analysis.

In terms of language distribution, 81.8% of the questions are in English and 18.2% in Chinese, enabling multilingual research. Question lengths can reach up to 495 words with an average of 77.2 words, reflecting the diversity and richness of the dataset. Explanations average 102.4 words and can extend up to 782 words, providing both concise justifications and detailed elaborations de-

Geography Geology Geomorphology Hydrology Climatology Geography Disaster Geography Geospatial Cartography Quantitative Geography Cosmogeography Population Geography Economic Geography Cultural Geography	Astronomy Astrophysics Cosmology Stellar Astronomy Galactic Astronomy Planetary Science Observational Astronomy	Mathematics Arithmetic Algebra Logic Statistics Combinatorics Graph Theory Probability Theory Mathematical Analysis Analytic Geometry Solid Geometry Metric Geometry	Biology Microbiology Cell Biology Genetics Zoology Botany Physiology Ecology Biophysics Biomedicine Biochemistry Molecular	Physics Optics Acoustics Electricity Magnetism Electrodynamics Atomic Physics Fluid Mechanics Thermodynamics
	Chemistry Electrochemistry Organic Chemistry Physical Chemistry Inorganic Chemistry Analytical Chemistry			

Figure 2: Domain diversity in SciVQR. Each color corresponds to one subject: mathematics, physics, chemistry, biology, astronomy, geography. Each subject consists of diverse subfields.

pending on context. In addition, the dataset includes 997 easy, 1,240 medium, and 1,017 hard questions, ensuring a balanced difficulty distribution.

Domain Diversity. Each question belongs to one of six subjects: mathematics, physics, chemistry, astronomy, geography, and biology. Within each subject, questions are categorized into fine-grained types (e.g., Arithmetic, Logic, Optics, Cosmology). SciVQR covers 6 subjects and 54 categories. The treemap in Figure 2 visualizes this distribution, showing that questions span a wide and diverse range of scientific domains.

3.4 Comparisons with Existing Benchmarks

To advance beyond existing limitations, SciVQR introduces critical innovations across four dimensions, as shown in Table 2. In disciplinary coverage, we encompass 6 subjects and 54 subfields across scientific domains, offering broader coverage compared to existing benchmarks such as ScienceQA (5 subjects, 41 subfields) and MMMU (5 subjects, 39 subfields). In difficulty design, we establish a multi-tiered structure ranging from high school to graduate-level multi-hop reasoning, avoiding the saturation of datasets restricted to K-12 or college levels. In reasoning transparency, we pioneer stepwise solution annotations that trace intermediate reasoning steps. This enables fine-grained evaluation of CoT reasoning across multiple dimensions (faithfulness, informativeness, redundancy, hallucination, and missing steps), moving far beyond prevalent answer-only assessments. In linguistic inclusivity, we support dual-language (English/Chinese) evaluation to mitigate monolingual bias and enable cross-lingual reasoning

studies, whereas benchmarks like ScienceQA and MMMU are limited to English only. Together, these features promote genuine multimodal reasoning and set a new standard for comprehensive scientific evaluation.

4 Experiments

4.1 Experimental Setup

Baselines. We selected representative multimodal large language models as baselines for the SciVQR benchmark, listed below.

Open-source models: These include Qwen2.5-VL-72B (Bai et al., 2025), LLaVA-OneVision-7B (Li et al., 2025), LLaVA-NeXT-7B (Liu et al., 2024c), InternVL3-8B (Zhu et al., 2025), and MiMo-VL-7B-RL-2508 (Xia et al., 2025), among others. These models represent the state-of-the-art in publicly available multimodal LLMs across various parameter scales and training strategies, offering a broad view of open-source capabilities in scientific visual question answering.

Proprietary models: These include GPT-4o (Hurst et al., 2024), o1 (Jaech et al., 2024), o4-mini (OpenAI, 2024), and Gemini-2.5-Flash (gemini-2.5-flash-preview-04-17) (Anil et al., 2023). Compared to open-source models, these proprietary models typically show stronger cross-modal understanding and reasoning capabilities, reflecting the latest advances in multimodal reasoning. Models such as o1 and o4-mini are optimized for complex reasoning tasks with targeted training strategies, while Gemini-2.5-Flash supports both reasoning-enabled and reasoning-disabled modes. We used the reasoning-enabled mode in our evaluation.

All models were evaluated under a zero-shot

Name	#Q	#I	AvgQ	AvgS	#Subjects	#Subfields	Level	Lang
ScienceQA(Lu et al., 2022)	21,208	10,332	12.1	47.7	5	41	K-12	en
MathVista(Lu et al., 2024)	6,141	5,487	15.6	-	1	12	K-12, college	en
MATH-V(Wang et al., 2024a)	3,040	3,472	42.3	-	1	16	K-12, college	en
MMMU(Yue et al., 2024a)	11,550	11,264	59.3	108.0	5	39	college	en
EMMA(Hao et al., 2025)	2,788	2,599	-	-	3	14	K-12, college	en
SciVQR (Ours)	3,254	3,254	77.2	102.4	6	54	high school, college, graduate	en, zh

Table 2: Statistics for SciVQR and comparisons with existing datasets. #Q: number of questions, #I: number of images, AvgQ: average question length; AvgS: average solution length. #Subjects: number of scientific subjects. #Subfields: number of subfields across different scientific subjects.

Models	Math	Physics	Chemistry	Biology	Geography	Astronomy	Overall
<i>Open-source Models</i>							
LLaVA-NeXT-7B (Liu et al., 2024c)	8.0	12.6	14.1	28.9	24.2	29.0	19.5
LLaVA-OneVision-7B (Li et al., 2025)	21.9	21.0	24.6	40.6	49.2	37.5	32.5
MiniCPM-o-2.6-8B (Yao et al., 2024)	26.8	23.0	22.7	35.3	38.1	27.0	28.8
Molmo-7B-D-0924 (Deitke et al., 2025)	20.3	19.9	23.4	38.1	41.8	35.0	29.8
InternVL3-8B-Instruct (Zhu et al., 2025)	36.4	30.9	38.4	50.3	52.9	52.5	43.6
InternVL3-14B-Instruct (Zhu et al., 2025)	36.9	39.2	49.6	57.2	60.7	55.5	49.9
InternVL3-78B-Instruct (Zhu et al., 2025)	47.7	46.7	52.0	60.6	57.0	57.5	53.6
Qwen2.5-VL-7B-Instruct (Bai et al., 2025)	35.4	32.4	43.2	55.8	55.1	48.0	45.0
Qwen2.5-VL-72B-Instruct (Bai et al., 2025)	54.3	51.8	66.6	65.3	70.3	62.5	61.8
Keye-VL-1.5-8B (Yang et al., 2025)	61.1	55.2	53.9	62.2	64.4	61.5	59.7
MiMo-VL-7B-RL-2508 (Xia et al., 2025)	65.5	67.5	74.7	71.1	70.6	66.0	69.2
<i>Proprietary Models</i>							
GPT-4o (Hurst et al., 2024)	47.4	50.3	60.1	72.2	67.8	66.5	60.7
o1 (Jaech et al., 2024)	79.2	73.0	79.5	80.6	76.2	74.5	77.2
o4-mini (OpenAI, 2024)	84.5	77.7	85.4	80.8	76.5	71.0	79.3
Gemini-2.5-Flash (Anil et al., 2023)	83.2	78.1	87.6	78.6	76.2	72.0	79.3

Table 3: Comparison of model performances across various subjects. The best-performing model in each category is highlighted in **bold**.

setting to assess their ability to answer questions without fine-tuning or few-shot examples.

For prompting, we designed two formats: one presenting direct questions, and another incorporating CoT reasoning to study its effect on understanding and inference. Results in Table 3 are reported under the CoT setting. In Section 4.3, we further investigate the impact of CoT prompting by comparing model performance with and without CoT prompts. The corresponding prompt templates are provided in Appendix C.

All experiments were conducted on an internal computing cluster equipped with NVIDIA A800 GPUs.

4.2 Overall Results

Table 3 presents the performance of various multimodal large language models on the SciVQR benchmark across different scientific subjects.

Comparison between proprietary and open-source models: Proprietary models markedly surpass open-source models in overall accuracy. GPT-4o, o1, o4-mini, and Gemini-2.5-Flash all achieve

accuracy rates above 60%, with o4-mini and Gemini-2.5-Flash achieving the highest overall score of 79.3%. In contrast, among open-source models, Qwen2.5-VL-72B ranks highest (61.8%). Although reasoning-oriented models such as Keye-VL-1.5-8B (59.7%) and MiMo-VL-7B-RL-2508 (69.2%) show clear improvements over previous open-source baselines, their performance still lags notably behind proprietary reasoning models, suggesting that closed-source reasoning optimization remains more advanced. Lightweight models such as LLaVA-NeXT-7B and LLaVA-OneVision-7B remain below 40%, indicating their limitations in complex scientific reasoning.

Importance of reasoning capability: Models with strong reasoning ability such as o1 and o4-mini achieve the best results across all subjects (77.2% and 79.3%), demonstrating that scientific problem-solving requires not only perception but also multi-step reasoning. The strong performance of MiMo-VL-7B-RL-2508 among open-source models further confirms that reasoning optimization, rather than scale, is key to fidelity

Models	CoT	Math	Physics	Chemistry	Astronomy	Geography	Biology	Overall
InternVL3-8B-Instruct (Zhu et al., 2025)	×	34.1	32.7	39.9	51.4	58.8	51.0	44.7
	✓	36.4	30.9	38.4	50.3	52.9	52.5	43.6
		+2.3	-1.8	-1.5	-1.1	-5.9	+1.5	-1.1
Qwen2.5-VL-72B-Instruct (Bai et al., 2025)	×	39.2	42.4	55.4	62.2	67.5	56.5	53.9
	✓	54.3	51.8	66.6	65.3	70.3	62.5	61.8
		+15.1	+9.4	+11.2	+3.1	+2.8	+6.0	+7.9
GPT-4o (Hurst et al., 2024)	×	33.9	42.6	52.5	68.9	68.1	61.5	54.6
	✓	47.4	50.3	60.1	72.2	67.8	66.5	60.7
		+13.5	+7.7	+7.6	+3.3	-0.3	+5.0	+6.1
Keye-VL-1.5-8B (Yang et al., 2025)	×	29.9	31.3	42.0	52.5	59.1	45.5	43.4
	✓	61.1	55.2	53.9	62.2	64.4	61.5	59.7
		+31.2	+24.1	+11.9	+9.7	+5.3	+16.0	+16.3

Table 4: Influence of CoT on model performances across various subjects. For Keye-VL-1.5-8B, “CoT on/off” corresponds to enabling or disabling its internal reasoning mode.

Models	Easy	Medium	Hard
<i>Open-source Models</i>			
LLaVA-OneVision-7B	40.1	25.2	18.8
InternVL3-8B-Instruct	59.3	37.6	22.6
InternVL3-14B-Instruct	65.1	43.0	27.2
InternVL3-78B-Instruct	70.0	50.0	33.5
Qwen2.5-VL-7B-Instruct	51.6	39.8	31.2
Qwen2.5-VL-72B-Instruct	75.1	58.2	42.9
Keye-VL-1.5-8B	73.4	59.1	45.6
MiMo-VL-7B-RL-2508	79.7	69.7	55.4
<i>Proprietary Models</i>			
GPT-4o	69.4	53.5	44.8
o1	87.0	77.4	67.8
o4-mini	89.0	82.0	72.0
Gemini-2.5-Flash	89.4	81.5	71.1

Table 5: Results across questions of different difficulty levels.

and completeness. The superior performance of reasoning-oriented models highlights the importance of explicit reasoning mechanisms in multi-modal model design for handling complex and abstract scientific tasks.

Performance trends across subjects: Model performance varies notably across subjects. Mathematics and physics remain most challenging, especially for open-source models, due to intensive reasoning and computation requirements. By contrast, biology and geography rely more on conceptual understanding. Models enhanced for reasoning exhibit more consistent performance across subjects. For example, o4-mini achieves over 70% accuracy in all subject areas.

Performance across questions of different difficulty levels: As shown in Table 5, all models perform best on easy questions and worst on hard ones, confirming the dataset’s valid difficulty strat-

ification. Open-source models consistently lag behind proprietary ones across all levels, and the performance gap further increases with question difficulty. This indicates that while reasoning-optimized models such as o4-mini and Gemini-2.5-Flash demonstrate strong robustness, open-source models still face substantial challenges in complex multi-step reasoning tasks. Additional analyses on language balance and resampling stability are provided in Appendices A and B, respectively.

4.3 Does CoT Help in Answering SciVQR Questions?

Table 4 compares model performance under two prompting settings: with and without CoT prompting.

For most models, CoT prompting leads to clear improvements. For example, Qwen2.5-VL-72B gains +7.9% overall, with substantial boosts in mathematics (+15.1%) and chemistry (+11.2%). Similarly, GPT-4o benefits from CoT prompting, showing a +6.1% overall increase, with notable gains in mathematics (+13.5%) and physics (+7.7%). These results indicate that CoT prompting can substantially assist models in solving multi-step reasoning and structured scientific problems.

For the reasoning-specialized Keye-VL-1.5-8B, disabling its built-in reasoning mode causes a substantial performance drop (−16.3% overall), especially in mathematics (−31.2%) and physics (−24.1%). This degradation reflects the model’s strong reliance on CoT-style reasoning, as it is primarily trained under such paradigms.

However, CoT prompting does not always bring

Models	Faithful	Informative	Redundancy	Hallucination	Missing	Overall
<i>Open-source Models</i>						
LLaVA-NeXT-7B (Liu et al., 2024c)	2.7	3.9	5.6	3.1	2.7	3.6
Qwen2.5-VL-7B-Instruct (Bai et al., 2025)	3.3	4.1	7.2	3.9	3.4	4.4
Qwen2.5-VL-72B-Instruct (Bai et al., 2025)	4.4	5.2	6.8	5.1	4.6	5.2
Keye-VL-1.5-8B (Yang et al., 2025)	5.5	6.2	4.3	5.5	5.5	5.4
MiMo-VL-7B-RL-2508 (Xia et al., 2025)	6.0	6.4	3.6	6.1	6.0	5.6
<i>Proprietary Models</i>						
GPT-4o (Hurst et al., 2024)	5.4	6.1	7.8	6.1	5.6	6.2
o1 (Jaech et al., 2024)	6.5	7.0	8.6	7.3	6.6	7.2
o4-mini (OpenAI, 2024)	7.1	7.5	8.6	7.8	7.3	7.7
Gemini-2.5-Flash (Anil et al., 2023)	6.2	6.7	7.3	6.9	6.4	6.8

Table 6: Results of reasoning quality evaluation.

benefits. InternVL3-8B shows a slight drop (−1.1%) after applying CoT, suggesting limitations in following complex instruction formats or maintaining concise, well-structured responses. The generation of verbose or off-format responses may hinder the intended benefits of CoT reasoning.

In summary, CoT prompting generally enhances reasoning for models with strong instruction adherence, while exposing weaknesses in those lacking such capability. These results underscore the importance of robust instruction-following capabilities as a key factor in determining the effectiveness of reasoning-oriented prompting strategies.

4.4 Reasoning Quality Evaluation

To systematically evaluate the quality of model-generated reasoning, we adopt a structured rubric assessing five core dimensions of CoT quality: *Faithfulness*, *Informativeness*, *Redundancy*, *Hallucination*, and *Missing Steps*:

- **Faithfulness:** Measures alignment between the model’s reasoning steps and the ground-truth solution trace.
- **Informativeness:** Assesses how much relevant and critical information is preserved in the model’s reasoning.
- **Redundancy:** Evaluates whether reasoning steps are unnecessarily repeated or rephrased without contributing to the logic.
- **Hallucination:** Detects fabricated or irrelevant reasoning steps unsupported by the problem context or solution.
- **Missing Steps:** Identifies whether essential in-

termediate steps are omitted, affecting the completeness of reasoning.

Each dimension is rated on a 10-point scale, and an overall score is computed by averaging the sub-scores. The rubric aligns with human evaluation standards, focusing on step-level fidelity and informativeness rather than final-answer correctness. We employ GPT-4o as an automated judge to compare model-generated CoTs with expert-authored solutions on a diverse SciVQR subset (see Appendix C.3 for detailed evaluation prompts).

As shown in Table 6, reasoning-optimized models such as o4-mini (7.7) and o1 (7.2) substantially outperform instruction-tuned models like Qwen2.5-VL-72B (5.2) and Qwen2.5-VL-7B (4.4) across all dimensions. For example, o4-mini exceeds Qwen2.5-VL-72B by +2.5 overall, demonstrating stronger faithfulness, completeness, and groundedness of reasoning. In contrast, scaling Qwen2.5-VL from 7B to 72B yields only a marginal +0.8 improvement despite a tenfold parameter increase.

Notably, the reasoning-finetuned open-source model MiMo-VL-7B-RL-2508 achieves 5.6 overall, surpassing both Qwen2.5-VL variants despite its smaller scale, which highlights that reasoning optimization, not size, drives better CoT quality. Nevertheless, a clear gap remains between open-source and proprietary systems: although MiMo-VL-7B-RL-2508 attains higher task accuracy than GPT-4o, its reasoning is more verbose and repetitive, lowering conciseness and quality. These findings underline the continued advantage of proprietary models in producing coherent, faithful reasoning traces.

5 Conclusion

The introduction of SciVQR represents a significant advancement in evaluating the scientific reasoning capabilities of MLLMs. By encompassing a broad range of scientific subjects, domain-specific visuals, and expert-authored solution traces, SciVQR not only assesses final answers but also scrutinizes the reasoning processes behind them. This unique design provides deeper insight into how models arrive at their conclusions, offering a clearer understanding of the structure and quality of their reasoning.

We contend that strong performance on SciVQR, particularly in solution traceability, should be a critical criterion for evaluating the scientific reasoning capabilities of MLLMs. Therefore, SciVQR will play a pivotal role in advancing multimodal scientific intelligence, driving models toward more transparent, structured, and sophisticated reasoning.

Limitations

While SciVQR provides a comprehensive and challenging framework, it does have limitations, including potential biases in the manual curation process and the inherent challenges of capturing the full complexity of scientific reasoning. Moreover, the reliance on expert-authored solutions might introduce subjective factors that could affect the evaluation.

These limitations highlight areas for improvement in future iterations of SciVQR, where addressing biases and ensuring a more objective and comprehensive evaluation framework will be essential to enhance its utility and accuracy.

Acknowledgements

This research is supported by the Strategic Priority Research Program of Chinese Academy of Sciences under Grant XDB1350103, the National Natural Science Foundation of China (62437001, 62436001, 62531026), the Key Research Development Program of Jiangsu Province under Grant BE2023016-3, and the Natural Science Foundation of Jiangsu Province under Grant BK20243051.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,

Diogo Moitinho de Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, and 260 others. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.

Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M. Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, Orhan Firat, and 1331 others. 2023. Gemini: A family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.

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.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025. Qwen2.5-VL technical report. *arXiv preprint arXiv:2502.13923*.

Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, and 31 others. 2025. Molmo and PixMo: Open weights and open data for state-of-the-art vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 91–104.

Bofei Gao, Feifan Song, Zhe Yang, Zefan Cai, Yibo Miao, Qingxiu Dong, Lei Li, Chenghao Ma, Liang Chen, Runxin Xu, Zhengyang Tang, Wang Benyou, Daoguang Zan, Shanghaoran Quan, Ge Zhang, Lei Sha, Yichang Zhang, Xuancheng Ren, Tianyu Liu, and Baobao Chang. 2025. Omni-MATH: A universal Olympiad level mathematic benchmark for large language models. In *International Conference on Learning Representations*, pages 100540–100569.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, and 175 others. 2025. DeepSeek-R1 incentivizes reasoning in LLMs through reinforcement learning. *Nature* 645, pages 633–638.

Yunzhuo Hao, Jiawei Gu, Huichen Will Wang, Linjie Li, Zhengyuan Yang, Lijuan Wang, and Yu Cheng. 2025. Can MLLMs reason in multimodality?

- EMMA: An enhanced multimodal reasoning benchmark. In *International Conference on Machine Learning*.
- Chaoqun He, Renjie Luo, Yuzhuo Bai, Shengding Hu, Zhen Leng Thai, Junhao Shen, Jinyi Hu, Xu Han, Yujie Huang, Yuxiang Zhang, Jie Liu, Lei Qi, Zhiyuan Liu, and Maosong Sun. 2024. Olympiad-Bench: A challenging benchmark for promoting AGI with Olympiad-level bilingual multimodal scientific problems. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3828–3850.
- Jie Huang and Kevin Chen-Chuan Chang. 2023. Towards reasoning in large language models: A survey. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1049–1065.
- Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. 2026. Vision-R1: Incentivizing reasoning capability in multimodal large language models. In *International Conference on Learning Representations*.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, Kai Dang, Yang Fan, Yichang Zhang, An Yang, Rui Men, Fei Huang, Bo Zheng, Yibo Miao, Shanghaoran Quan, and 5 others. 2024. Qwen2.5-Coder technical report. *arXiv preprint arXiv:2409.12186*.
- Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, and 399 others. 2024. GPT-4o system card. *arXiv preprint arXiv:2410.21276*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helvar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, Ally Bennett, Ananya Kumar, and 242 others. 2024. OpenAI o1 system card. *arXiv preprint arXiv:2412.16720*.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2025. LLaVA-OneVision: Easy visual task transfer. *Transactions on Machine Learning Research*.
- Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. 2024. SEED-Bench: Benchmarking multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13299–13308.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning*, pages 19730–19742. PMLR.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. 2023b. Evaluating object hallucination in large vision-language models. In *Conference on Empirical Methods in Natural Language Processing*, pages 292–305.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision*, pages 740–755.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, and 180 others. 2024a. DeepSeek-V3 technical report. *arXiv preprint arXiv:2412.19437*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26286–26296.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024c. LLaVA-NeXT: Improved reasoning, OCR, and world knowledge.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, volume 36, pages 34892–34916.
- Zikang Liu, Longteng Guo, Yepeng Tang, Tongtian Yue, Junxian Cai, Kai Ma, Qingbin Liu, Xi Chen, and Jing Liu. 2025a. VRoPE: Rotary position embedding for video large language models. In *Conference on Empirical Methods in Natural Language Processing*.
- Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. 2025b. Visual-RFT: Visual reinforcement fine-tuning. In *International Conference on Computer Vision*, pages 2034–2044.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. MathVista: Evaluating mathematical reasoning of foundation models in visual contexts. In *International Conference on Learning Representations*.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science

- question answering. In *Advances in Neural Information Processing Systems*, volume 35, pages 2507–2521.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. 2024. Video-ChatGPT: Towards detailed video understanding via large vision and language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12585–12602.
- OpenAI. 2023. [GPT-4V\(ision\) System Card](#). Accessed: 2025-04-18.
- OpenAI. 2024. [OpenAI o3 and o4-mini System Card](#). Accessed: 2025-04-18.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Driani, Julian Michael, and Samuel R. Bowman. 2024. GPQA: A graduate-level Google-proof QA benchmark. In *Conference on Language Modeling*.
- Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jijia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao, Qianqian Zhang, Ruochen Xu, and Tiancheng Zhao. 2025. VLM-R1: A stable and generalizable R1-style large vision-language model. *arXiv preprint arXiv:2504.07615*.
- Trieu H. Trinh, Yuhuai Wu, Quoc V. Le, He He, and Thang Luong. 2024. Solving Olympiad geometry without human demonstrations. *Nature*, 625(7995):476–482.
- Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Houxing Ren, Aojun Zhou, Mingjie Zhan, and Hongsheng Li. 2024a. Measuring multimodal mathematical reasoning with the MATH-Vision dataset. In *Advances in Neural Information Processing Systems*, volume 37, pages 95095–95169.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024b. Qwen2-VL: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. Chain-of-Thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837.
- Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, Zhenda Xie, Yu Wu, Kai Hu, Jiawei Wang, Yaofeng Sun, Yukun Li, Yishi Piao, Kang Guan, Aixin Liu, and 8 others. 2024. DeepSeek-VL2: Mixture-of-experts vision-language models for advanced multimodal understanding. *arXiv preprint arXiv:2412.10302*.
- Bingquan Xia, Bowen Shen, Cici, Dawei Zhu, Di Zhang, Gang Wang, Hailin Zhang, Huaqiu Liu, Jiebao Xiao, Jinhao Dong, Liang Zhao, Peidian Li, Peng Wang, Shihua Yu, Shimao Chen, Weikun Wang, Wenhan Ma, Xiangwei Deng, Yi Huang, and 44 others. 2025. MiMo: Unlocking the reasoning potential of language model—from pretraining to posttraining. *arXiv preprint arXiv:2505.07608*.
- Huajian Xin, Z. Z. Ren, Junxiao Song, Zhihong Shao, Wanxia Zhao, Haocheng Wang, Bo Liu, Liyue Zhang, Xuan Lu, Qiushi Du, Wenjun Gao, Haowei Zhang, Qihao Zhu, Dejian Yang, Zhibin Gou, Z. F. Wu, Fuli Luo, and Chong Ruan. 2025. DeepSeek-Prover-V1.5: Harnessing proof assistant feedback for reinforcement learning and Monte-Carlo tree search. In *International Conference on Learning Representations*, pages 72274–72303.
- Cheng Xu, Xiaofeng Hou, Jiacheng Liu, Chao Li, Tianhao Huang, Xiaozhi Zhu, Mo Niu, Lingyu Sun, Peng Tang, Tongqiao Xu, Kwang-Ting Cheng, and Minyi Guo. 2023. MMBench: Benchmarking end-to-end multi-modal DNNs and understanding their hardware-software implications. In *2023 IEEE International Symposium on Workload Characterization (IISWC)*, pages 154–166. IEEE.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, Kai Dang, and 22 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- Biao Yang, Bin Wen, Boyang Ding, Changyi Liu, Chenglong Chu, Chengru Song, Chongling Rao, Chuan Yi, Da Li, Dunju Zang, Fan Yang, Guorui Zhou, Guowang Zhang, Han Shen, Hao Peng, Haojie Ding, Hao Wang, Haonan Fan, Hengrui Ju, and 42 others. 2025. [Kwai Keye-VL 1.5 technical report](#). Preprint, arXiv:2509.01563.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, Qianyu Chen, Huarong Zhou, Zhensheng Zou, Haoye Zhang, Shengding Hu, Zhi Zheng, Jie Zhou, Jie Cai, Xu Han, and 4 others. 2024. MiniCPM-V: A GPT-4V level MLLM on your phone. *arXiv preprint arXiv:2408.01800*.
- Zhenfei Yin, Jiong Wang, Jianjian Cao, Zhelun Shi, Dingning Liu, Mukai Li, Xiaoshui Huang, Zhiyong Wang, Lu Sheng, Lei Bai, Jing Shao, and Ouyang Wanli. 2023. LAMM: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark. In *Advances in Neural Information Processing Systems*, volume 36, pages 26650–26685.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan

- Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, and 3 others. 2024a. 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 (CVPR)*, pages 9556–9567.
- Xiang Yue, Tianyu Zheng, Yuansheng Ni, Yubo Wang, Kai Zhang, Shengbang Tong, Yuxuan Sun, Botao Yu, Ge Zhang, Huan Sun, Yu Su, Wenhao Chen, and Graham Neubig. 2024b. MMMU-Pro: A more robust multi-discipline multimodal understanding benchmark. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15134–15186.
- Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Yu Qiao, Peng Gao, and Hongsheng Li. 2024. MATHVERSE: Does your multi-modal LLM truly see the diagrams in visual math problems? In *European Conference on Computer Vision*, pages 169–186.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2025. LLaVA-Video: Video instruction tuning with synthetic data. *Transactions on Machine Learning Research*.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2024. MiniGPT-4: Enhancing vision-language understanding with advanced large language models. In *International Conference on Learning Representations*.
- Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, and 32 others. 2025. InternVL3: Exploring advanced training and test-time recipes for open-source multimodal models. *arXiv preprint arXiv:2504.10479*.

A Language Distribution Analysis

To assess the potential impact of language distribution on model evaluation, we conducted a language-specific performance analysis. Table 7 shows the accuracy of selected models on the English and Chinese subsets of SciVQR.

Model	English (%)	Chinese (%)
LLaVA-NeXT-7B	16.4	9.1
MiniCPM-o-2.6-8B	29.3	19.1
InternVL3-14B-Instruct	45.7	40.8
GPT-4o	57.9	45.7
Qwen2.5-VL-72B-Instruct	59.6	54.1
o1	79.2	69.1
o4-mini	83.0	72.2
Gemini-2.5-Flash	81.5	76.9

Table 7: Accuracy of selected models on English and Chinese subsets of SciVQR.

The results reveal that most multilingual models exhibit consistent performance trends across both languages. For most models, accuracy on English questions is slightly higher than on Chinese ones. This pattern may arise from differences in the models' multilingual training coverage or optimization, as many models have stronger English capabilities due to larger and more diverse English pretraining corpora.

B Robustness under Subsampling and Bootstrapping

To test evaluation stability, we conducted subsampling and bootstrapping analysis (100 trials each). Table 8 shows that standard deviations do not exceed 1.33%, confirming the robustness of the evaluation.

Model	Subsample (%)	Bootstrap (%)
LLaVA-NeXT-7B	0.78	0.63
MiniCPM-o-2.6-8B	1.10	0.78
InternVL3-14B-Instruct	1.33	0.87
Qwen2.5-VL-72B-Instruct	1.22	0.84
GPT-4o	1.24	0.89
o1	1.05	0.73
o4-mini	1.03	0.69
Gemini-2.5-Flash	0.95	0.68

Table 8: Standard deviations of model accuracy under subsampling and bootstrapping. "Subsample" refers to the standard deviation of model accuracy under subsampling, and "Bootstrap" refers to the standard deviation under bootstrapping.

C Prompt Templates

C.1 Prompt for Response Generation

The prompt for response generation is as follows:

System Prompt used for Response Generation with CoT

```
Please solve the problem step by step and put your answer in one "boxed{}". If it is a multiple choice question, only one letter is allowed in the "boxed{}"
{question}
{options}
```

System Prompt used for Response Generation without CoT

```
Answer the question directly without explanations or reasoning and put your answer in one "boxed{}". If it is a multiple choice question, only one letter is allowed in the "boxed{}"
{question}
{options}
```

C.2 Prompt for Evaluating Open Questions

The prompt for open questions evaluation is as follows:

System Prompt used for Open Questions Evaluation

```
You are given a response from a model and the correct answer. Your task is to determine if the model's response is correct. You should only return 'true' if the response matches the answer. If the answer is a floating-point number greater than 1, when it is represented in scientific notation, a difference of up to 0.1 is allowed. Otherwise, return 'false'.
Response: {example['response']}
Correct Answer: {example['answer']}
Is the response correct? (true/false)
```

C.3 Prompt for Evaluating Reasoning Quality

The prompt for reasoning quality evaluation is as follows:

System Prompt used for Reasoning Quality Evaluation (Part 1)

You are a reasoning evaluator tasked with analyzing the alignment, clarity, and completeness of reasoning steps in generated text. Your role is to assess the reasoning path between the ground truth and the LLM-generated response based on the criteria below:

** Faithfulness (1-10) :**

Description: Assesses how closely the LLM's reasoning steps reflect those in the source (ground truth).

Scoring Guide:

9-10: Reasoning steps almost entirely mirror the original.
7-8: Most reasoning is aligned, with only minor deviations.
5-6: Partial alignment, but several key differences or omissions.
3-4: Few reasoning elements match; many are missing or incorrect.
1-2: Majority of reasoning does not correspond to the source.

** Repetition and Redundancy (1-10) :**

Description: Evaluates repeated or unnecessary reasoning, including paraphrased points that do not add substance.

Scoring Guide:

9-10: No or minimal redundant content.
7-8: Slight repetition that does not hinder comprehension.
5-6: Noticeable redundancy without meaningful contribution.
3-4: Frequent repetition that disrupts the logical flow.
1-2: Excessive and disruptive duplication of reasoning steps.

** Informativeness (1-10) :**

Description: Measures whether the reasoning includes all essential details derived from the source.

Scoring Guide:

9-10: Nearly all important information is captured and accurately conveyed.
7-8: Most major points are included; only minor items are overlooked.
5-6: Some significant pieces of information are lacking or underdeveloped.

System Prompt used for Reasoning Quality Evaluation (Part 2)

3-4: Many critical points are missing.
1-2: Very poor coverage of key information.

** Hallucination (1-10) :**

Description: Identifies reasoning steps that are fabricated or unsupported by the source material.

Scoring Guide:

9-10: All reasoning is well-grounded in the original content.
7-8: One or two minor hallucinations are present.
5-6: Several steps introduce irrelevant or false information.
3-4: Many hallucinated steps, with partial grounding.
1-2: Reasoning is largely based on invented content.

** Missing Step (1-10) :**

Description: Checks whether any necessary components of the reasoning chain are absent.

Scoring Guide:

9-10: No critical reasoning steps are missing.
7-8: Only minor omissions that don't impact the overall conclusion.
5-6: Several important steps are left out, affecting the output's quality.
3-4: Multiple essential steps are not included.
1-2: Major reasoning gaps make the logic incomplete.

Consistency and Evaluation Guidelines:

1. Adhere strictly to the above scoring scales.
2. Always re-read both the ground truth and the LLM response thoroughly before scoring.
3. Perform a direct comparison of the reasoning sequences to assess agreement or deviation.
4. Use reference examples, if provided, to maintain evaluation consistency.
5. Avoid subjective judgment; stay within the defined thresholds.
6. After assigning scores for each category, compute the final Overall score as the average.

System Prompt used for Reasoning Quality Evaluation (Part 3)

Output must be returned in the following Python dictionary format exactly as shown do not add or modify anything, as this output feeds directly into a processing pipeline and any deviation will cause a system failure:

```
# Example output: {'Faithfulness': 8.0, 'Informativeness': 8.5, 'Repetition&Redundancy': 9.0, 'Hallucination': 9.5, 'Missing': 8.5, 'Overall': 8.65}
```

D Performance on Multiple-Choice and Direct-Response Questions

To provide a more granular understanding of model capabilities, we report separate performance on multiple-choice (MC) and direct-response (DR) questions. Results for a representative subset of models are shown in Table 9.

Overall Trends:

- **Higher accuracy for MC than DR.** All models perform better on MC questions, confirming that open-ended reasoning and generation are inherently more challenging than option selection.
- **Larger gaps for mid-scale open-source models.** Models such as Qwen2.5-VL-7B and InternVL3-8B show substantial performance drops from MC to DR, indicating weaker robustness in free-form reasoning.
- **Smaller gaps for stronger models.** Keye-VL-1.5-8B and GPT-4o exhibit more stable DR performance, suggesting better generalization beyond fixed-choice settings.

Model-wise Observations:

- Qwen2.5-VL-72B-Instruct demonstrates notable strengths in chemistry and geography, achieving the highest MC and DR scores among the listed models in these subjects (e.g., 67.7 MC / 61.3 DR in chemistry; 70.0 MC / 71.7 DR in geography).
- Keye-VL-1.5-8B performs particularly well in mathematics and physics, showing competitive

or leading results in these areas (e.g., 61.2 MC / 60.3 DR in math; 50.1 MC / 65.5 DR in physics).

E Reasoning Quality across Difficulty Levels

Reasoning evaluation is a central focus of our work. While Section 4.4 reports the overall quantitative evaluation of model-generated CoTs using our five-dimension rubric (faithfulness, informativeness, redundancy, hallucination, and missing steps), we additionally include an expanded evaluation here to examine how reasoning performance varies with question difficulty. Table 10 reports the results across Easy, Medium, and Hard questions.

The results demonstrate the following trends:

- Reasoning-optimized models achieve higher scores across all dimensions, showing better alignment with expert solution traces.
- Scores decrease progressively from lower- to higher-difficulty questions, indicating that complex problems lead to less faithful and complete CoTs.
- This trend supports the reliability of our difficulty annotations and demonstrates that SciVQR captures meaningful differences in reasoning complexity.

F Error Analysis

To analyze the limitations of current large-scale multimodal models, we perform both quantitative and qualitative error analyses based on empirical results. In our study, error instances are classified into categories including Image Perception Error, Refusal to Answer, and Lack of Knowledge, alongside an overarching reasoning category. To further disentangle reasoning bottlenecks, we subdivide reasoning errors into four fine-grained categories: Reasoning Inference Error, Reasoning Premise Error, Reasoning Misused Evidence, and Reasoning Step-wise Omission.

F.1 Quantitative Error Distribution

We conducted a quantitative error-type analysis to better understand model failure patterns. Table 11 reports the distribution of error types from 100 randomly sampled incorrect predictions for four representative models.

Model	Math		Physics		Chem.		Biology		Geo.		Astro.		Overall	
	MC	DR	MC	DR	MC	DR	MC	DR	MC	DR	MC	DR	MC	DR
InternVL3-8B-Instruct	37.1	32.9	34.1	24.6	39.5	33.3	52.5	38.6	54.2	45.7	65.0	23.3	47.1	33.1
InternVL3-78B-Instruct	48.5	44.3	48.9	42.5	54.1	42.7	62.7	49.1	57.0	56.5	71.4	25.0	57.1	39.6
Qwen2.5-VL-7B-Instruct	35.9	32.9	35.7	25.8	44.5	37.3	57.8	45.6	54.5	58.7	60.7	18.3	48.2	36.4
Qwen2.5-VL-72B-Instruct	56.5	44.3	55.7	44.0	67.7	61.3	67.7	52.6	70.0	71.7	77.1	28.3	65.8	50.4
Keye-VL-1.5-8B	61.2	60.3	50.1	65.5	54.9	49.3	64.0	52.6	64.3	65.2	66.4	50.0	60.2	54.6
GPT-4o	48.0	44.7	50.8	49.2	62.5	49.3	73.3	66.7	68.6	63.0	81.4	31.7	64.1	50.8

Table 9: Comparison of Multiple-Choice (MC) and Direct-Response (DR) accuracy across scientific subjects.

Model	Diff.	Fai.	Inf.	Red.	Hal.	Mis.	Over.
Qwen2.5-VL (7B-Inst.)	Easy	4.5	5.3	7.4	5.2	4.7	5.4
	Med.	4.3	5.2	7.3	5.0	4.5	5.3
	Hard	4.0	4.9	7.1	4.6	4.1	4.9
Keye-VL (1.5-8B)	Easy	6.2	6.8	5.1	6.3	6.2	6.1
	Med.	5.5	6.2	4.2	5.4	5.4	5.4
	Hard	5.1	5.8	3.9	5.0	5.0	5.0
GPT-4o	Easy	5.9	6.5	8.0	6.7	6.0	6.6
	Med.	5.4	6.1	7.8	6.2	5.6	6.2
	Hard	5.0	5.8	7.7	5.7	5.2	5.9
o4-mini	Easy	7.0	7.4	8.5	7.8	7.1	7.6
	Med.	6.4	7.0	8.4	7.2	6.6	7.1
	Hard	6.0	6.6	8.2	6.8	6.2	6.7

Table 10: Results of reasoning quality evaluation by difficulty level. Abbreviations: Diff. (Difficulty), Fai. (Faithfulness), Inf. (Informativeness), Red. (Redundancy), Hal. (Hallucination), Mis. (Missing), Over. (Overall). Qwen2.5-VL and Keye-VL refer to the Qwen2.5-VL-7B-Instruct and Keye-VL-1.5-8B versions, respectively.

Based on the quantitative results, we identify three key observations:

- **Reasoning errors dominate overall.** Across all evaluated models, reasoning failures remain the primary source of errors. This indicates that models frequently struggle to construct complete and coherent reasoning chains even when visual perception is relatively accurate.
- **Standard models show higher premise-level errors.** Models without explicit Long CoT capabilities (e.g., Qwen2.5-VL and GPT-4o) exhibit significantly more Reasoning Premise Errors. Without a structured reasoning process to re-evaluate initial assumptions, their early grounding mistakes directly lead to wrong answers.
- **Reasoning models show higher inference-level errors.** Models that produce explicit reasoning traces (e.g., MiMo-VL and o4-mini) typically ground visual premises correctly. However,

Model	Reasoning						
	Inf.	Pre.	Evi.	Omi.	Per.	Ref.	Kno.
Qwen2.5	14	43	6	8	28	1	0
GPT-4o	14	37	14	10	23	0	2
MiMo	49	15	4	8	23	0	1
o4-mini	29	15	15	24	16	0	1

Table 11: Error-type distribution (counts) from 100 sampled incorrect cases. Categories include Inf. (Reasoning Inference), Pre. (Reasoning Premise), Evi. (Misused Evidence), Omi. (Step-wise Omission), Per. (Image Perception), Ref. (Refusal), and Kno. (Knowledge). Qwen2.5 and MiMo denote the Qwen2.5-VL-7B-Instruct and MiMo-VL-7B-RL-2508 models, respectively.

their extended reasoning chains introduce more opportunities for logical inconsistencies and incorrect step transitions, making Reasoning Inference Errors their dominant failure mode.

F.2 Qualitative Error Examples

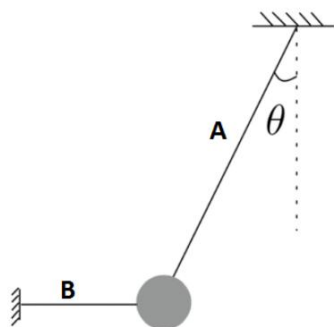
Following the quantitative breakdown, we provide concrete qualitative examples to illustrate these specific limitations. Figures 3 to 7 present typical failure cases of GPT-4o, spanning Reasoning Errors, Image Perception Errors, and Lack of Knowledge.

Reasoning Error

Question: As shown in the picture, a ball is attached to a ceiling and a wall with massless ropes A and B. Rope A is at angle θ from the vertical direction, and rope B is horizontal. The system is static. If rope B is cut, what is the ratio of the tension in rope A immediately after it is cut to the tension in A before it is cut?

Options: (A) 1 (B) $\cos \theta$ (C) $1 / \cos \theta$ (D) $\cos^2 \theta$ (E) $\cos \theta \sin \theta$

Ground Truth: (D) $\cos^2 \theta$



GPT-4o: To solve this problem, we'll analyze the situation before and after rope B is cut. Before rope B is cut, for forces in the horizontal direction, tension in rope B, T_B , is balanced by the horizontal component of the tension in rope A, $T_A \sin \theta$, so $T_B = T_A \sin \theta$. For forces in the vertical direction, the vertical component of the tension in rope A, $T_A \cos \theta$, balances the gravitational force on the ball, mg , so $T_A \cos \theta = mg$. After rope B is cut, with rope B gone, the ball can accelerate horizontally. However, since the question asks for the immediate tension in rope A, we look at only the vertical forces at that moment. The entire tension in rope A must now support the weight of the ball, so $T'_A = mg$. To determine the ratio, initially, from vertical force balance, $T_A \cos \theta = mg$. After rope B is cut, $T'_A = mg$. The ratio of the tension in rope A immediately after the cut to the tension in A before the cut is $\frac{T'_A}{T_A} = \frac{mg}{\frac{mg}{\cos \theta}} = \cos \theta$. Thus, the ratio is $\cos \theta$.

Error Reason: When analyzing the force conditions on the ball after the rope is cut, the model failed to correctly infer the ball's state of motion in the next moment. Therefore, it did not perform the force analysis correctly, which led to the final calculation of an incorrect result.

Figure 3: Reasoning error case of GPT-4o.

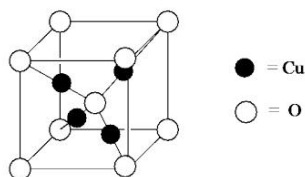
Perceptual Error 01

Question: A mineral containing only copper and oxygen adopts the cubic unit cell shown below. What is the formula of this mineral?

Options: (A) Cu_2O (B) CuO (C) Cu_3O_2 (D) Cu_4O_9

GPT-4o: To determine the formula of the mineral, we need to count the number of copper (Cu) and oxygen (O) atoms per unit cell. 1. Oxygen Atoms (White): - Corner Atoms: There are 8 corners in the cube, and each corner is shared by 8 adjacent cubes in a 3D lattice. Therefore, each corner atom contributes $\frac{1}{8}$ of an atom to the unit cell. - Total contribution from corners: $8 \times \frac{1}{8} = 1$ oxygen atom. 2. Copper Atoms (Black): - Face-Centered Atoms: There are 6 face-centered positions, and each face-centered atom is shared by 2 cubes. Therefore, each face-centered atom contributes $\frac{1}{2}$ of an atom to the unit cell. - Total contribution from face-centers: $6 \times \frac{1}{2} = 3$ copper atoms. Based on this counting, the unit cell contains: - 1 oxygen (O) atom. - 3 copper (Cu) atoms. **Conclusion:** The formula of the mineral, based on the copper-to-oxygen ratio, is Cu_3O . Comparing with the choices: - (A) Cu_2O - (B) CuO - (C) Cu_3O_2 - (D) Cu_4O_9 . The correct choice is not listed; however, based on the given options and typical chemical notation, the closest might be Cu_3O_2 but it's incorrect based on our detailed calculation. The closest choice according to typical unit cell counting and empirical formula conventions would resemble Cu_3O . But none of them perfectly match the analysis within the context provided. If forced to choose, Cu_3O_2 is typical for a smaller cubic cell if the lattice complexity is increased, so ideal choice is none of the above or revise the image interpretation. In conclusion based on provided options: - If acknowledging the typical possible counting errors or interpretations in generalized exams C is selected cautiously with errors or context arguments, or textbook mistake. More information or contextual clarity, as stated, gives exact match choices like Cu_3O .

Ground Truth:
(A) Cu_2O



Error Reason: The model incorrectly identified the positions of the copper atoms inside the unit cell as face-centered positions and also miscalculated the number of copper atoms.

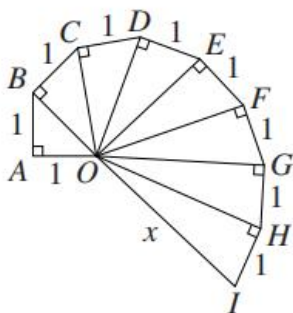
Figure 4: Perception error case of GPT-4o.

Perceptual Error 02

Question: What is the length of the line segment marked x ?

Options: (A) $\sqrt{2}$ (B) 2 (C) $2\sqrt{2}$ (D) 3 (E) 4

Ground Truth: (D) 3



GPT-4o: To solve this problem, we need to count the number of distinct paths that start at town S, visit each road exactly once, and end at town Q. The roads form a network that can be represented as a graph with the towns as vertices and roads as edges. Based on the diagram: - The towns are S, P, Q, and R. - The connections are: - S is connected to P and R. - P is connected to Q. - R is connected to Q. - P is connected to R. The paths must use each of the five roads once, starting at S and ending at Q. We can approach this problem by considering all possible orderings of towns while ensuring all roads are used once and starting/ending points are fixed. There are two routes from S (either to P or R), and then the paths continue as follows: 1. Route 1: $S \rightarrow P$ - After reaching P, go to R, as the path must eventually reach Q. - Then, from R, you can go to Q. - This path is $S \rightarrow P \rightarrow R \rightarrow Q$. 2. Route 2: $S \rightarrow R$ - After reaching R, go to P, as the path must eventually reach Q. - Then, from P, you can go to Q. - This path is $S \rightarrow R \rightarrow P \rightarrow Q$. So there are exactly two distinct paths that meet the criteria. Thus, the number of possible routes is 2.

Error Reason: The model incorrectly identified several distinct triangles in the diagram as congruent triangles. In the subsequent reasoning process, some inferences actually contradicted the identified information from the diagram, but the model failed to recognize the existence of these contradictions.

Figure 5: Perception error case of GPT-4o.

Lack of Knowledge

Question: What is the standard reduction potential of hydrogen selenate ion to form elemental selenium under acidic conditions at 298 K ?

Options: (A) 0.88 V (B) 0.94 V (C) 1.56 V (D) 1.89 V

GPT-4o: To find the standard reduction potential of the reaction $\text{HSeO}_4^-(\text{aq}) + 7\text{H}^+(\text{aq}) + 6e^- \rightarrow \text{Se}(\text{s}) + 4\text{H}_2\text{O}(\text{l})$, we can use the given potentials for intermediate reactions and apply the principle of adding reactions. Step 1: Identify the given reactions and their potentials: 1. $\text{HSeO}_4^-(\text{aq}) + 3\text{H}^+(\text{aq}) + 2e^- \rightarrow \text{H}_2\text{SeO}_3(\text{aq}) + \text{H}_2\text{O}(\text{l})$; $E^\circ = 1.15\text{V}$ 2. $\text{H}_2\text{SeO}_3(\text{aq}) + 4\text{H}^+(\text{aq}) + 4e^- \rightarrow \text{Se}(\text{s}) + 3\text{H}_2\text{O}(\text{l})$; $E^\circ = 0.74\text{V}$ Step 2: Add the equations directly: The total reaction involves the combination of these two steps: - The first half-reaction involves 2 electrons: $\text{HSeO}_4^-(\text{aq}) + 3\text{H}^+(\text{aq}) + 2e^- \rightarrow \text{H}_2\text{SeO}_3(\text{aq}) + \text{H}_2\text{O}(\text{l})$ - The second half-reaction involves 4 electrons: $\text{H}_2\text{SeO}_3(\text{aq}) + 4\text{H}^+(\text{aq}) + 4e^- \rightarrow \text{Se}(\text{s}) + 3\text{H}_2\text{O}(\text{l})$ Step 3: Sum the half-reactions: To balance the electrons in both half-reactions, we add them directly: $\text{HSeO}_4^-(\text{aq}) + 7\text{H}^+(\text{aq}) + 6e^- \rightarrow \text{Se}(\text{s}) + 4\text{H}_2\text{O}(\text{l})$ Step 4: Find the overall potential: Since the E° values given are standard reduction potentials, add them directly because the reactions and their stoichiometry are such that electrons are conserved: $E_{\text{total}}^\circ = E_1^\circ + E_2^\circ = 1.15\text{V} + 0.74\text{V} = 1.89\text{V}$ The standard reduction potential for the given overall reaction is 1.89V, so the correct choice is D.

Half-reaction	E° , V
$\text{H}_2\text{SeO}_3(\text{aq}) + 4\text{H}^+(\text{aq}) + 4e^- \rightarrow \text{Se}(\text{s}) + 3\text{H}_2\text{O}(\text{l})$	0.74
$\text{HSeO}_4^-(\text{aq}) + 3\text{H}^+(\text{aq}) + 2e^- \rightarrow \text{H}_2\text{SeO}_3(\text{aq}) + \text{H}_2\text{O}(\text{l})$	1.15
$\text{HSeO}_4^-(\text{aq}) + 7\text{H}^+(\text{aq}) + 6e^- \rightarrow \text{Se}(\text{s}) + 4\text{H}_2\text{O}(\text{l})$???

Ground Truth: (A) 0.88 V

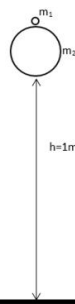
Error Reason: The model lacks knowledge of standard reduction potentials and thus failed to perform a weighted average during the calculation. Instead, it simply added the standard electrode potentials of the two reactions directly, leading to an incorrect result.

Figure 6: Knowledge gap error case of GPT-4o.

Textual Understanding

Question: As shown in the figure below, two balls of radii $R_1 = 2$ cm and $R_2 \gg R_1$, and with masses $m_2 \gg m_1$ fall from a height $h = 1$ m. Ignoring friction, assuming that all impacts are perfectly elastic and that the initial distance between the balls is small, which of the choices below is closest to the maximum height of ball 2 after one bounce?

Options: (A) 1 m (B) 2 m (C) 3 m (D) 4 m (E) 9 m



Error Reason: The problem asks for the maximum height reached by ball 2 after one bounce, but the model incorrectly assumed it was asked to find the maximum height reached by ball 1 after one bounce, leading to an incorrect answer.

Ground Truth: (A) 1 m

GPT-4o: To solve this problem, we will consider the collision physics between two balls and make use of conservation laws for elastic collisions. Step 1: Initial Velocities Before Collision - Both balls are dropped from a height $h = 1$ m. - The potential energy at height converts to kinetic energy just before impact. Using the equation for free fall: $v = \sqrt{2gh}$ where $g = 9.81$ m/s², the velocity of both balls just before hitting the ground is: $v = \sqrt{2 \times 9.81 \times 1} = \sqrt{19.62} \approx 4.43$ m/s. So, $v_1 = v_2 = 4.43$ m/s downward. Step 2: Collision with the Ground - Ball m_2 hits the ground first and reverses its velocity because it is elastic. Thus, m_2 will have $v_1' = v_2 = -4.43$ m/s after bouncing upward. Step 3: **Collision Between m_1 and m_2 - Since, we treat m_2 as effectively stationary during the collision from m_1 's viewpoint (after it starts moving upward). - Since the collisions are perfectly elastic, the relative velocity of approach equals the relative velocity of separation: $v_1' - (-v_2') = -(v_1 - (-v_2))$** Since v_2 is very large compared to v_1 and m_2 is moving upwards with speed v_2 , the equations simplify assuming perfect reversal for the small mass: $v_1' = -(v_1 - 2v_2) = -(4.43 - 2(-4.43)) \approx 13.29$ m/s. Step 4: Calculate Maximum Height of m_1 Now, find the height using v_1' : $h_{max} = \frac{(v_1')^2}{2g} = \frac{(13.29)^2}{2 \times 9.81} \approx 9$ m. Answer: The maximum height ball can reach is closest to: 9.

Figure 7: Textual understanding error case of GPT-4o.

G More Examples of the Dataset

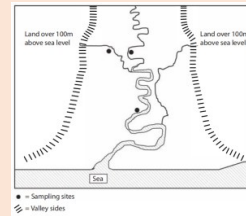
Geometry Case

Question:

Why is group work often better suited to river channel investigations than an individual working alone?

Choices:

- A. rivers are polluted.
- B. One person can carry the refreshments.
- C. One person can record while others use the equipment.
- D. The measuring equipment is heavy.



Answer: C

Solution:

Answer: C one person can record while others use the equipment

Distractors:

- A - irrelevant to the question and not a point of distinction between individual and group
- B - too non-geographical to be serious answer
- D - some geographical truth but there is a better fieldwork reason.

Figure 8: An example question from geography.

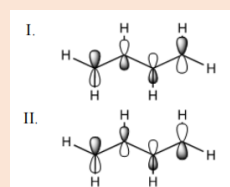
Chemistry Case

Question:

Which molecular orbitals are occupied in the ground state of 1,3-butadiene?

Choices:

- A. I only
- B. II only
- C. Both I and II
- D. Neither I nor II



Answer: B

Solution:

The molecular orbital diagram for the pi bonds of a conjugated diene can be made by drawing p orbitals on each carbon atom with different orientations. The lowest energy MO is made when the diagram has no nodes (all p orbitals have the same orientation), the second lowest when there is 1 node (there is one plane at which the two p orbitals directly adjacent to it have opposite orientations), etc. For 1,3butadiene, there are 4 carbon atoms and thus 4 p-orbitals drawn in each MO. 4 MO configurations exist: one with 0 nodes, one with 1 node, one with 2 nodes, and one with 3 nodes (be aware that for four p-orbitals, there can only be 3 "transitions" between p-orbital orientation for adjacent orbitals, thus 3 nodal planes). Each MO, like an atomic orbital, can fit 2 pi bonding electrons. Since there are 2 pi bonds and thus 4 pi bonding electrons, the two MOs with the least amount of energy (0 nodes and 1 node) are occupied (2 and 3 node MOs are empty). Looking at the answer choices, I has 2 nodes and II has 1 node. Thus, only MO II is filled, making B the correct answer.

Figure 9: An example question from chemistry.

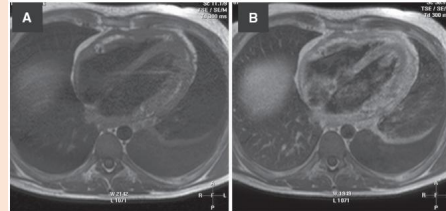
Biometry Case

Question:

The T1-weighted images were obtained from a young man with dyspnea, before (A) and after (B) administration of paramagnetic contrast agent (gadolinium). What is the likely diagnosis?

Choices:

- A. Cardiac tumor
- B. Pulmonary embolism
- C. Atrial septal defect
- D. Arrhythmogenic right ventricular dysplasia
- E. Constrictive pericarditis



Answer: E

Solution:

The images in the figure demonstrate the following abnormal findings:

- i. Circumferentially thickened pericardium (dark rim around the heart indicated by white arrows)
- ii. Pericardial contrast uptake (white arrowheads)
- iii. Left pleural effusion (asterisk)

These findings are suggestive of constrictive pericarditis with active inflammation. The pericardium is considered thickened when it is $>4\text{mm}$ [2, 3, 14, 15]. In this case, the thickness of the pericardium adjacent to the left ventricular wall was 10 mm. The pericardial contrast uptake suggests increased blood flow consistent with active inflammation. The right heart size is normal, arguing against intracardiac shunt, arrhythmogenic right ventricle, or pulmonary embolism. There is no evidence of cardiac mass.

Figure 10: An example question from biology.

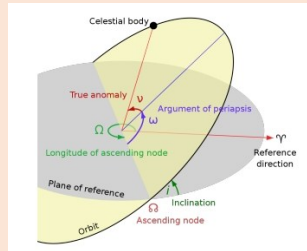
Astronomy Case

Question:

A new interstellar comet is about to enter the solar system, and Erez is really excited to calculate some Keplerian elements on its approach and figure out where it's going! The comet, dubbed Scratcher Crab by the astronomical community (for no particular reason), is approaching the solar system with a velocity of 29.8km/s from the direction of Altair, which has equatorial coordinates $\alpha = 19^h51^m, \delta = 8^\circ51'$. If the comet were to continue traveling in a straight line with its current velocity, its closest approach to the sun would be at a distance of 1au, and on March 21st, the comet would briefly be in superior conjunction

For reference, here is a diagram of orbital elements. Note that the longitude of the ascending node is measured from Aries and the argument of the periapsis is measured from the ascending node:

What are the inclination of the orbit and the longitude of the ascending node, both measured with reference to the equatorial plane (not the ecliptic)?



Answer: Orbits are planar

Solution:

Orbits are planar. This is a very useful fact, because it allows us to deduce that the orbit should lie in the same plane as the unique plane that contains both the sun and the straightline trajectory described in the problem. If on the straight-line trajectory, on March 21st the comet is behind the sun (at Aries, since this is the equinox), this means that there is a point in time on the straight-line trajectory where the comet has equatorial coordinates $(0,0)$. Thus, the orbit lies in the unique plane through the origin, Altair, and Aries. First, this means that the descending node is at Aries, i.e. the longitude of the ascending node is exactly 180° . Second, this allows us to deduce the inclination by four-parts on the spherical triangle with Altair, Aries, and the projection of Aries onto the equatorial plane

$$\begin{aligned} \cos 90 \cos(-\alpha) &= \sin 90(-\alpha)\cot\delta - \sin 90 \cot i \\ i &= \cot^{-1}(-\cot \delta \sin\alpha) \\ &= 9.98^\circ \end{aligned}$$

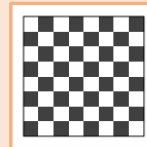
Figure 11: An example question from astronomy.

Mathematics Case

Question:

A chessboard is formed from an 8×8 grid of alternating black and white squares, as shown. The side of each small square is 1 cm.

What is the largest possible radius of a circle that can be drawn on the board in such a way that the circumference is entirely on white squares or corners?



Answer: $\frac{1}{2}\sqrt{10}$

Solution:

To go from one white square to another, it is necessary to pass through a corner.

But it is not possible for a circle to pass through the point where two white squares meet (such as Y in the diagram alongside) and both the diagonally opposite corners of this point (X and Z), since these points form a straight line and not an arc of a circle.

Is it possible for a circle to pass through exactly one of X and Z, for every choice of Y on the circle? In other words, the circle is alternately "close" to a diagonal and an edge of white squares.

The diagram alongside shows one such circle, whose centre is as shown. The circle has diameter $\sqrt{1^2 + 3^2}$, from Pythagoras' theorem. A circle passing through Y and neither X nor Z, for every point Y on the circle, such as the dashed circle in the diagram alongside, has a smaller radius.

Therefore the largest possible radius of the circle is $\frac{1}{2}\sqrt{10}$.

Figure 12: An example question from mathematics.

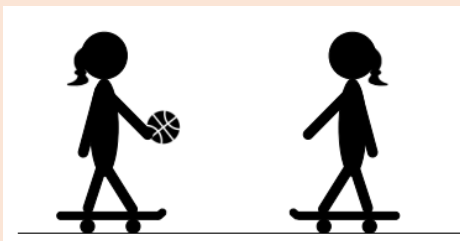
Physics Case

Question:

Ann and Betty are good at throwing and catching balls. They can always throw a ball for the other to catch it with only negligible motion. This time they want to have more fun and play the game on skateboards as shown. Initially they are both at rest. Ann first throws the ball to Betty at speed u . After catching the ball, Betty throws the ball back to Ann at the same speed. After Ann catches the ball, Betty is moving faster than Ann by how much? Assuming
 mass of the ball = m
 mass of Ann and the skateboard = mass of Betty and the skateboard = M .

Choices:

- A. $\frac{Mm+2m^2}{M(M+m)}u$
- B. $\frac{2m^2}{M(M+m)}u$
- C. $\frac{2m}{M+m}u$
- D. $-\frac{2m}{M+m}u$
- E. $\frac{m}{M}u$



Answer: B

Solution:

After Ann threw the ball to Betty:

$$\begin{aligned} mu &= Mv_A && \text{where } v_A \text{ is speed of Ann} \\ (M+m)v_B &= mu && \text{where } v_B \text{ is speed of Betty} \end{aligned}$$

After Betty threw the ball back to Ann:

$$\begin{aligned} Mv'_B - mu &= (M+m)v_B && \text{where } v'_B \text{ is new speed of Betty} \\ (M+m)v_A^{prime} &= mu + Mv_A && \text{where } v_A^{prime} \text{ is new speed of Ann} \end{aligned}$$

Solving,

$$\begin{aligned} v'_B &= \frac{2mu}{M} \\ v'_A &= \frac{2mu}{M+m} \end{aligned}$$

Hence,

$$v'_B - v'_A = 2mu \left(\frac{1}{m} - \frac{1}{M+m} \right) = 2mu \frac{m}{M(M+m)} = \frac{2m^2u}{M(M+m)}$$

Figure 13: An example question from physics.