

MV-MATH: Evaluating Multimodal Math Reasoning in Multi-Visual Contexts

Peijie Wang^{1,2}, Zhongzhi Li^{1,2}, Fei Yin^{1,2}, Dekang Ran^{1,2}, Chenglin Liu^{1,2*}

MAIS, Institute of Automation of Chinese Academy of Sciences¹

School of Artificial Intelligence, University of Chinese Academy of Sciences²

{wangpeijie2023, lizhongzhi2022, randekang2025}@ia.ac.cn

{fyin, liucl}@nlpr.ia.ac.cn

Abstract

Multimodal Large Language Models (MLLMs) have shown promising capabilities in mathematical reasoning within visual contexts across various datasets. However, most existing multimodal math benchmarks are limited to single-visual contexts, which diverges from the multi-visual scenarios commonly encountered in real-world mathematical applications. To address this gap, we introduce MV-MATH: a meticulously curated dataset of 2,009 high-quality mathematical problems. Each problem integrates multiple images interleaved with text, derived from authentic K-12 scenarios, and enriched with detailed annotations. MV-MATH includes multiple-choice, free-form, and multi-step questions, covering 11 subject areas across 3 difficulty levels, and serves as a comprehensive and rigorous benchmark for assessing MLLMs' mathematical reasoning in multi-visual contexts. Through extensive experimentation, we observe that MLLMs encounter substantial challenges in multi-visual math tasks, with a considerable performance gap relative to human capabilities on MV-MATH. Furthermore, we analyze the performance and error patterns of various models, providing insights into MLLMs' mathematical reasoning capabilities within multi-visual settings. The data and code: <https://eternal8080.github.io/MV-MATH.github.io/>.

1. Introduction

"The more you see, the more you know." — Leonardo da Vinci

Large Language Models (LLMs) have demonstrated exceptional performance across a wide range of tasks including natural language understanding, question answering and code generation [5, 29, 31, 38, 47]. Building on LLMs, MLLMs excel at various multimodal tasks by learning specialized features from different modalities, demonstrating a strong ability to understand different modalities,

and have achieved remarkable results in various visual-language tasks such as image captioning, visual question answering, and visual reasoning [4, 8, 15, 43, 52].

As MLLMs advance in various tasks, their mathematical reasoning abilities have garnered significant attention. Progressing from early text-based benchmarks like GSM8K [7] and MATH [14] to more recent multimodal benchmarks that incorporate visual elements, such as MathVista [28], MMMU [46], and GeoEval [48], MLLMs have demonstrated exceptional proficiency in mathematical reasoning, with the best-performing model on MathVista even surpassing human performance [34]. Most existing multimodal mathematics evaluation datasets, such as MathVista, MathVision [40], and MathVerse [50], are limited to single-visual contexts, requiring models to reason based on only a single image input. This raises an important question: are single-visual scenarios sufficient to fully capture the reasoning capabilities of MLLMs? Evidently, they are not.

Several multi-visual datasets in general domains have been introduced [12, 20, 39, 42, 45], which provide valuable resources. However, research specifically focused on mathematical evaluation within multi-visual scenarios remains limited. MathVerse-mv [18] extends the original MathVerse dataset by adapting problems and manually adding supplementary images to transform MathVerse into a multi-visual dataset. However, this approach to data construction limits the diversity of visual information and introduces potential biases. Additionally, MathVerse-mv is relatively small, containing only 788 samples. Another dataset, CMM-Math [26], focuses on Chinese contexts and includes 765 multi-image samples in its test set, but some images are of suboptimal quality, affecting the clarity of certain samples. Both MathVerse-mv and CMM-Math lack fine-grained categorizations and a diverse range of question types.

To overcome the limitations of existing multi-image mathematical reasoning datasets, which are constrained in both quantity and diversity, and to enable a more comprehensive evaluation of MLLMs in multi-visual contexts,

*Corresponding author

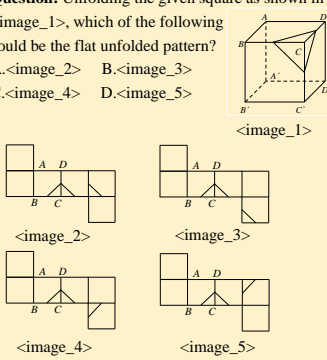
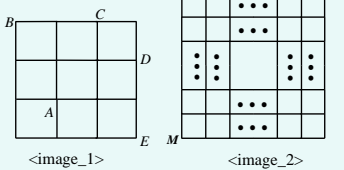
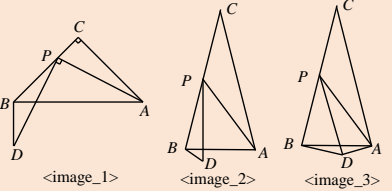
Multipile Choice	Free-Form:Single-Step	Free-Form:Multi-Step
<p>Question: Unfolding the given square as shown in <image_1>, which of the following could be the flat unfolded pattern?</p> <p>A.<image_2> B.<image_3> C.<image_4> D.<image_5></p>  <p><image_1> <image_2> <image_3> <image_4> <image_5> <image_6></p> <p>Answer: B Subject:Solid Geometry Difficulty: Medium Image Relevance: Mutually Dependent</p>	<p>Question: On a grid composed of squares with side length 1, each vertex of squares is called a <i>lattice point</i>. A movement from one <i>lattice point</i> to another is $\sqrt{5}$, units away is called a <i>knight's move</i>. In a 3×3 square grid (<image_1>), from point A, a <i>knight's move</i> can reach points B, C, D, E. Given a 25×25 square grid (<image_2>), the minimum number of <i>knight's moves</i> to reach the opposite vertex N from the M is _____.</p>  <p><image_1> <image_2></p> <p>Answer: 12 Subject:Combinatorics Difficulty: Medium Image Relevance: Mutually Dependent</p>	<p>Question: In an isosceles triangle ABC with $AC = BC$, point P is a point on side BC. Connect PA, and rotate segment PA clockwise around point P by an angle equal to $\angle C$, resulting in segment PD.</p> <p>(1) As shown in <image_1>, if $\angle C = 90^\circ$, measure $\angle DBA$. (2) As shown in <image_2>, if $\angle C = \alpha$, measure $\angle DBA$. (3) As shown in <image_3>, If $\angle C = \alpha$, $\tan \alpha = \frac{5}{12}$, $AC = 13$, $\angle APC = 135^\circ$. Find the length of AD.</p>  <p><image_1> <image_2> <image_3></p> <p>Answer: (1) 90° (2) $\angle DBA = \alpha$ (3) $10\sqrt{13}/13$ Subject:Analytic Geometry Difficulty: Hard Image Relevance: Independent</p>

Figure 1. Sampled MV-MATH examples from each question type. Each sample contains a multi-visual contexts. The specifics of image relevance are discussed in detail in Sec. 3.1

we present a novel dataset: MV-MATH. This dataset consists of 2,009 high-quality mathematics problems, each presented in interleaved multi-visual contexts with more than two images, covering a variety of grade 12 scenarios. Every problem has been meticulously curated from authentic K-12 sources. To ensure the quality and reliability of MV-MATH, each sample has undergone cross-validation by at least two annotators, ensuring the accuracy and precision of the problems, answers, and images. The dataset includes 1,109 multiple-choice questions and 900 free-form questions, of which 100 are multi-step problems, adding complexity and presenting a greater challenge for the models. The problems are categorized into three difficulty levels and span 11 distinct mathematical subjects, each annotated with image relevance, allowing for finer-grained analysis.

Extensive experiments are conducted on MV-MATH to comprehensively evaluate model performance in multi-visual mathematical contexts. On our MV-MATH dataset, Claude [2] achieves the highest performance with a score of 33.9%. Notably, the open-source model LLaVA-OneVision [17] achieves a competitive performance with a score of 26.2%, surpassing GPT-4v [30] and ranking just below Qwen-vl-max [3]. Nevertheless, the performance of all these models remains well below human-level capability. Through analysis of these results, we offer insights into the strengths and limitations of open-source models. In conclusion, the contributions of our study are as follows:

- We propose the MV-MATH benchmark, comprising 2,009 interleaved multi-visual mathematics problems derived from real K-12 scenarios, covering multiple-choice, free-form, and multi-step formats. All problems are classified across 3 difficulty levels and annotated with 11 fine-grained subjects.

- We evaluate 25 MLLMs on MV-MATH, revealing a substantial gap between current model performance and human-level capabilities in multi-visual mathematical tasks. Through an in-depth study of LLaVA-OneVision, we provide insights for enhancing model capabilities in multi-image mathematical reasoning.
- Our further investigation on MV-MATH reveals that MLLMs perform less effectively on image-dependent tasks compared to image-independent ones. Additionally, sequential image input outperforms the merged approach, and chain-of-thought (CoT) prompting does not consistently yield improvements.
- With the fine-grained categorization, we perform a comprehensive error analysis of current MLLMs, offering in-depth insights to inform future advancements.

2. Related Work

2.1. Benchmark for Mathematical Reasoning

Pure Text Benchmarks MathQA [1] is a large-scale benchmark focused on math word problems, designed to evaluate models' capabilities in solving complex arithmetic and algebraic tasks through natural language. GSM8K [7], developed by OpenAI, assesses LLMs' mathematical reasoning with 8,500 grade-school-level problems requiring 2 to 8 steps, covering arithmetic, fractions, and basic algebra. MATH [14] offers a more challenging dataset of 12,500 high school competition problems, surpassing GSM8K in difficulty. MathBench [25] includes 3,709 problems, ranging from basic arithmetic to college-level questions and spanning multiple difficulty levels. SuperCLUE-Math [44], a Chinese benchmark, evaluates multi-step reasoning with over 2,000 problems, each requiring multiple steps and ac-

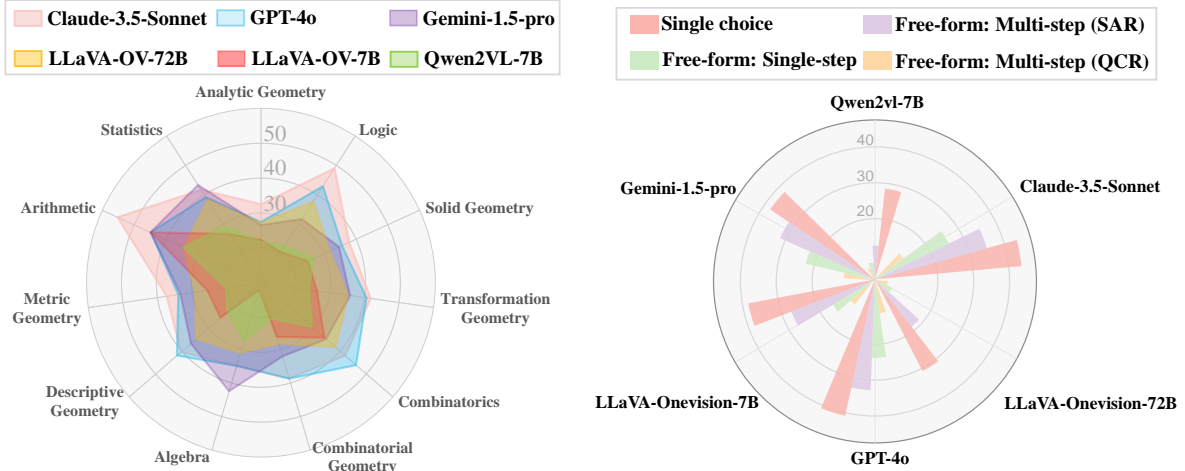


Figure 2. Performance comparison of six Multimodal Large Language Models (MLLMs) on our proposed MV-MATH dataset across 11 subjects (left) and 3 question types (right). SAR: Step Accuracy Rate, QCR: Question Completeness Rate.

accompanied by natural language solutions.

Multimodal Benchmarks With the advancement of MLLMs, several high-quality benchmarks have been developed to evaluate mathematical problem-solving abilities in visual contexts. MathVista [28] is the first multimodal dataset aimed at assessing MLLMs’ mathematical reasoning, specifically focusing on visual math QA tasks. MATH-Vision [40] targets multimodal mathematical understanding, featuring questions primarily sourced from mathematics competitions and requiring deep reasoning. MathVerse [50] tests whether current MLLMs can effectively leverage diagram information in mathematical reasoning tasks, offering a collection of 2,612 visual diagrams. GeoEval [48] focuses on geometry, examining the capabilities of large models in this field through 2,000 geometric samples. CMMath [21] primarily evaluates mathematical reasoning abilities in the context of Chinese K-12 education.

All of the above datasets are restricted to single-visual contexts, relying on a single image input per problem. In contrast, our MV-MATH dataset emphasizes multi-visual mathematical reasoning by incorporating multiple images per problem. Addressing multi-image problems requires models to develop a more comprehensive understanding of relationships between mathematical visual elements, aligning more closely with real-world mathematical scenarios.

2.2. Multimodal Foundation Models

The progress in visual-language alignment and rapid evolution of LLMs have endowed these models with visual capabilities, with notable achievements including CLIP [35] and BLIP [19]. CLIP aligns image-text pairs via contrastive learning, while BLIP combines contrastive and generative learning to align and generate image-text pairs. Inspired by these, MiniGPT-4 [53], LLaMA-Adapter [49],

and LLaVA [24] extend text-only models into multimodal ones through vision-language alignment and instruction tuning. Subsequently, large multimodal models emerged, such as the closed-source GPT-4V [30], Gemini [37], Claude [2], GPT-4o [32] and open-source DeepSeek-VL [27], Qwen-VL [3], InternLMXComposer-VL [10], LLaVA-NeXT [23], and LLaVA-OneVision [17], which perform well on benchmarks like MMMU [46] and MathVista [28]. In mathematics, G-LLaVA [13] and Math-LLaVA [36] enhance performance on mathematical tasks using large instruction datasets Geo170K [13] and MathV360K [36]. MAVIS [51] improves visual perception with three-stage fine-tuning, while MultiMath [33] refines visual-language alignment through a four-stage training process.

3. The MV-MATH Benchmark

3.1. Overview

We introduce the MV-MATH benchmark, a meticulously annotated dataset designed to evaluate the mathematical reasoning capabilities of MLLMs in multi-visual contexts. Each sample in MV-MATH consists of interleaved multi-image and text, closely reflecting the multimodal distribution found in real-world scenarios. This interleaved format imposes greater demands on modal fusion and comprehension in MLLMs, providing a more realistic assessment of their performance [16, 39].

The MV-MATH dataset comprises 2,009 multi-image questions, with some questions containing up to 8 images. It includes three types: multiple-choice, free-form and multi-step questions. MV-MATH is organized into 11 subjects over 3 difficulty levels, including *Analytic Geometry*, *Algebra*, *Metric Geometry*, *Combinatorics*, *Transformation Ge-*

Statistics	Number
Total Questions	2009
*multiple-choice questions	1109
*Free-form questions	900
-one-step questions	800
-multi-step questions	100
Questions in the testmini set	200
Difficulties (Easy: Medium: Hard)	27%:48%:25%
Total images	6061
* question with 2 images	979 (48.73%)
* question with 3 images	312 (15.53%)
* question with 4 images	453 (22.55%)
* question with 5 images and above	265 (13.39%)
Image Relevance	
* Mutually Dependent questions	1412
* Independent questions	597
Minimum question length	14
Maximum question length	383
Average question length	80.17

Table 1. Key statistics of MV-MATH. The unit of question length is words.

ometry, Logic, Solid Geometry, Arithmetic, Combinatorial Geometry, Descriptive Geometry and Statistics, covering a range of scenarios from the K-12 mathematics curriculum. Based on image relevance, we categorize MV-MATH into two subsets: a mutually dependent set (**MD**), where images are interrelated and understanding one image necessitates information from another (left and middle of Figure 1); and an independent set (**ID**), where images are unrelated and can be interpreted independently without reference to other images (right of Figure 1). To facilitate analysis, we provide a test-mini version of MV-MATH containing 200 samples, preserving similar distribution of difficulty levels and question types as the full MV-MATH dataset. Detailed statistics and coverage of MV-MATH are presented in Table 1.

3.2. Data Construction

Data collection. We initially collect over 300,000 mathematics problems spanning grade 12 scenarios, comprising multiple-choice and free-form questions, including plain text, single-image, and multi-image questions, all stored in PDF format. Using the Mathpix API¹, we extract text content—such as questions, answers, and analyses—as well as images, and organized this data in JSON format. Based on image counts, we identified a preliminary set of 49,538 multi-image problems.

Data filtering. To ensure the high quality of the multi-image mathematical data, we employ a three-stage data screening strategy. In the first stage, we verify the alignment between text and images within each question, this is

¹<https://mathpix.com/convert>

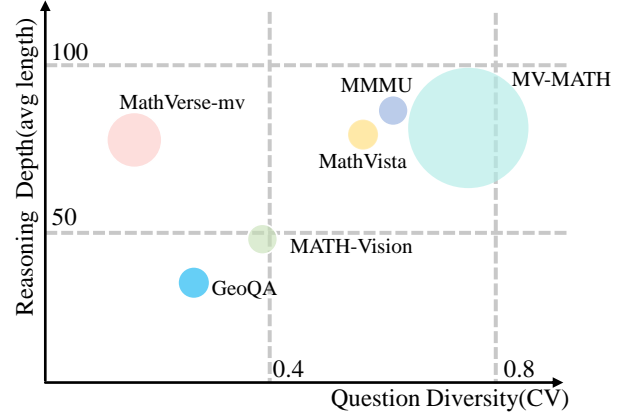


Figure 3. Comparison of MV-MATH with other existing benchmarks. The size of each circle represents the richness of image distribution per question. The definition and calculation method of CV are provided in Section 3.3

a crucial step since Mathpix does not always yield the correct number of images. This initial filtering retains 35,562 samples. In the second stage, we inspect the remaining samples for missing text fields or semantic inaccuracies, then categorize them into two subsets: 10,110 multiple-choice and 8,933 free-form questions for further screening. In the third stage, we manually filter out samples with low-quality images (e.g., blurry images or images containing text content). This process ultimately yields 1,109 multiple-choice questions, 800 one-step free-form questions, and 100 multi-step free-form questions. Each step was cross-verified by at least two graduate students to ensure the reliability of the screening outcomes.

Data labeling. Using the above samples, we conduct difficulty grading, subject categorization, and image relevance classification. For each problem, difficulty is categorized by analyzing the lengths of the question and analysis fields, assigning weights of 0.4 and 0.6 respectively. Based on the weighted lengths, MV-MATH is divided into three levels: easy, medium, and hard. For subject categorization and image relevance, we apply a majority voting approach, leveraging classification results from GPT-4o, Claude-3.5-Sonnet, and Qwen-VL-Max to determine each problem’s subject and its degree of image dependence. We then manually validate the assigned difficulty levels, subjects, and image relevance, with two annotators independently reviewing each problem and making adjustments to correct any misclassifications, ensuring the accuracy of the final classification.

3.3. Comparison with Existing Benchmarks

Most existing multimodal mathematical reasoning datasets are limited to single-visual contexts, where the model interprets only a single image without the need for information interaction between images. To further distinguish the

Models	Easy (546)	Medium (968)	Hard (495)	Overall (2009)
Claude-3.5-sonnet [2]	35.7	37.5	26.6	33.9
GPT-4o [32]	40.3	32.7	22.9	32.1
Gemini-1.5-pro [37]	34.1	30.0	20.3	29.0
Qwen-vl-max [3]	32.7	27.7	20.1	26.9
GPT-4V [30]	31.6	24.7	17.1	24.5
LLaVA-OneVision-Chat-72B [17]	34.6	26.0	19.2	26.2
LLaVA-OneVision-Chat-7B [17]	20.6	21.2	14.8	19.0
Qwen2VL-Instruct-7B [41]	18.8	17.1	13.9	16.5
LLaVA-NeXT-Interleave-7B [18]	15.3	16.1	10.9	14.7
Deepseek-VL-Chat-7B [27]	17.1	14.1	12.5	14.5

Table 2. Result decomposition across question difficulty levels.

difference between MV-MATH and other existing ones, we elaborate the benchmark details in Figure 3. Here, we primarily compare our MV-MATH with MathVerse-mv [18] and CMM-MATH [28].

Comparing with MathVerse-mv MathVerse-mv [18] is derived by reformatting original problems in MathVerse through question rephrasing and image combination, resulting in 788 generated problems. However, all questions in MathVerse-mv are multiple-choice and primarily focus on a single task type: *‘select the correct paired math diagram from given images.’* This dataset lacks differentiation in difficulty levels and subject categories. Such artificial rephrasing and image augmentation introduce distribution bias, limiting its representation of authentic multi-image mathematical scenarios. In comparison, our MV-MATH dataset provides 2,009 high-quality samples from real K-12 contexts, including both multiple-choice and free-form questions, each classified by fine-grained difficulty and subject. Sample lengths in MathVerse-mv range from 62 to 246, with an average of 81.91, whereas MV-MATH samples range from 14 to 383, with an average of 80.17. To assess the diversity in sample length distribution, we calculate the coefficient of variation (CV), with a higher CV indicating greater distribution richness. The CVs for MathVerse-mv and MV-MATH are 0.19 and 0.74 respectively, highlighting the broader distribution diversity of MV-MATH. The CV calculation is as follows:

$$CV = \frac{\sigma}{\mu} \times 100\% \quad (1)$$

where σ denotes the standard deviation of sample lengths, and μ denotes the mean sample length.

Comparing with CMM-Math CMM-Math [26] is a dataset designed for Chinese-language scenarios, with the multi-image portion containing 765 samples. Some images in CMM-Math are of lower quality, due to Mathpix’s filtering inconsistencies, some images contain excessive text. In contrast, we performed manual checks on each image in our samples to ensure that every sample in MV-MATH is of reliable quality. Additionally, MV-MATH primarily

Models	Original	CoT	CoT & 2-shot
Claude-3.5-sonnet [2]	29.2	32.6(+3.4)	33.9(+1.3)
GPT-4o [32]	31.8	30.9(-0.9)	32.1(+1.2)
Gemini-1.5-pro [37]	29.8	28.3(-1.5)	29.1(+0.8)
Qwen-vl-max [3]	28.4	29.1(+0.7)	26.9(-2.2)
GPT-4V [30]	23.4	25.1(+1.7)	24.5(-0.6)
LLaVA-OneVision-Chat-72B [17]	27.3	26.7(-0.6)	26.2(-0.5)
LLaVA-OneVision-Chat-7B [17]	20.5	20.1(-0.4)	19.1(-1.0)
LLaVA-NeXT-Interleave-7B [18]	17.9	16.3(-1.6)	14.7(-1.6)
Qwen2VL-Instruct-7B [41]	19.7	16.7(-2.0)	16.5(-0.2)
Deepseek-VL-Chat-7B [27]	17.2	15.7(-1.5)	14.5(-1.2)

Table 3. Model Performance Evaluation for Original, CoT, and CoT with 2-shot configurations.

focuses on English-language scenarios. To investigate the impact of language on multi-image mathematical reasoning performance, we also provide a translated Chinese version of MV-MATH.

4. Experiments

We conduct extensive experiments on MV-MATH to evaluate a diverse range of models, including 18 open-source and 7 API-based models, assessing both LMMs and MLLMs. Our evaluation encompasses three conditions: text-only, text with image captions, and text with images. Results indicate that models perform better when images are input sequentially rather than merged, with mutually dependent image tasks proving more challenging. Experimental outcomes show that even the top-performing model, Claude 3.5-Connect, falls significantly short of human-level performance, underscoring considerable room for improvement in mathematical multi-visual reasoning capabilities.

Our evaluation is conducted under three settings: original answering, Chain of Thought (CoT), and CoT with 2-shot. To establish a human performance baseline, we use the testmini set and recruit 50 high school students to independently complete the questions. For multiple-choice, single-step, and multi-step free-form questions, we carefully design specific prompts to ensure that models generated responses in the correct format. Answer extraction and evaluation are performed using the Deepseek API², with prompts specifically crafted for each question type to optimize model evaluation accuracy, more detailed information can be found in the Appendix.

4.1. Main Results

Challenging Nature of MV-MATH: The data presented in Table 4 illustrates the inherent challenges of MV-MATH. Notably, the overall accuracy of the leading model, Claude-3.5-Sonnet, is only 33.9%, which falls significantly short of the human level accuracy of 76.5%. Among the open-source models, only LLaVA-OneVision-72B (LLaVA-OV-

²<https://www.deepseek.com/>

Model	Overall	AG	Algebra	MG	Combinatorics	TG	Logic	SG	Arithmetic	CG	DG	Statistics
LLMs(Text-only, CoT with 2-shot)												
Deepseek-Chat[9]	16.1	14.7	10.1	16.3	18.7	25.3	21.8	18.3	6.2	19.0	8.2	13.2
LLMs(Text + Image Caption, CoT with 2-shot)												
Deepseek-Chat[9]	16.9	14.1	8.0	15.3	13.4	13.2	12.7	9.4	14.4	12.8	18.0	17.3
Open-source MLLMs (Text + Image, CoT with 2-shot)												
Math-LLaVA-13B[36]	3.0	1.6	6.9	4.7	4.8	2.9	0.0	3.2	18.7	6.6	2.1	5.8
LLaVA-v1.5-13B[24]	5.0	4.8	6.8	4.1	4.8	8.7	9.0	3.5	12.5	5.1	5.0	11.7
LLaVA-v1.5-7B[24]	10.3	9.3	11.7	11.2	9.7	12.8	13.6	10.2	0.0	7.7	23.7	11.7
VILA-13B[22]	12.0	11.5	11.0	11.0	12.1	14.4	18.1	13.2	37.5	10.6	20.8	5.8
InternLM-XComposer2.5-VL-7B[10]	13.1	12.2	12.6	13.2	24.3	20.6	36.3	9.4	18.7	11.1	23.7	17.6
InternVL-Chat-8B[6]	14.4	14.1	20.4	17.5	19.5	19.6	27.2	13.0	31.2	9.9	20.1	23.5
Llama-3.2-Vision-Instruct-11B[11]	14.4	15.0	15.4	16.2	23.1	15.6	18.1	11.9	31.2	14.3	25.1	17.6
Deepseek-VL-7B[27]	14.5	14.8	20.2	10.8	17.0	19.8	9.0	15.1	18.7	10.9	26.6	29.4
LLaVA-NeXT-Interleave-7B[18]	14.7	14.0	15.5	15.2	17.0	18.2	18.1	16.3	6.2	14.1	24.4	23.5
Mantis-Idedics2-8B[16]	15.5	13.3	17.7	19.2	14.6	20.9	22.7	12.3	6.2	16.0	26.6	11.7
Mantis-siglip-8B[16]	15.8	17.9	17.7	17.9	14.6	20.4	22.7	12.1	18.7	10.8	32.3	17.6
Qwen2VL-Instruct-7B[41]	16.5	14.2	18.6	14.8	17.0	21.9	22.7	17.2	31.2	16.1	25.1	23.5
LLaVA-OneVision-SI-7B[17]	17.2	16.1	19.5	13.2	16.0	19.5	12.6	15.0	36.5	13.2	31.3	13.6
LLaVA-OneVision-SFT-7B[17]	18.8	18.2	20.3	22.3	17.3	20.1	9.0	15.8	43.1	15.8	27.3	23.5
LLaVA-OneVision-Chat-7B[17]	19.1	19.6	20.4	21.4	14.6	18.8	4.5	20.4	43.7	16.7	28.9	29.4
LLaVA-OneVision-SI-72B[17]	25.0	24.7	24.3	27.6	27.0	25.3	37.9	24.4	37.1	20.4	31.2	23.5
LLaVA-OneVision-SFT-72B[17]	25.9	24.2	31.3	21.1	23.1	28.9	31.8	32.8	18.7	21.5	39.5	29.4
LLaVA-OneVision-Chat-72B[17]	26.2	25.1	32.4	23.9	35.3	28.1	27.2	31.6	31.2	22.6	35.9	35.2
Closed-source MLLMs (Text + Image, CoT with 2-shot)												
Qwen-vl-plus[3]	19.7	17.9	24.1	22.0	16.0	19.9	24.8	15.9	15.2	18.7	31.4	29.4
GPT-4V[30]	24.5	18.7	31.6	32.4	25.6	26.3	36.3	26.8	43.7	19.3	33.8	35.2
Qwen-vl-max[3]	26.9	27.6	32.1	24.7	36.5	29.6	31.8	30.9	37.5	23.7	32.3	23.5
Gemini-1.5-Pro[37]	29.1	29.9	32.9	28.3	28.0	30.5	40.5	33.9	42.7	21.7	30.6	35.2
GPT-4o[32]	32.1	28.7	36.7	34.4	39.4	30.6	29.8	38.2	41.7	20.8	44.3	47.0
Claude-3.5[2]	33.9	32.7	38.1	34.3	46.7	33.3	29.8	36.3	54.2	27.0	38.2	41.1
Human Performance												
Human (testmini)	76.5	70.2	74.9	91.2	83.1	67.4	80.2	62.5	83.2	71.7	68.5	85.6

Table 4. Comparison of model performances across various mathematical subjects. AG: Analytic Geometry, MG: Metric Geometry, TG: Transformation Geometry, SG: Solid Geometry, CG: Combinatorial Geometry, DG: Descriptive Geometry. The first and second highest accuracy of LLMs are marked in red and blue, respectively.

72B) achieves an accuracy above 20% on MV-MATH, while all others fall below this threshold. Notably, Math-LLaVA, a specialized mathematical MLLM, demonstrates limited performance on the MV-MATH dataset, with an accuracy of only 3.0%. Our analysis suggests that this under-performance stems from the distribution of its training data. Specifically, Math-LLaVA trained on the Math360K dataset do not effectively generalize to the MV-MATH dataset, underscoring a significant limitation in its generalization capabilities. A more detailed analysis can be found in the Appendix. This finding underscores a broader challenge for current mathematical MLLMs: the need to achieve robust performance on both in-distribution and out-of-distribution data. Future advancements should prioritize strengthening generalization capabilities over simply memorizing training data.

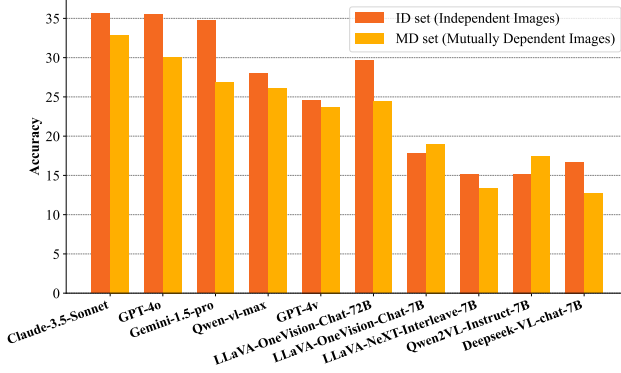
Disparity of Closed-source and Open-source Models:

Most open-source models exhibit a substantial performance gap compared to closed-source counterparts. Notably, Claude-3.5-Sonnet achieves the highest accuracy on the MV-MATH dataset at 33.9%, while LLaVA-OV-Chat-72B

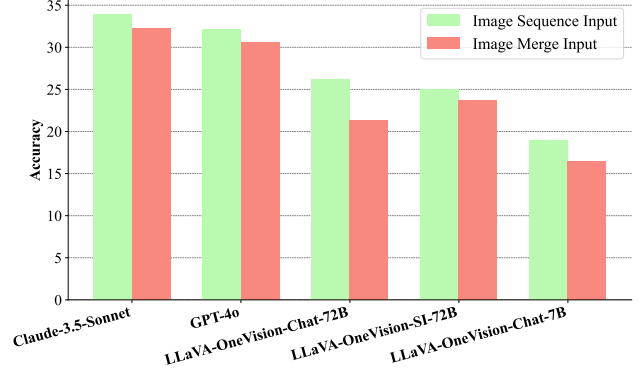
attains a commendable 26.2%, thereby narrowing the performance disparity between open-source and closed-source models. Models with smaller parameter sizes, such as the 7B and 13B variants, generally attain accuracy levels between 10% and 20%, with the top-performing LLaVA-OV-Chat-7B reaching 19.0%.

Comparison among Different Subjects: The dataset encompasses 11 distinct mathematical subjects, ranging from algebra to geometry, each characterized by unique features. Model performance tends to be weaker in subjects requiring a higher level of image understanding, such as combinatorial geometry, where even the most advanced model, Claude-3.5-Sonnet, achieves an accuracy of only 27.0%. In contrast, Claude’s accuracy in arithmetic reaches 54.2%. This disparity indicates that current MLLMs lack the necessary capabilities to effectively process complex images and comprehend relationships between multiple intricate images. Significant improvements are still needed in fine-grained visual tasks, particularly in geometric perception and understanding.

Failure of CoT/few-shot: As shown in Table 3, we ob-



(a) Performance Evaluation of Models on ID and MD Sets.



(b) Comparison of Accuracy with Merge and Sequence Image Inputs.

Figure 4. Evaluation of Model Performance across Different Image Relevance (ID vs. MD) and Input Methods (Merged vs. Sequential).

serve that CoT and few-shot do not always improve model performance on MV-MATH. While Claude performs 3.4% better with CoT and 1.3% better with 2-shot based on CoT, GPT4o performs best under CoT with 2-shot setting, while Qwen-vl-max and gpt-4v perform best under CoT setting. For Gemini-1.5-pro and all open-source models, the original versions outperform CoT and CoT with 2-shot prompts. Notably, for all open-source models, performance declines steadily with the introduction of CoT and 2-shot prompting. We conduct further analysis on different types of questions in Appendix, and we find that original prompts perform better for multiple-choice questions, while CoT tends to produce better results for free-form questions. Adding 2-shot on top of CoT tends to degrade model performance on MV-MATH. Similar results are observed in Math-Vision [40].

4.2. Analysis of Question type and Difficulty

Different Question Types. We compare the performance of various models on three types of questions. As shown in Figure 2(right), the accuracy of the model on multiple-choice questions is significantly higher than on free-form questions. We design two evaluation metrics for multi-step questions: Step Accuracy Rate (SAR) and Question Completeness Rate (QCR). SAR refers to the proportion of correctly answered steps out of the total steps, while QCR is the proportion of questions for which all steps were answered correctly. The model with the best performance on the QCR metric is GPT-4o, achieving only 6%, while its corresponding SAR is 32.6%. This reflects the model’s insufficient ability to perform complex reasoning on open and multi-step questions, and also reveals that the multiple-choice question format may not truly reflect the model’s actual reasoning and problem-solving ability, as they usually rely on prompts to identify options.

Different Difficulty Levels. Table 2 compares the performance of selected models across three difficulty levels. GPT-4o demonstrates significantly higher proficiency with a success rate of 40.3% in the “Easy” category. For

the “Medium” category, Claude achieves the best performance at 37.5%. The performance disparity further decreases in the “Hard” category, where the leading Claude model reaches only 26.6% accuracy. As the difficulty level increases, the performance gap between models becomes smaller. Of particular note, Claude is the only model whose performance on medium difficulty exceeds that on low difficulty, whereas the performance of other models progressively declines as the difficulty level increases.

4.3. Impact of Image Relevance and Input Methods

Image Relevance, Mutually Dependent vs. Independent: To facilitate a more detailed analysis of MLLMs on mathematical multi-visual tasks, we conduct experiments on mutually dependent sets (MD) and independent sets (ID) (as detailed in Section 3.1) in MV-MATH. While the ID set includes multiple images, they are not interrelated. In contrast, the MD set comprises mutually dependent images, demanding a higher level of cross-image understanding. As shown in Figure 4a, with the exception of Qwen2VL-7B and LLaVA-OV-Chat-7B, all models exhibit lower performance on the MD set compared to the ID set, with Gemini-1.5-pro demonstrating the largest performance gap between the two, reaching 7.8%. This observation suggests that most models face challenges in effectively handling interdependent image tasks in mathematical scenarios, highlighting the potential limitations of handling cross-image interdependencies in mathematical multi-visual contexts.

Input Methods, Merge vs. Sequential: To investigate the impact of image input methods on model performance, we apply both merge and sequence input approaches on the MV-MATH dataset, with results presented in Figure 4b. The findings consistently demonstrate that the sequence input method outperforms the merge input method across all tested models, indicating that preserving the positional and sequential information of images is crucial for effective multi-image mathematical reasoning. This superior performance of sequence input underscores the importance of

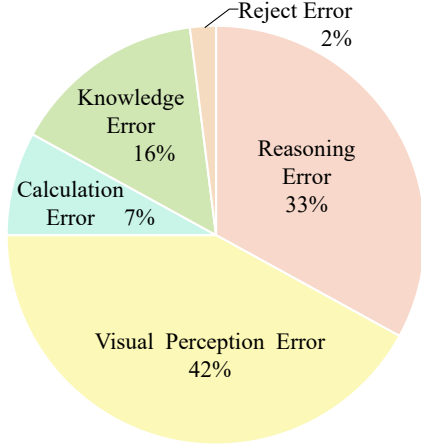


Figure 5. Error distribution over 100 annotated Claude errors.

structured visual information in enhancing model capabilities for interpreting and processing complex mathematical scenarios. An interesting observation is that, while the LLaVA-OV-SI-72B model underperforms compared to the LLaVA-OV-Chat-72B model on sequence input method, it outperforms the latter on merge input method. This shift may be attributed to the initial single-image training of LLaVA-OV-SI, which enhances its ability to retain distinct visual features within individual images when processed in a merged format.

5. Error Analysis and Model Insights

Error Analysis We conduct a detailed error analysis of the best-performing model, Claude-3.5-Sonnet. We randomly sample 100 errors, examine their distribution in detail, and provide corresponding qualitative examples. Based on this analysis, we classify the errors into five types, with the specific categories and their distribution shown in Figure 5.

Visual perception error is the most common type, accounting for 42%, which is higher than that of most single-image datasets such as MATH-Vision and MMMU [40, 46]. This suggests that perceiving multi-visual contexts in the MV-MATH dataset presents a greater challenge compared to single-visual contexts. An example of a visual perception error in MV-MATH is shown in Figure 6. Multi-visual perception requires MLLMs to not only accurately perceive and understand each image individually but also to comprehend the relationships between them. Reasoning errors also constitute a significant proportion. Even when the model correctly interprets text and images, it often fails to apply logical and mathematical reasoning effectively, leading to incorrect answers due to flaws in reasoning. Knowledge errors, stemming from the model’s lack of relevant expertise, represent the third most common error type on MV-MATH,

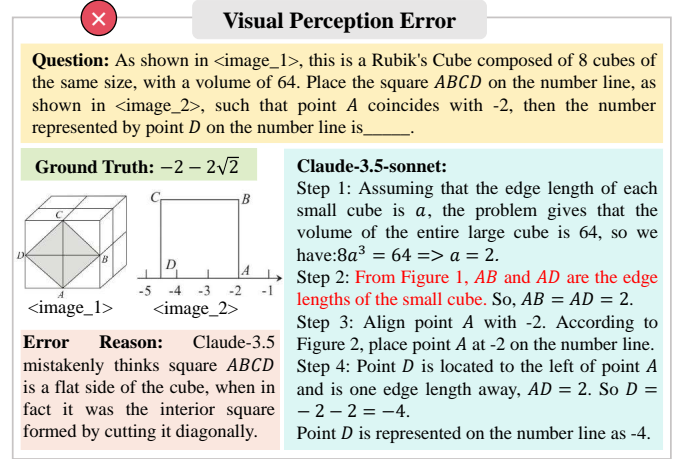


Figure 6. A basic visual perception error, with the error highlighted in red. More examples can be found in the appendix.

accounting for 7% of errors. Calculation errors and reject errors make up 7% and 2%, respectively. Notably, we observe that all calculation errors are associated with symbolic computation rather than numerical computation.

Model Insights We analyze LLaVA-OV, the best-performing open-source model on MV-MATH, to gain insights for enhancing model capabilities in multi-visual mathematical contexts. LLaVA-OV’s strong performance is attributed to its staged training strategy and innovative architecture, which progressively increase the complexity of visual signals, thereby enhancing generalization across multi-image and multimodal tasks. As shown in Table 4, LLaVA-OV-SI consistently underperforms compared to LLaVA-OV-SFT, while LLaVA-OV-Chat achieves the highest accuracy across both 7B and 72B model sizes. LLaVA-OV-SI is trained with a single-image stage; LLaVA-OV-SFT incorporates additional multi-image training to enhance image-text comprehension; and LLaVA-OV-Chat, building on LLaVA-OV-SFT, leverages Direct Preference Optimization (DPO) and human feedback to further improve generalization and reasoning in multi-visual contexts. The combination of multi-image training, DPO, and human feedback significantly advances multi-image mathematical reasoning, strengthening both generalization and inference capabilities in multi-visual scenarios.

6. Conclusion

The introduction of the MV-MATH dataset represents a major advancement in evaluating MLLMs’ multi-visual mathematical reasoning capabilities, addressing a critical gap in the assessment of multi-image mathematical tasks. MV-MATH encompasses a diverse and challenging array of multi-visual mathematical content. Using this dataset, we conduct a systematic and comprehensive evaluation of nu-

merous open-source and closed-source models, establishing an extensive benchmark. This benchmark not only reveals a substantial performance gap between human capabilities and current MLLMs in multi-visual mathematical reasoning but also highlights the impact of image relevance, image input methods and question types on multi-visual math performance. Additionally, insights derived from the LLaVA-OV model offer guidance for enhancing performance in multi-image mathematical tasks, providing valuable direction for future advancements in this field.

References

- [1] Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. Mathqa: Towards interpretable math word problem solving with operation-based formalisms. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2357–2367, 2019. 2
- [2] Anthropic. The claude 3 model family: Opus, sonnet, haiku. <https://www.anthropic.com/claude-3-model-card>, 2024. Claude-3 Model Card. 2, 3, 5, 6, 23, 24, 25, 26, 27
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond, 2023. 2, 3, 5, 6, 23, 24, 25, 26, 27
- [4] Tianyi Bai, Hao Liang, Binwang Wan, Ling Yang, Bozhou Li, Yifan Wang, Bin Cui, Conghui He, Binhang Yuan, and Wentao Zhang. A survey of multimodal large language model from a data-centric perspective. *arXiv preprint arXiv:2405.16640*, 2024. 1
- [5] Jiaqi Chen, Jianheng Tang, Jinghui Qin, Xiaodan Liang, Lingbo Liu, Eric P. Xing, and Liang Lin. Geoqa: A geometric question answering benchmark towards multimodal numerical reasoning. In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, pages 513–523. Association for Computational Linguistics, 2021. 1
- [6] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*, 2023. 6, 23, 24, 25, 26
- [7] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021. 1, 2
- [8] Can Cui, Yunsheng Ma, Xu Cao, Wenqian Ye, Yang Zhou, Kaizhao Liang, Jintai Chen, Juanwu Lu, Zichong Yang, Kuei-Da Liao, Tianren Gao, Erlong Li, Kun Tang, Zhipeng Cao, Tong Zhou, Ao Liu, Xinrui Yan, Shuqi Mei, Jianguo Cao, Ziran Wang, and Chao Zheng. A survey on multimodal large language models for autonomous driving. In *WACV (Workshops)*, pages 958–979. IEEE, 2024. 1
- [9] DeepSeek-AI. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model, 2024. 6, 23, 24, 25, 26
- [10] Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Bin Wang, Linke Ouyang, Xilin Wei, Songyang Zhang, Haodong Duan, Maosong Cao, Wenwei Zhang, Yining Li, Hang Yan, Yang Gao, Xinyue Zhang, Wei Li, Jingwen Li, Kai Chen, Conghui He, Xingcheng Zhang, Yu Qiao, Dahua Lin, and Jiaqi Wang. Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. *CoRR*, abs/2401.16420, 2024. 3, 6, 23, 24, 25, 26
- [11] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 6, 23, 24, 25, 26
- [12] Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A. Smith, Wei-Chiu Ma, and Ranjay Krishna. BLINK: multimodal large language models can see but not perceive. In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XXIII*, pages 148–166. Springer, 2024. 1
- [13] Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wanjun Zhong, Yufei Wang, Lanqing Hong, Jianhua Han, Hang Xu, Zhenguo Li, and Lingpeng Kong. G-llava: Solving geometric problem with multi-modal large language model. *CoRR*, abs/2312.11370, 2023. 3
- [14] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*, 2021. 1, 2
- [15] Yupan Huang, Zaiqiao Meng, Fangyu Liu, Yixuan Su, Nigel Collier, and Yutong Lu. Sparkles: Unlocking chats across multiple images for multimodal instruction-following models. *CoRR*, abs/2308.16463, 2023. 1
- [16] Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max Ku, Qian Liu, and Wenhui Chen. MANTIS: interleaved multi-image instruction tuning. *CoRR*, abs/2405.01483, 2024. 3, 6, 23, 24, 25, 26
- [17] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *CoRR*, abs/2408.03326, 2024. 2, 3, 5, 6, 23, 24, 25, 26, 27
- [18] Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv preprint arXiv:2407.07895*, 2024. 1, 5, 6, 23, 24, 25, 26, 27
- [19] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Internat-*

- tional conference on machine learning, pages 12888–12900. PMLR, 2022. 3
- [20] Juncheng Li, Kaihang Pan, Zhiqi Ge, Minghe Gao, Wei Ji, Wenqiao Zhang, Tat-Seng Chua, Siliang Tang, Hanwang Zhang, and Yueting Zhuang. Fine-tuning multimodal llms to follow zero-shot demonstrative instructions. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 1
- [21] Zhong-Zhi Li, Ming-Liang Zhang, Fei Yin, Zhi-Long Ji, Jin-Feng Bai, Zhen-Ru Pan, Fan-Hu Zeng, Jian Xu, Jia-Xin Zhang, and Cheng-Lin Liu. Cmmath: A chinese multimodal math skill evaluation benchmark for foundation models. *CoRR*, abs/2407.12023, 2024. 3
- [22] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. VILA: on pre-training for visual language models. In *CVPR*, pages 26679–26689. IEEE, 2024. 6, 23, 24, 25, 26
- [23] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024. 3
- [24] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 3, 6, 23, 24, 25, 26
- [25] Hongwei Liu, Zilong Zheng, Yuxuan Qiao, Haodong Duan, Zhiwei Fei, Fengzhe Zhou, Wenwei Zhang, Songyang Zhang, Dahua Lin, and Kai Chen. Mathbench: Evaluating the theory and application proficiency of llms with a hierarchical mathematics benchmark. In *ACL (Findings)*, pages 6884–6915. Association for Computational Linguistics, 2024. 2
- [26] Wentao Liu, Qianjun Pan, Yi Zhang, Zhuo Liu, Ji Wu, Jie Zhou, Aimin Zhou, Qin Chen, Bo Jiang, and Liang He. Cmm-math: A chinese multimodal math dataset to evaluate and enhance the mathematics reasoning of large multimodal models, 2024. 1, 5
- [27] Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yaofeng Sun, Chengqi Deng, Hanwei Xu, Zhenda Xie, and Chong Ruan. Deepseek-vl: Towards real-world vision-language understanding, 2024. 3, 5, 6, 23, 24, 25, 26, 27
- [28] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In *ICLR*. OpenReview.net, 2024. 1, 3, 5
- [29] Shaghayegh Mobasher, Ghazal Zamaninejad, Maryam Hashemi, Melika Nobakhtian, and Sauleh Eetemadi. Parsvqa-caps: A benchmark for visual question answering and image captioning in persian. *people*, 101:404. 1
- [30] OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023. 2, 3, 5, 6, 23, 24, 25, 26, 27
- [31] OpenAI. Chatgpt, 2023. 1
- [32] OpenAI. Hello gpt-4o, 2024. 3, 5, 6, 23, 24, 25, 26, 27
- [33] Shuai Peng, Di Fu, Liangcai Gao, Xiuqin Zhong, Hongguang Fu, and Zhi Tang. Multimath: Bridging visual and mathematical reasoning for large language models. *CoRR*, abs/2409.00147, 2024. 3
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, pages 8748–8763. PMLR, 2021. 1
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3
- [36] Wenhao Shi, Zhiqiang Hu, Yi Bin, Junhua Liu, Yang Yang, See-Kiong Ng, Lidong Bing, and Roy Ka-Wei Lee. Mathllava: Bootstrapping mathematical reasoning for multimodal large language models. *CoRR*, abs/2406.17294, 2024. 3, 6, 23, 24, 25, 26
- [37] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 3, 5, 6, 23, 24, 25, 26, 27
- [38] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 1
- [39] Fei Wang, Xingyu Fu, James Y Huang, Zekun Li, Qin Liu, Xiaogeng Liu, Mingyu Derek Ma, Nan Xu, Wenxuan Zhou, Kai Zhang, et al. Muirbench: A comprehensive benchmark for robust multi-image understanding. *arXiv preprint arXiv:2406.09411*, 2024. 1, 3
- [40] Ke Wang, Juntong Pan, Weikang Shi, Zimu Lu, Mingjie Zhan, and Hongsheng Li. Measuring multimodal mathematical reasoning with math-vision dataset. *CoRR*, abs/2402.14804, 2024. 1, 3, 7, 8
- [41] 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. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 5, 6, 23, 24, 25, 26, 27
- [42] Xiyao Wang, Yuhang Zhou, Xiaoyu Liu, Hongjin Lu, Xuancheng Xu, Feihong He, Jaehong Yoon, Taixi Lu, Fuxiao Liu, Gedas Bertasius, Mohit Bansal, Huaxiu Yao, and Furong Huang. Mementos: A comprehensive benchmark for multimodal large language model reasoning over image sequences. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, *ACL 2024, Bangkok, Thailand, August 11-16, 2024*, pages 416–442. Association for Computational Linguistics, 2024. 1
- [43] Jiayang Wu, Wensheng Gan, Zefeng Chen, Shicheng Wan, and Philip S. Yu. Multimodal large language models: A survey. In *IEEE Big Data*, pages 2247–2256. IEEE, 2023. 1

- [44] Liang Xu, Hang Xue, Lei Zhu, and Kangkang Zhao. Superclue-math6: Graded multi-step math reasoning benchmark for llms in chinese. *CoRR*, abs/2401.11819, 2024. 2
- [45] Kaining Ying, Fanqing Meng, Jin Wang, Zhiqian Li, Han Lin, Yue Yang, Hao Zhang, Wenbo Zhang, Yuqi Lin, Shuo Liu, Jiayi Lei, Quanfeng Lu, Runjian Chen, Peng Xu, Renrui Zhang, Haozhe Zhang, Peng Gao, Yali Wang, Yu Qiao, Ping Luo, Kaipeng Zhang, and Wenqi Shao. Mmt-bench: A comprehensive multimodal benchmark for evaluating large vision-language models towards multitask AGI. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. 1
- [46] Xiang Yue, Yuansheng Ni, Tianyu Zheng, Kai Zhang, Ruochi 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, Huan Sun, Yu Su, and Wenhua Chen. MMMU: A massive multi-discipline multimodal understanding and reasoning benchmark for expert AGI. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 9556–9567, 2024. 1, 3, 8
- [47] Chao Zhang, Qingfeng Xu, Yongrui Yu, Guanghui Zhou, Keyan Zeng, Fengtian Chang, and Kai Ding. A survey on potentials, pathways and challenges of large language models in new-generation intelligent manufacturing. *Robotics Comput. Integr. Manuf.*, 92:102883, 2025. 1
- [48] Jiaxin Zhang, Zhongzhi Li, Ming-Liang Zhang, Fei Yin, Cheng-Lin Liu, and Yashar Moshfeghi. Geoeval: Benchmark for evaluating llms and multi-modal models on geometry problem-solving. In *ACL (Findings)*, pages 1258–1276. Association for Computational Linguistics, 2024. 1, 3
- [49] Renrui Zhang, Jiaming Han, Chris Liu, Aojun Zhou, Pan Lu, Yu Qiao, Hongsheng Li, and Peng Gao. Llama-adapter: Efficient fine-tuning of large language models with zero-initialized attention. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*, 2024. 3
- [50] Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Peng Gao, and Hongsheng Li. Mathverse: Does your multi-modal LLM truly see the diagrams in visual math problems? *CoRR*, abs/2403.14624, 2024. 1, 3
- [51] Renrui Zhang, Xinyu Wei, Dongzhi Jiang, Yichi Zhang, Ziyu Guo, Chengzhuo Tong, Jiaming Liu, Aojun Zhou, Bin Wei, Shanghang Zhang, Peng Gao, and Hongsheng Li. MAVIS: mathematical visual instruction tuning. *CoRR*, abs/2407.08739, 2024. 3
- [52] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023. 1
- [53] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*, 2024. 3

MV-MATH: Evaluating Multimodal Math Reasoning in Multi-Visual Contexts

Supplementary Material

Table of Contents in Appendix

A More Details about MV-MATH	2
A.1. Question Distribution	2
A.2. Image Distribution	2
A.3. Division of Difficulty Levels	2
A.4. Image by Subjects	3
A.5. Introduction of Subjects	3
B Evaluation Details	15
B.1. Prompt for Image Caption Generation	15
B.2. Prompt for Response Generation	15
B.3. Prompt for Answer Evaluation	15
B.4. Model Details	15
C Main Results Across 3 Question Types	21
D Results of CoT, 2-shot on 3 Question Types	21
E Image Relevance	21
F. Data Collection and Annotation	21
F.1. Data Collection	21
F.2. Annotation	21
G Comparison with Existing Benchmarks	22
H Case Study	22

A. More Details about MV-MATH

In this chapter, we will introduce MV-MATH in more detail.

A.1. Question Distribution

All questions in MV-MATH are presented in English. As shown in Table 1, the longest question in MV-MATH spans 383 words, while the shortest contains 14 words, with an average length of 80.17 words. Figure 7 further illustrates the distribution of text lengths, highlighting the diversity of MV-MATH.

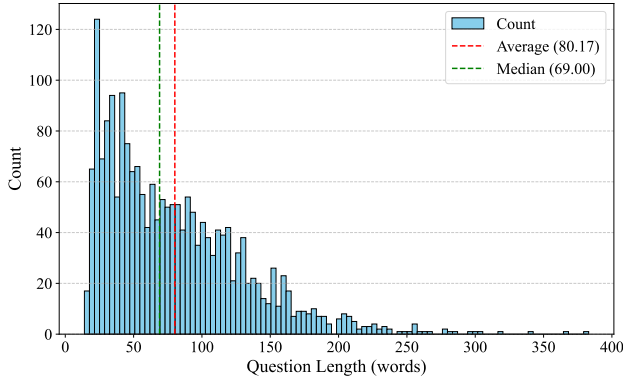


Figure 7. The distribution of the number of words per question in MV-MATH.

A.2. Image Distribution

Figure 8 illustrates the distribution of the number of images associated with each question in the MV-MATH dataset, highlighting its multimodal diversity. The dataset features an average of 3.02 images per question, with a median of 3, demonstrating a balanced and realistic allocation of visual resources. Most questions include 2, 3 or 4 images, reflecting the dataset’s emphasis on providing sufficient visual context for reasoning tasks. The presence of questions with 5 or more images showcases the dataset’s capability to handle complex, multi-visual scenarios. This diversity ensures MV-MATH is suitable for evaluating the ability of models to integrate and reason across multiple interconnected visual elements, further enhancing its value as a benchmark for multimodal mathematical reasoning.

A.3. Division of Difficulty Levels

Since MV-MATH is a meticulously annotated dataset containing both answers and an analysis field, it offers a unique advantage for difficulty classification based on the lengths of the question and analysis. The difficulty classification is conducted in two steps.

Step 1: A weighted average is calculated using the lengths of the *question* and *analysis* fields, assigning a weight of 0.4 to the *question* length and 0.6 to the *analysis*




Figure 8. The distribution of the number of images per question in MV-MATH.

length. This weighting reflects the assumption that the solution process (captured in the *analysis* field) is more indicative of a question’s difficulty. The weighted length distribution is shown in Figure 8. Using the K-means clustering algorithm, we cluster the weighted lengths into three difficulty levels: easy (0–150), medium (150–500), and hard (> 500).

Step 2: Manual verification is performed on the clustering results from Step 1 by two graduate students. This process adjusts the difficulty classification to account for cases where the weighted length alone may not fully reflect complexity. For instance, a question with a short weighted length but involving intricate formula derivation is reclassified as harder, while a question with a long weighted length but straightforward reasoning is reclassified as easier.

The detailed distribution of weighted lengths can be seen in Figure 9. After the two-step process described above, we classified the questions into 542 as *easy*, 964 as *medium*, and 503 as *hard*.




Figure 9. Weighted Length Distribution of Questions and Analysis for Difficulty Classification in MV-MATH

A.4. Image by Subjects

This section presents examples of images from various categories in the MV-MATH dataset. Figures 10 to 20 correspond to images under the categories of *Algebra*, *Arithmetic*, *Combinatorial Geometry*, *Analytic Geometry*, *Combinatorics*, *Descriptive Geometry*, *Logic*, *Metric Geometry*, *Solid Geometry*, *Statistics*, and *Transformation Geometry*, respectively.

The diversity of images across categories reflects the wide range of mathematical concepts covered in MV-MATH. For instance, images in the *Statistics* category primarily include charts, graphs, and tables, providing rich visual representations of statistical data. In contrast, images under *Arithmetic*, often targeted at younger students, feature elements with real-world objects or playful designs to aid comprehension. Similarly, the *Combinatorics* and *Combinatorial Geometry* categories include diagrams with intricate arrangements, requiring detailed reasoning. On the other hand, images in *Analytic Geometry* and *Metric Geometry* are more abstract and geometric, often involving coordinate systems, vectors, or precise measurements.

This variety of visual styles within and across categories highlights the versatility of MV-MATH in evaluating models' ability to interpret and reason over diverse visual contexts, making it a comprehensive benchmark for multi-visual mathematical reasoning tasks.

A.5. Introduction of Subjects

Analytic Geometry. Analytic Geometry integrates algebraic techniques with geometry through the use of a coordinate system. It provides a systematic method to describe geometric shapes using equations and to interpret these equations visually. This branch of mathematics enables a detailed study of geometric properties, such as distance, angles, and tangents.

Algebra. Algebra is a fundamental area of mathematics that explores the use of symbols to represent numbers and quantities in formulas and equations. It encompasses a wide range of topics, from solving linear equations to studying abstract structures such as groups, rings, and fields, playing a crucial role in generalizing mathematical principles and solving problems where specific values are unknown.

Metric Geometry. Metric Geometry focuses on the study of geometric figures based on distances and angles. It examines properties that remain invariant under transformations such as rotations, translations, and reflections. Metric Geometry serves as a foundation for applications in computer graphics, engineering, and physics.

Combinatorics. Combinatorics is the branch of mathematics concerned with counting, arrangement, and combination of objects. It investigates problems related to discrete struc-

tures, such as permutations, combinations, and graph theory. Combinatorics is widely used in areas like cryptography, algorithm design, and network theory.

Transformation Geometry. Transformation Geometry studies geometric transformations, such as translations, rotations, reflections, and dilations, to understand how shapes change while preserving certain properties. It provides insights into symmetry, congruence and has significant applications in computer vision.

Logic. Logic in mathematics involves the study of principles of reasoning, including the formulation and analysis of valid arguments. It lays the groundwork for proof techniques and the development of formal systems. Mathematical logic is essential for understanding the structure of mathematical theories.

Solid Geometry. Solid Geometry is the study of three-dimensional figures such as spheres, cubes, cones, and cylinders. It explores their properties, measurements, and spatial relationships, often involving volume, surface area, and intersections.

Arithmetic. Arithmetic is the oldest and most fundamental branch of mathematics, dealing with numbers and basic operations such as addition, subtraction, multiplication, and division. It serves as the foundation for other mathematical disciplines and is essential for everyday problem-solving and quantitative reasoning.

Combinatorial Geometry. Combinatorial Geometry combines concepts from combinatorics and geometry to study the arrangement and interaction of geometric objects. It addresses problems involving configurations of points, lines, and planes, often focusing on optimization and enumeration. This field has applications in graph theory, computational geometry, and optimization.

Descriptive Geometry. Descriptive Geometry is a method for representing three-dimensional objects in two dimensions using projections. It involves techniques to visualize and solve spatial problems through precise drawings. This field is widely used in engineering and architecture for designing and visualizing complex structures.

Statistics. Statistics is the study of data collection, analysis, interpretation, and presentation. It involves mathematical techniques to summarize and infer conclusions from data, often focusing on patterns, variability, and uncertainty. Statistics is indispensable in research, economics, and decision-making processes across various domains.




Figure 10. Some images from Analytic Algebra.

Arithmetic








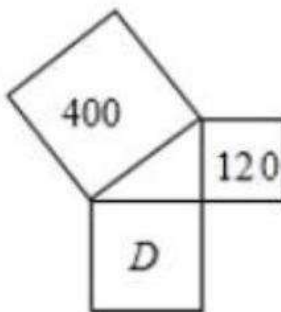

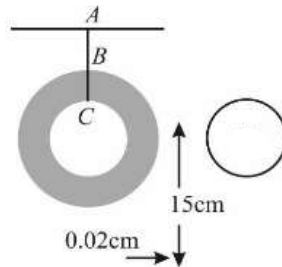


Figure 11. Some images from Arithmetic.

Combinatorial Geometry

Figure 12. Some images from Combinatorial Geometry.

Analytic Geometry

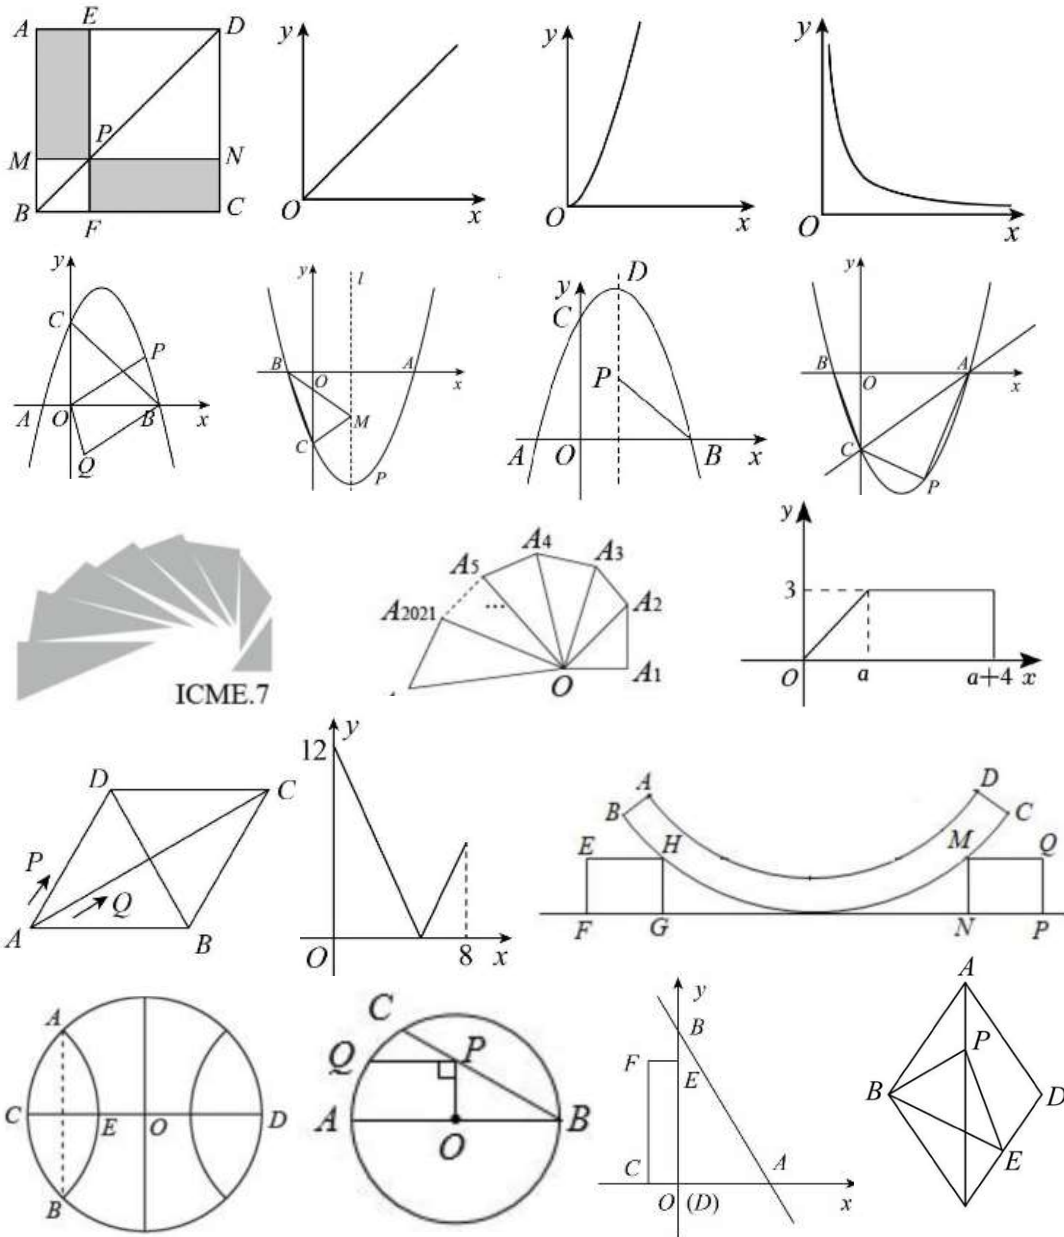


Figure 13. Some images from Analytic Geometry.

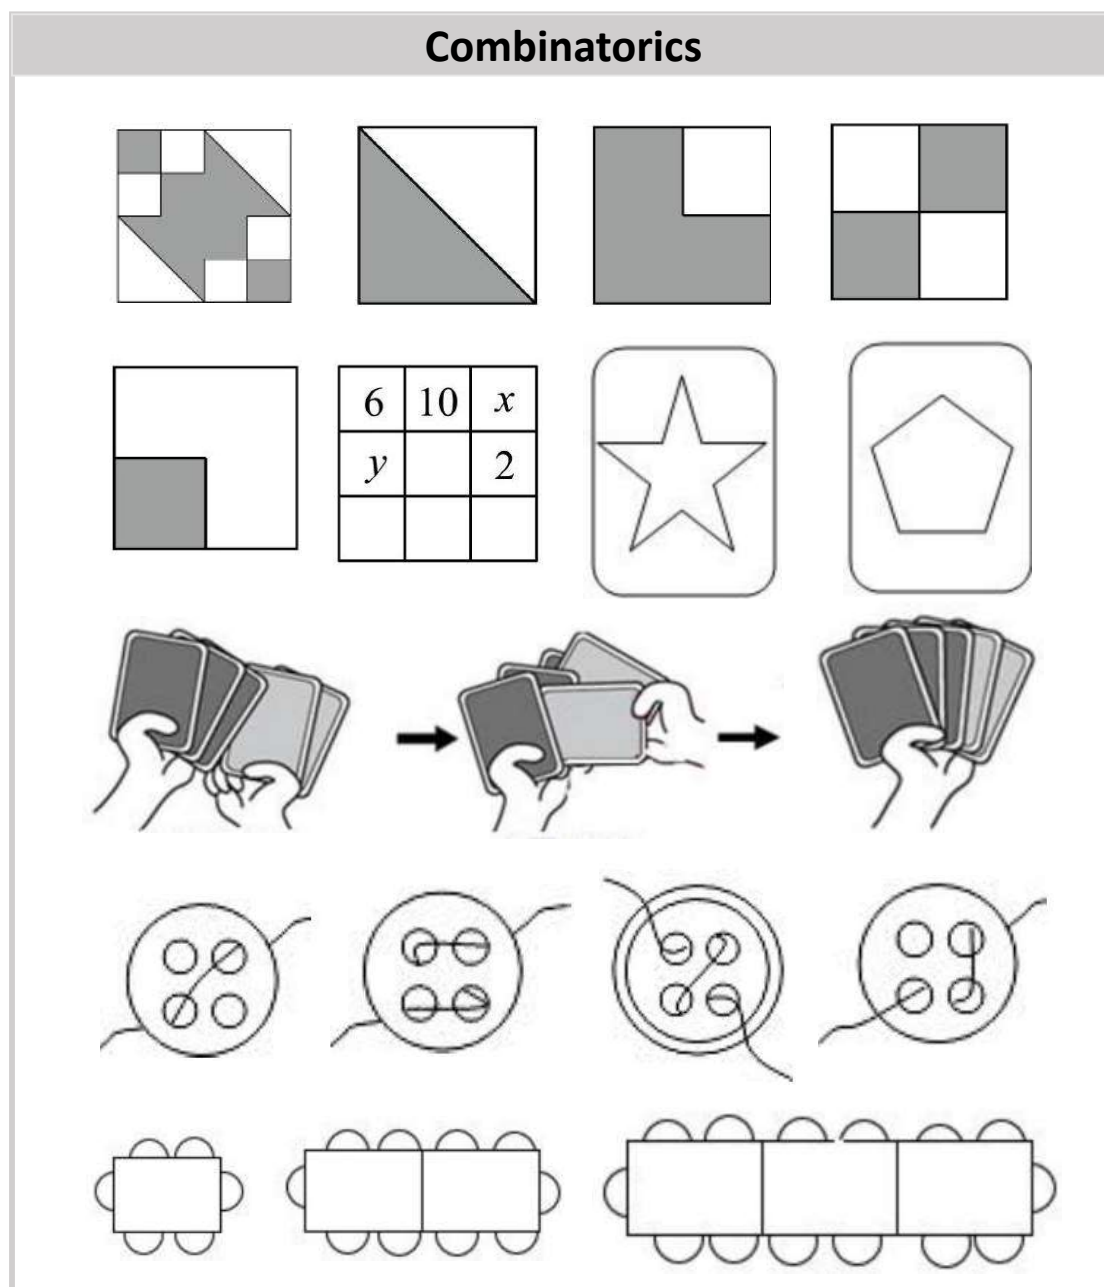


Figure 14. Some images from Combinatorics.

Descriptive Geometry




Figure 15. Some images from Descriptive Geometry.

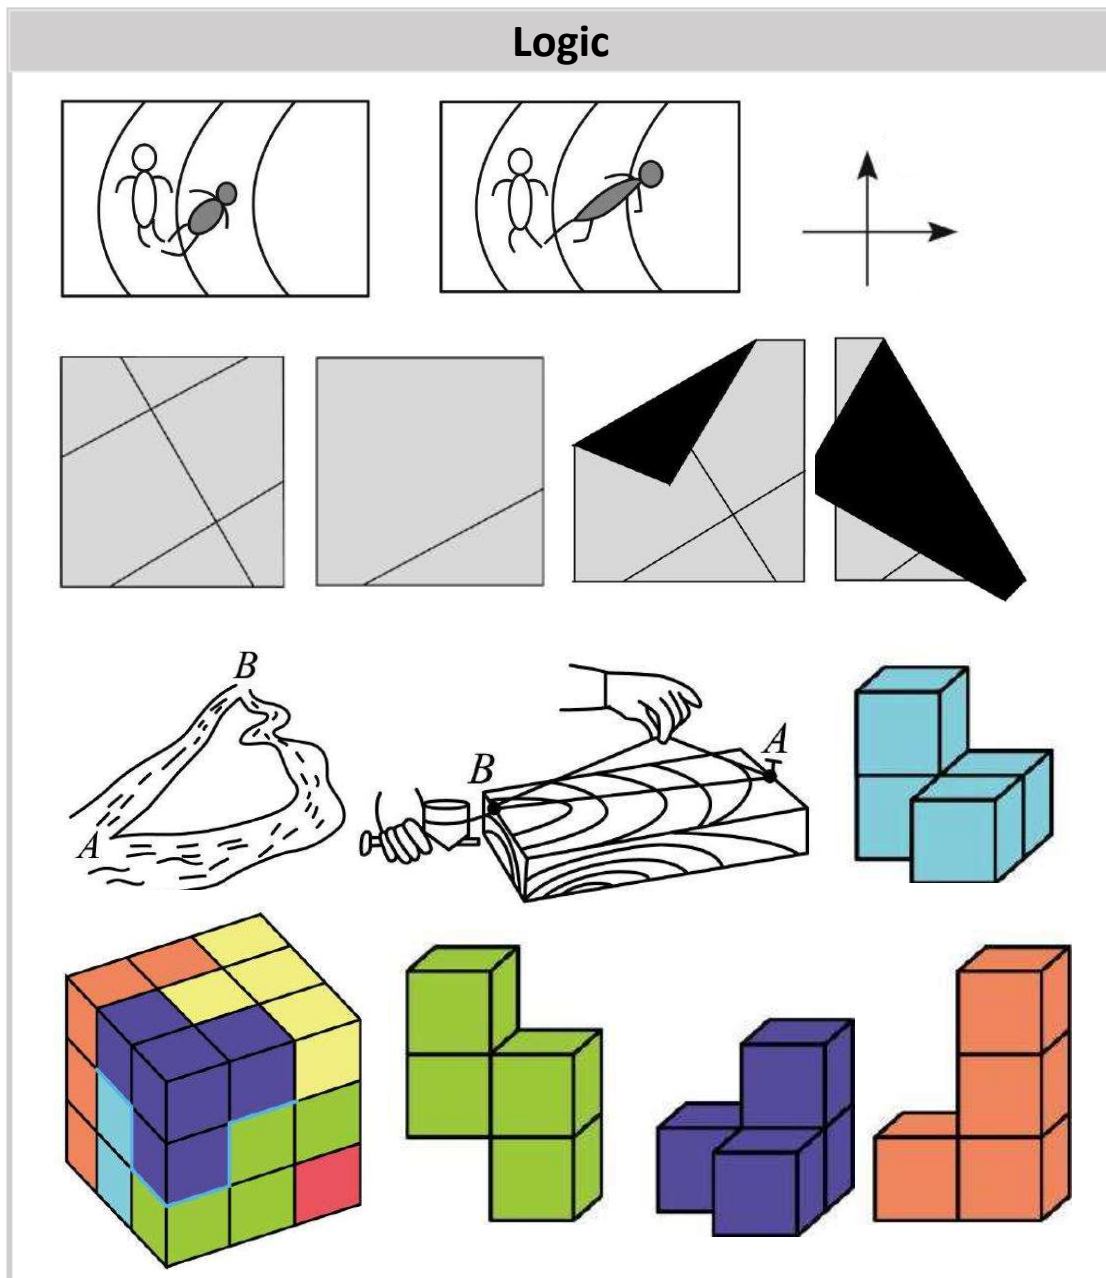


Figure 16. Some images from Logic.

Metric Geometry

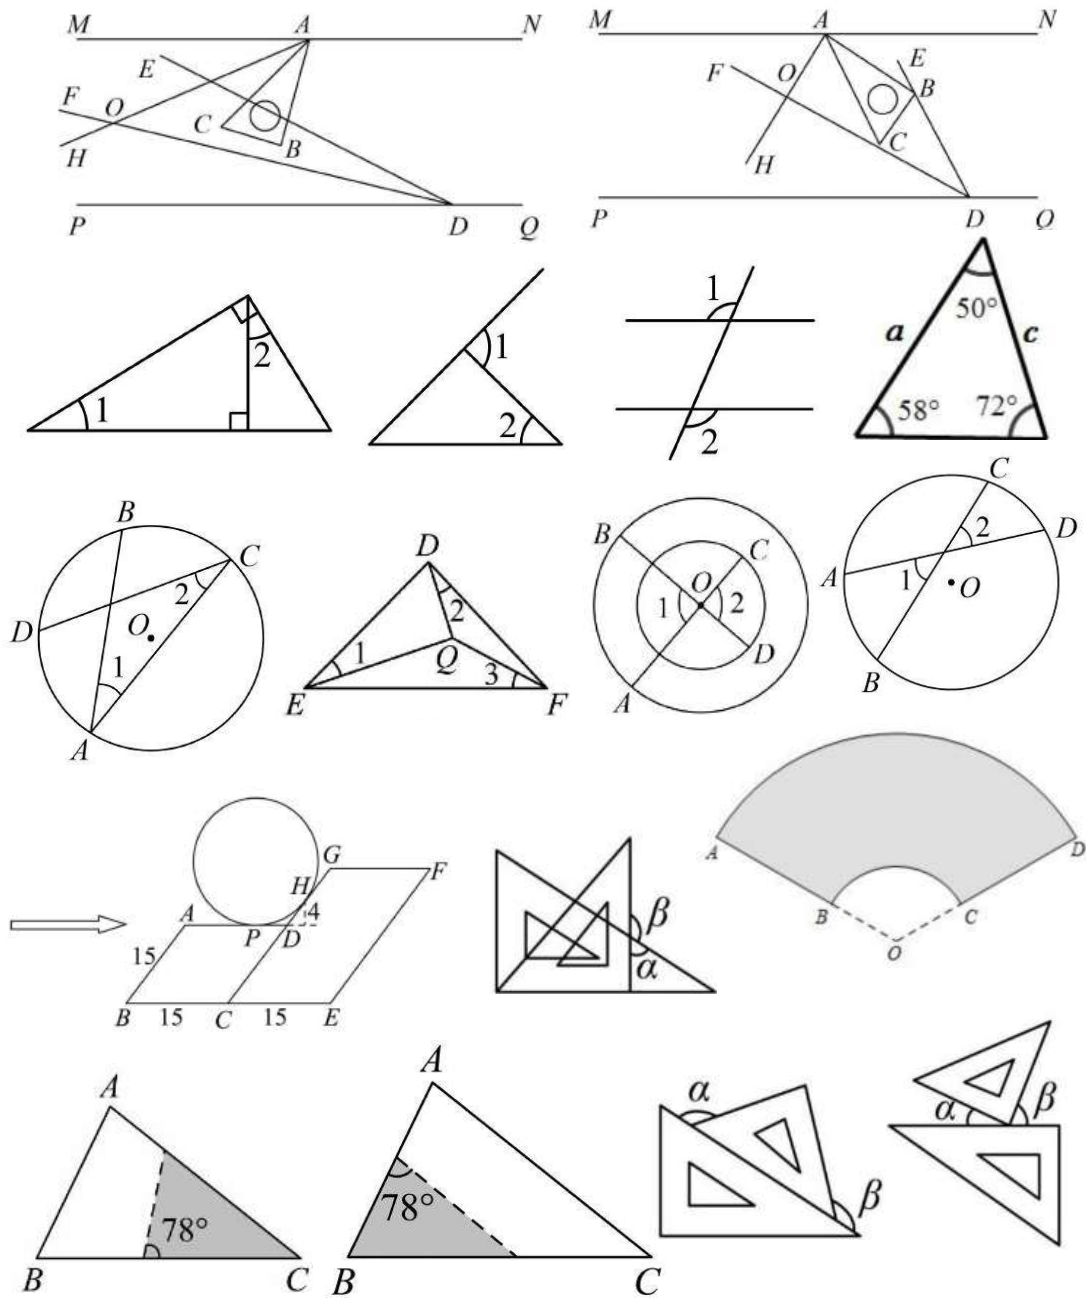


Figure 17. Some images from Metric Geometry.

Solid Geometry

The image displays a collection of 20 geometric figures and nets arranged in a 5x4 grid, intended for a matching exercise. The figures include:

- Row 1: Four cubes. The first has a black circle on top and a shaded triangle on the front. The second has a black circle on top and a shaded triangle on the front. The third has a white circle on top and a shaded triangle on the front. The fourth has a white circle on top and a black circle on the front.
- Row 2: Four groups of three objects. The first group consists of a cylinder and two cubes. The second group consists of a cone and two pyramids. The third group consists of a hexagonal prism and two rectangular prisms. The fourth group consists of a frustum and two rectangular prisms.
- Row 3: Four 3D shapes made of cubes. The first is a 2x2x1 shape. The second is a 2x2x2 shape. The third is a 2x2x1 shape. The fourth is a 2x2x1 shape.
- Row 4: Four basic solids. The first is a triangular prism. The second is a sphere. The third is a cube. The fourth is a square pyramid.
- Row 5: Four nets. The first is a cube net with shaded faces. The second is a cube net with shaded faces and circles. The third is a cross net with shaded faces and circles. The fourth is a cross net with shaded faces and circles.
- Row 6: Four items. The first is a cube net. The second is a cylinder and a rectangular prism. The third is a cylinder with a surface area of 18.84. The fourth is a cylinder with a surface area of 24.

Figure 18. Some images from Solid Geometry.

Statistics

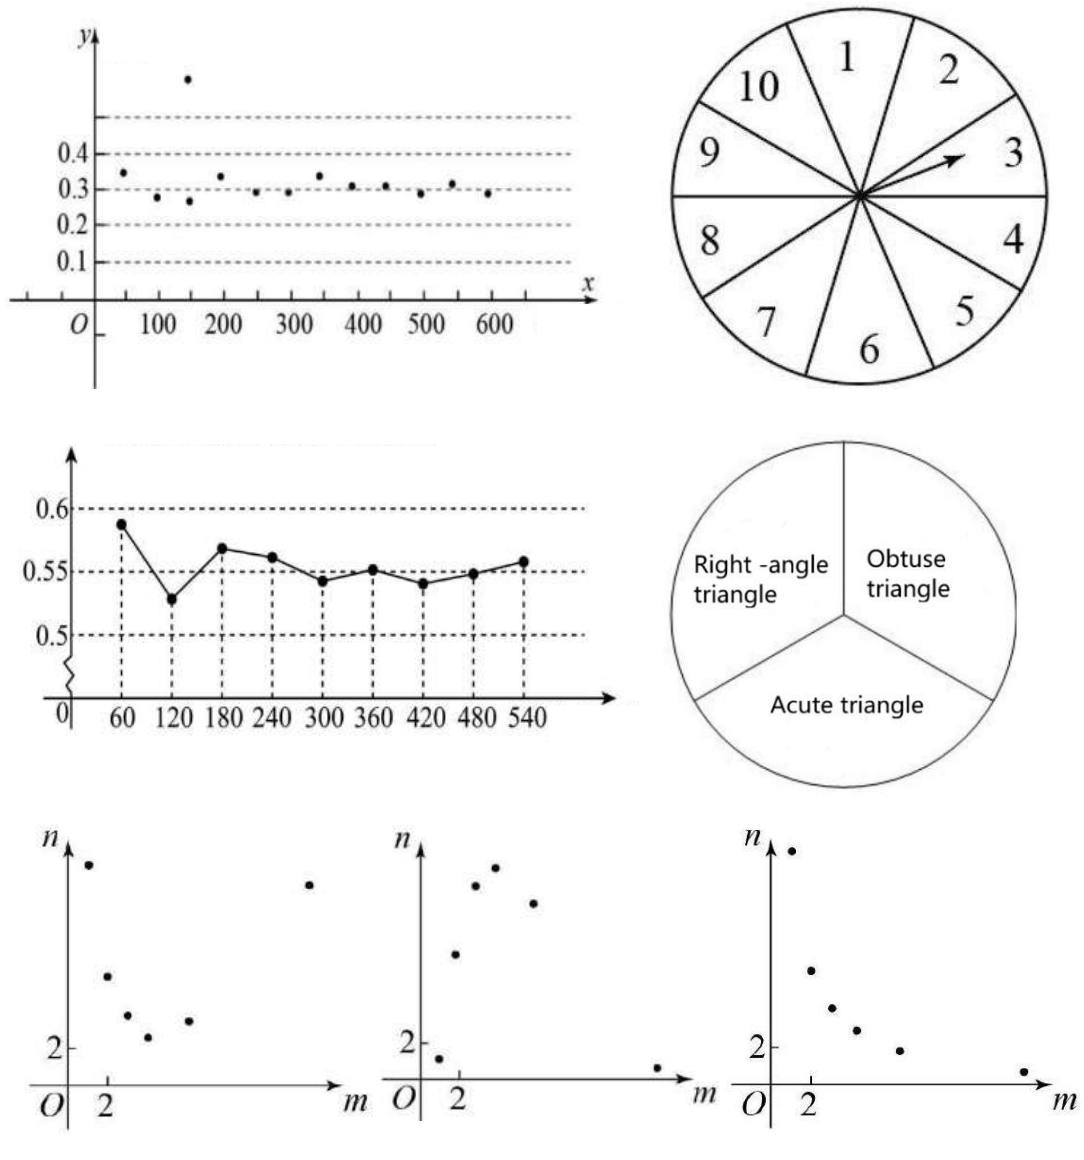


Figure 19. Some images from Statistics.

Transformation Geometry

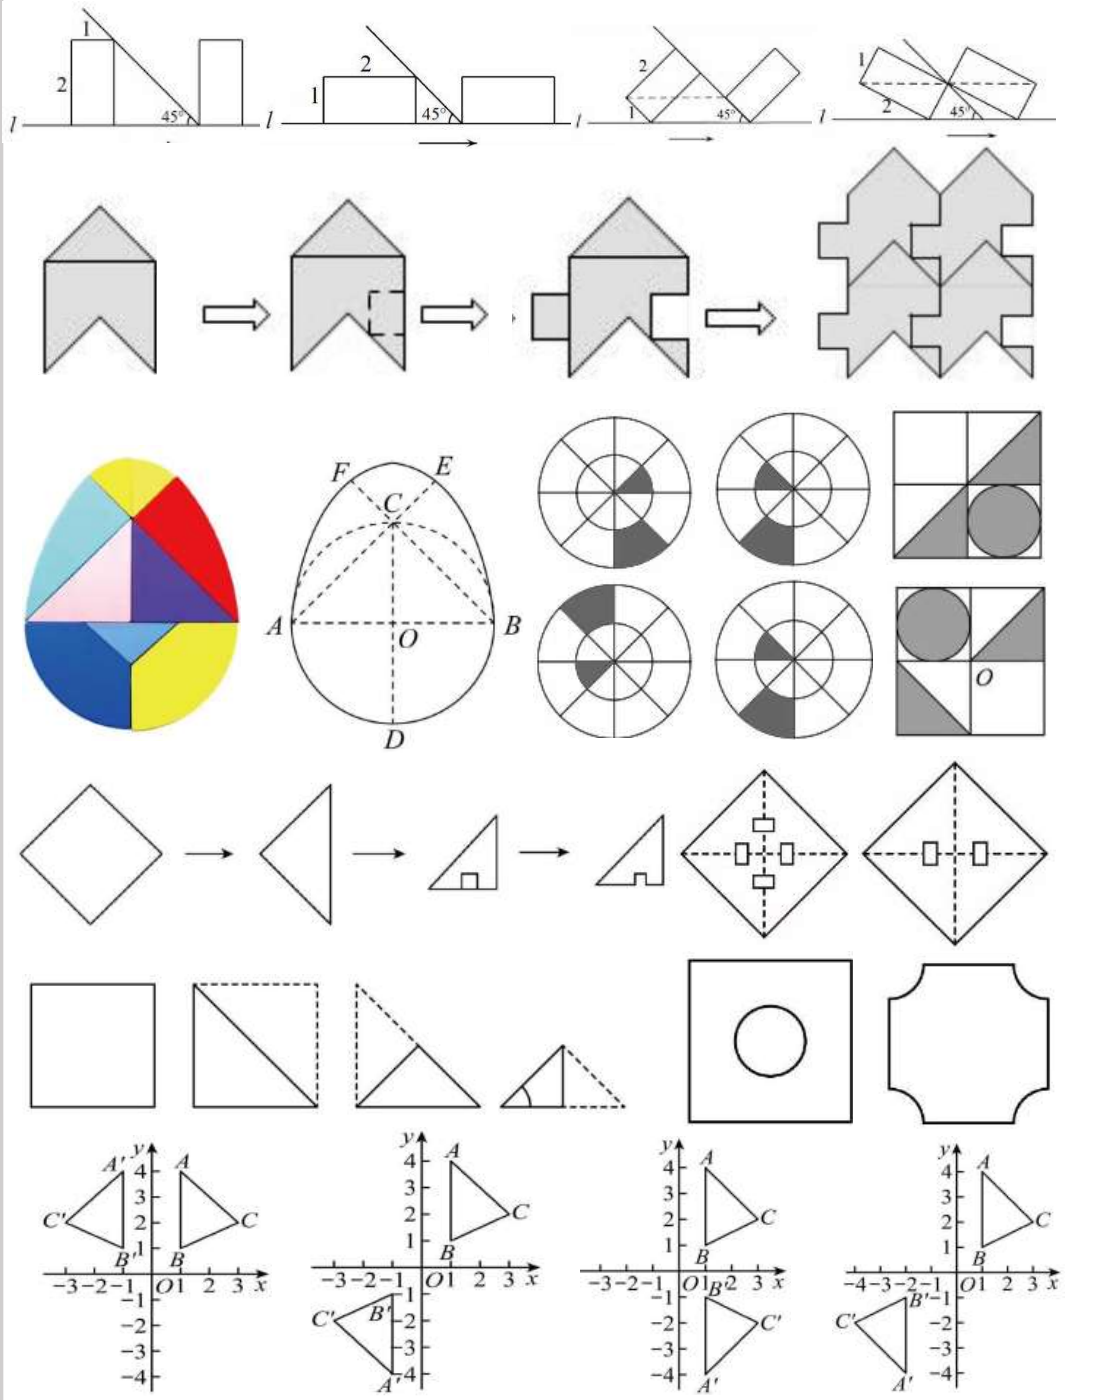


Figure 20. Some images from Transformation Geometry.

B. Evaluation Details

For open-source models, all experiments for 72B-scale models are conducted on H-800 GPUs, while the remaining models are conducted on A6000 GPUs.

B.1. Prompt for Image Caption Generation

We use Claude-3.5-Sonnet to generate image caption, the prompt is as follows: *You are an expert in image description, here is a diagram of a math problem, you need to convert the image into text description in detail so that your description can be used to replace the diagram.*

B.2. Prompt for Response Generation

To ensure the model provides accurate responses, we design distinct CoT and 2-shot prompts tailored for multiple-choice, single-step, and multi-step free-form questions. The original prompt directly instructs the model to generate the final answer without intermediate reasoning. Detailed information can be found in Table Tables 5 to 7.

B.3. Prompt for Answer Evaluation

Our evaluation is conducted using the Deepseek API. For the evaluation of multiple-choice, single-step, and multi-step free-form questions, different prompts are designed to ensure accuracy in answer extraction and assessment. We first use the Deepseek API to extract the model’s answers and then it compares the extracted answers with the ground truth to determine the correctness of the answers. The specific prompts are shown in Table 8 below.

B.4. Model Details

All experiments are conducted using models configured with a temperature of 0.2 and a max_new_token limit of 2048 for text generation. Comprehensive details regarding the models utilized in the evaluation are presented in Table 9.

Original
You are an assistant for solving math problems. Your input consists of a math question and images, give your answer directly, without any intermediate steps.
CoT (Chain of Thought)
You are an assistant for solving math problems. Your input consists of a math question and images. Your task is to output the solution steps and the answer. The output format should be a step-by-step approach. Each question is multiple choice with one correct answer. Your final answer must be one of A, B, C, or D, and it should be placed within {}. For example: {A}, {B}, {C}, or {D}.
CoT with 2-shot
<p>Example 1:</p> <p>Question: If a triangle has two sides of length 3 and 4, what is the length of the hypotenuse? A.10 B.8 C.5 D.4</p> <p>Answer:</p> <p>Step 1 (Mathematical theorem used: Pythagorean theorem): The Pythagorean theorem states that in a right triangle, the square of the hypotenuse is equal to the sum of the squares of the other two sides. The formula is: $c^2 = a^2 + b^2$, where a and b are the legs, and c is the hypotenuse.</p> <p>Step 2 (Substitute the known values): Given $a = 3$ and $b = 4$. Substituting these values into the formula: $c^2 = 3^2 + 4^2 = 9 + 16 = 25$</p> <p>Step 3 (Calculate the hypotenuse): Taking the square root gives: $c = \sqrt{25} = 5$</p> <p>Answer: {C}</p> <p>Example 2:</p> <p>Question: In the right triangle ABC, AB is perpendicular to BC. It is known that AC=5 and AB=4. Find the area of the right triangle. A.20 B.10 C.5 D.6</p> <p>Answer:</p> <p>Step 1 (Mathematical theorem used: Pythagorean theorem): We first use the Pythagorean theorem to find the length of BC. The formula is: $AC^2 = AB^2 + BC^2$, where AC is the hypotenuse, and AB and BC are the legs.</p> <p>Step 2 (Substitute the known values): Given $AC = 5$ and $AB = 4$. Substituting these values: $5^2 = 4^2 + BC^2 \implies 25 = 16 + BC^2$</p> <p>Step 3 (Solve for BC): $BC^2 = 25 - 16 = 9 \implies BC = \sqrt{9} = 3$</p> <p>Step 4 (Calculate the area): The area of the right triangle is given by $\frac{1}{2} \times AB \times BC$. Substituting the known values: Area = $\frac{1}{2} \times 4 \times 3 = 6$</p> <p>Answer: {D}</p> <p>Your final answer must be one of A, B, C, or D, and it should be placed within {}</p>

Table 5. The prompts used for choice questions in the evaluation for response generation.

Original Prompt
You are an assistant for solving math problems. Your input consists of a math question and images. Give your answer directly, without any intermediate steps.
CoT (Chain of Thought)
You are an assistant for solving math problems. Your input consists of a math question and images. Your task is to output the solution steps and the answer. The output format should be a step-by-step approach.
CoT with 2-shot
<p>Example 1:</p> <p>Question: If a triangle has two sides of length 3 and 4, what is the length of the hypotenuse?</p> <p>Answer:</p> <p>Step 1: (Mathematical theorem used: Pythagorean theorem): The Pythagorean theorem states that in a right triangle, the square of the hypotenuse is equal to the sum of the squares of the other two sides. The formula is: $c^2 = a^2 + b^2$, where a and b are the legs, and c is the hypotenuse.</p> <p>Step 2: (Substitute the known values): Given $a = 3$ and $b = 4$. Substituting these values into the formula: $c^2 = 3^2 + 4^2 = 9 + 16 = 25$.</p> <p>Step 3: (Calculate the hypotenuse): Taking the square root gives: $c = \sqrt{25} = 5$.</p> <p>Answer: 5</p> <p>Example 2:</p> <p>Question: In the right triangle ABC, AB is perpendicular to BC. It is known that $AC = 5$ and $AB = 4$. Find the area of the right triangle.</p> <p>Answer:</p> <p>Step 1: (Mathematical theorem used: Pythagorean theorem): We first use the Pythagorean theorem to find the length of BC. The formula is: $AC^2 = AB^2 + BC^2$, where AC is the hypotenuse, and AB and BC are the legs.</p> <p>Step 2: (Substitute the known values): Given $AC = 5$ and $AB = 4$. Substituting these values: $5^2 = 4^2 + BC^2 \implies 25 = 16 + BC^2$.</p> <p>Step 3: (Solve for BC): $BC^2 = 25 - 16 = 9 \implies BC = \sqrt{9} = 3$.</p> <p>Step 4: (Calculate the area): The area of the right triangle is given by $\frac{1}{2} \times AB \times BC$. Substituting the known values: Area = $\frac{1}{2} \times 4 \times 3 = 6$.</p> <p>Answer: 6</p> <p>Please reason step by step. Each step is placed on a new line, using the following format: Step X (Mathematical theorem/basis used): Detailed solution steps. Answer: {}</p>

Table 6. Prompts used for single-step free-form questions in the evaluation for response generation.

Original Prompt
You are an assistant for solving math problems. Your input consists of a math question and images. Each problem is a multi-step problem. Give your answer directly, without any intermediate steps.
CoT (Chain of Thought)
You are a math problem-solving assistant. Your input is a math problem and a picture of the problem. Each problem is a multi-step problem. Your task is to output the solution ideas and answers for each step. The output format is step-by-step.
CoT with 2-shot Examples
<p>Example 1:</p> <p>Question: If a triangle has two sides of length 3 and 4, (1) what is the length of the hypotenuse? (2) what is the area of this triangle?</p> <p>Answer:</p> <p>(1) Step 1: (Mathematical theorem used: Pythagorean theorem): The Pythagorean theorem states that in a right triangle, the square of the hypotenuse is equal to the sum of the squares of the other two sides. The formula is: $c^2 = a^2 + b^2$, where a and b are the legs, and c is the hypotenuse.</p> <p>Step 2: (Substitute the known values): Given $a = 3$ and $b = 4$. Substituting these values into the formula: $c^2 = 3^2 + 4^2 = 9 + 16 = 25$.</p> <p>Step 3: (Calculate the hypotenuse): Taking the square root gives: $c = \sqrt{25} = 5$.</p> <p>So the length of the hypotenuse is 5.</p> <p>(2) Step 1: The area of a right triangle is half the product of its two sides.</p> <p>Step 2: So the area of this triangle is $3 \times 4/2 = 6$.</p> <p>So the area of this triangle is 6.</p> <p>Example 2:</p> <p>Question: In the right triangle ABC, AB is perpendicular to BC. It is known that $AC = 5$ and $AB = 4$. (1) Find the area of the right triangle. (2) What is the height of the hypotenuse of this right triangle?</p> <p>Answer:</p> <p>(1) Step 1: (Mathematical theorem used: Pythagorean theorem): We first use the Pythagorean theorem to find the length of BC. The formula is: $AC^2 = AB^2 + BC^2$, where AC is the hypotenuse, and AB and BC are the legs.</p> <p>Step 2: (Substitute the known values): Given $AC = 5$ and $AB = 4$. Substituting these values: $5^2 = 4^2 + BC^2 \implies 25 = 16 + BC^2$.</p> <p>Step 3: (Solve for BC): $BC^2 = 25 - 16 = 9 \implies BC = \sqrt{9} = 3$.</p> <p>Step 4: (Calculate the area): The area of the right triangle is given by $\frac{1}{2} \times AB \times BC$. Substituting the known values: Area = $\frac{1}{2} \times 4 \times 3 = 6$.</p> <p>So the area of the right triangle is 6.</p> <p>(2) Step 1: According to the equal area method, the area of a right triangle is equal to half the product of the two right-angled sides, which is also equal to half the product of the hypotenuse and the corresponding height.</p> <p>Step 2: According to the above principle and the conclusion of the first step, we can get $AB \times BC/2 = AC \times h/2$.</p> <p>Step 3: Substituting the values, we get $h = 3 \times 4/5 = 2.4$.</p> <p>So the height of the hypotenuse of this right triangle is 2.4.</p> <p>Please reason step by step. Each step is placed on a new line, using the following format: Step X (Mathematical theorem/basis used): Detailed solution steps. Answer: {}</p>

Table 7. Prompts used for multi-step free-form questions in the evaluation for response generation.

Multiple-Choice Prompt
<p>You are an assistant for evaluating math problems. Your task is to extract the model’s answer to the given multiple-choice question and compare it with the ground truth.</p> <p>Steps:</p> <ol style="list-style-type: none"> 1. Extract the model’s answer. The answer must be one of A, B, C, or D. 2. Compare the extracted answer with the ground truth. 3. Indicate whether the model’s answer is correct or incorrect. <p>Output format:</p> <ul style="list-style-type: none"> - Extracted Answer: {A}, {B}, {C}, or {D}. - Correctness: [true/false].
Single-Step Free-Form Prompt
<p>You are an assistant for evaluating math problems. Your task is to extract the model’s answer to the given single-step free-form question and compare it with the ground truth.</p> <p>Steps:</p> <ol style="list-style-type: none"> 1. Extract the model’s final answer. 2. Compare the extracted answer with the ground truth. 3. Indicate whether the model’s answer is correct or incorrect. <p>Output format:</p> <ul style="list-style-type: none"> - Extracted Answer: [Final Answer]. - Correctness: [true/false].
Multi-Step Free-Form Prompt
<p>You are an assistant for evaluating math problems. Your task is to extract the model’s answers to each sub-question of a multi-step free-form problem and compare them with the ground truth.</p> <p>Steps:</p> <ol style="list-style-type: none"> 1. Extract the final answers for each sub-question. 2. Compare the extracted answers with the corresponding ground truth. 3. Indicate whether each answer is correct or incorrect. <p>Output format:</p> <ul style="list-style-type: none"> - Sub-Question 1: Extracted Answer: [Answer]. Correctness: [true/false]. - Sub-Question 2: Extracted Answer: [Answer]. Correctness: [true/false].

Table 8. Prompts used for evaluating different types of math problems with the Deepseek API.

Model	Source	URL
Deepseek-chat	Deepseek-chat	https://api-docs.deepseek.com/
Math-LLaVA-13B	local checkpoint	https://huggingface.co/Zhiqiang007/Math-LLaVA
LLaVA-v1.5-13B	local checkpoint	https://huggingface.co/liuhaotian/llava-v1.5-13b
LLaVA-v1.5-7B	local checkpoint	https://huggingface.co/liuhaotian/llava-v1.5-7b
VILA-13B	local checkpoint	https://huggingface.co/Efficient-Large-Model/VILA-13b
InternLM-XComposer2.5-VL-7B	local checkpoint	https://huggingface.co/internlm/internlm-xcomposer2d5-7b
InternVL-Chat-8B	local checkpoint	https://huggingface.co/OpenGVLab/InternVL2-8B
Llama-3.2-Vision-Instruct-11B	local checkpoint	https://huggingface.co/meta-llama/Llama-3.2-11B-Vision-Instruct
Deepseek-VL-7B	local checkpoint	https://huggingface.co/deepseek-ai/deepseek-vl-7b-chat
LLaVA-NeXT-Interleave-7B	local checkpoint	https://huggingface.co/lmms-lab/llava-next-interleave-qwen-7b
Mantis-Idefics2-8B	local checkpoint	https://huggingface.co/TIGER-Lab/Mantis-8B-Idefics2
Mantis-siglip-8B	local checkpoint	https://huggingface.co/TIGER-Lab/Mantis-8B-siglip-llama3
Qwen2VL-Instruct-7B	local checkpoint	https://huggingface.co/Qwen/Qwen2-VL-7B-Instruct
LLaVA-OneVision-SI-7B	local checkpoint	https://huggingface.co/lmms-lab/llava-onevision-qwen2-7b-si
LLaVA-OneVision-SFT-7B	local checkpoint	https://huggingface.co/lmms-lab/llava-onevision-qwen2-7b-ov
LLaVA-OneVision-Chat-7B	local checkpoint	https://huggingface.co/lmms-lab/llava-onevision-qwen2-7b-ov-chat
LLaVA-OneVision-SI-72B	local checkpoint	https://huggingface.co/lmms-lab/llava-onevision-qwen2-72b-si
LLaVA-OneVision-SFT-72B	local checkpoint	https://huggingface.co/lmms-lab/llava-onevision-qwen2-72b-ov-sft
LLaVA-OneVision-Chat-72B	local checkpoint	https://huggingface.co/lmms-lab/llava-onevision-qwen2-72b-ov-chat
InternLM-XComposer2-VL	local checkpoint	https://huggingface.co/internlm/internlm-xcomposer2-vl-7b
Qwen-VL-Plus	qwen-vl-plus	https://help.aliyun.com/zh/dashscope/developer-reference/vl-plus-quick-start
GPT-4V	gpt-4-vision-2023-05-15	https://platform.openai.com/
Qwen-VL-Max	qwen-vl-max	https://help.aliyun.com/zh/dashscope/developer-reference/vl-plus-quick-start
Gemini-1.5-Pro	gemini-1.5-Pro-2023-05-15	https://ai.google.dev/
GPT-4o	gpt-4o-2024-05-14	https://platform.openai.com/
Claude-3.5-Sonnet	claude-3.5-sonnet-2024-05-24	https://www.anthropic.com/news/claude-3-5-sonnet

Table 9. The source of the models used in the evaluation.

C. Main Results Across 3 Question Types

In this section, we present a detailed analysis of the experimental results across three question types: multiple-choice, single-step free-form, and multi-step free-form. Each subsection highlights the model’s performance on these distinct question formats, providing insights into the effectiveness and limitations of the evaluated methods. The results are compared and discussed to showcase how different models handle varying levels of complexity and reasoning requirements.

The result of *choice*, *single-step* and *multi-step* are shown in Tables 10 to 13. We observe that there is a performance gap between multiple-choice questions and free-form question types. Claude-3.5-Sonnet perform the best on multiple-choice questions with a score of 44.0%, followed by GPT-4o and LLaVA-OneVision-Chat-72B. On single-step free-form questions, Claude still leads with an accuracy of 24.1%, and the gap between open-source models and closed-source models has widened. For multi-step free-form questions, we calculate two metrics: Step Accuracy Rate (SAR) and Question Completeness Rate (QCR). GPT-4o achieves the highest scores in both SAR and QCR, with 32.0% and 6.0%, respectively. The models’ performance on QCR reflects their inability to perform complex multi-visual mathematical reasoning tasks. The multi-step part of the overall accuracy in Table 4 is calculated using the QCR metric.

D. Results of CoT, 2-shot on 3 Question Types

We observe a distinct difference in how prompting strategies affect model performance across different question types as shown in Table 14. For multiple-choice questions, the addition of Chain-of-Thought (CoT) and 2-shot examples tend to decrease performance for most models. Specifically, out of the ten models we test, eight models perform best with the original prompt. This suggests that for multiple-choice questions, which typically require selecting an answer from given options, simpler prompts lead to better outcomes as they reduce potential confusion or overthinking induced by extra information.

In contrast, for free-form question types, CoT and 2-shot strategies bring about more significant performance improvements, particularly for multi-step problems. The complex nature of these questions benefits from the step-by-step reasoning facilitated by CoT and the illustrative examples provided by 2-shot prompting. Models like Claude-3.5-sonnet show a substantial increase in performance on single-step free-form questions when CoT and 2-shot examples are used, improving from 19.6% to 25.6%. Similarly, on multi-step free-form questions, models such as GPT-4o improved from 25.4% to 32.6% with the addition of these strategies.

These findings highlight that while CoT and few-shot prompting strategies may not universally enhance performance across all question types, they are particularly effective for free-form questions that require detailed reasoning and problem-solving steps. Incorporating these strategies can aid models in navigating the complexities of open-ended mathematical problems, thereby improving their overall reasoning capabilities.

E. Image Relevance

To the best of our knowledge, this study is the first to analyze multi-image mathematical tasks from the perspective of image relevance. Beyond relevance, we also investigate the impact of image quantity on model performance. However, no clear patterns are observed, suggesting that the number of images may not be a critical factor in determining performance. Instead, relevance appears to play a more significant role, as the interdependence among images demands a model’s ability to perform cross-image understanding, posing a greater challenge to its reasoning capabilities.

F. Data Collection and Annotation

F.1. Data Collection

We have crawled a large number of multimodal math test questions from ”Zujuan” totaling around 380k multimodal math questions. To format all questions for use, we process them by OCR engine like Mathpix interface. Due to inherent errors in the OCR engine, we introduce manual checks to ensure the accuracy of parsing results and to verify whether the questions belong to multimodal math problems. Specifically, we use the three-stage strategy outlined in Section 3 to complete the final filtering of the data, ensuring not only the high quality of the questions but also the high quality of the images.

In the 1st stage, text-image alignment refers to ensuring the number of referenced images matches those returned by the Mathpix API. For instance, if the question text mentions ’as shown in figure 4’ but the Mathpix returns 3 images, it is filtered out by our rule-based system. The 2nd stage applies rule-based filtering to detect missing text fields (e.g., missing answers or analysis) and categorize samples into multiple-choice and free-form subsets for further screening. The 3rd stage involves manual verification, where we filter out blurred images or images with text overlays. Our annotation team includes graduate students and a field expert.

F.2. Annotation

For the categorization of subjects and image relevance, we first obtain preliminary results through majority voting among three models: GPT-4o, Claude-3.5-Sonnet, and Qwen-VL-Max. For questions without a consensus from

the voting, no annotations are initially assigned. Subsequently, each question is independently reviewed by two graduate students specializing in relevant fields, with each producing an individual annotation result. Finally, questions with conflicting annotations are adjudicated by domain experts to determine the final MV-MATH annotations.

G. Comparison with Existing Benchmarks

We provide a detailed comparison between our MV-MATH and existing mathematical benchmarks in Table 15.

H. Case Study

In this section, we provide more detailed error examples of Claude-3.5-Sonnet. We classify Claude’s errors into five categories: Visual Perception Error, Reasoning Error, Calculation Error, Knowledge Error, and Reject Error. Detailed examples can be seen in Figures 21 to 30.

Model	Overall	AG	Algebra	MG	Combinatorics	TG	Logic	SG	Arithmetic	CG	DG	Statistics
LLMs(Text-only, CoT with 2-shot)												
Deepseek-Chat[9]	27.0	26.3	32.6	19.4	31.8	28.5	9.0	29.9	27.2	24.8	28.1	13.2
LLMs(Text + Image Caption, CoT with 2-shot)												
Deepseek-Chat[9]	27.5	27.7	30.1	18.4	36.8	30.5	10.9	32.9	26.8	26.8	28.7	17.3
Open-source MLLMs (Text + Image, CoT with 2-shot)												
Math-LLaVA-13B[36]	3.6	2.3	7.9	3.7	9.0	2.0	0.0	1.7	18.1	5.8	2.3	5.8
LLaVA-v1.5-13B[24]	8.6	7.0	10.8	7.4	9.0	10.0	9.0	5.9	18.1	13.1	5.4	11.7
LLaVA-v1.5-7B[24]	18.0	15.7	18.8	20.3	9.0	18.0	13.6	21.4	0.0	16.7	25.7	11.7
VILA-13B[22]	21.5	20.4	18.1	20.3	22.7	20.0	18.1	23.3	54.5	23.3	22.6	5.8
InternLM-XComposer2.5-VL-7B[10]	23.0	21.0	20.2	24.0	31.8	27.0	36.3	16.7	27.2	24.0	25.0	5.8
InternVL-Chat-8B[6]	25.9	23.9	31.8	25.9	31.8	27.5	27.2	25.1	36.3	22.6	21.8	23.5
Llama-3.2-Vision-Instruct-11B[11]	24.1	23.9	22.4	24.0	31.8	21.5	18.1	19.7	36.3	32.1	27.3	17.6
Deepseek-VL-7B[27]	26.0	26.3	32.6	19.4	31.8	28.5	9.0	29.9	27.2	24.8	28.1	29.4
LLaVA-NeXT-Interleave-7B[18]	26.6	23.9	23.1	26.8	31.8	25.5	18.1	31.7	9.0	30.6	26.5	23.5
Mantis-Idetics2-8B[16]	26.6	22.8	27.5	31.4	27.2	27.0	22.7	23.9	9.0	29.9	28.9	11.7
Mantis-siglip-8B[16]	27.8	29.8	27.5	28.7	22.7	27.0	22.7	26.9	27.2	21.8	35.1	17.6
Qwen2VL-Instruct-7B[41]	27.8	21.0	28.2	24.0	31.8	30.0	22.7	32.3	27.2	32.1	26.5	23.5
LLaVA-OneVision-SI-7B[17]	26.3	23.9	28.2	21.2	22.7	26.5	13.6	27.5	45.4	23.3	34.3	13.6
LLaVA-OneVision-SFT-7B[17]	30.1	30.4	28.6	54.5	26.2	27.2	28.1	9.0	37.9	31.7	30.0	23.5
LLaVA-OneVision-Chat-7B[17]	31.5	32.7	29.7	34.2	22.7	26.0	4.5	37.7	54.5	32.8	32.0	29.4
LLaVA-OneVision-SI-72B[17]	35.8	34.5	34.0	38.8	50.0	36.0	40.9	40.1	45.4	27.0	36.7	23.5
LLaVA-OneVision-SFT-72B[17]	37.0	29.8	42.0	28.7	36.3	37.5	31.8	47.9	18.1	32.8	39.8	29.4
LLaVA-OneVision-Chat-72B[17]	38.0	32.1	42.0	33.3	54.5	37.5	27.2	47.3	36.3	35.0	35.9	35.2
Closed-source MLLMs (Text + Image, CoT with 2-shot)												
Qwen-vl-plus[3]	27.7	30.4	25.3	25.0	18.1	29.0	31.8	33.5	45.4	22.6	24.2	29.4
GPT-4V[30]	32.5	21.0	34.7	44.4	36.3	32.5	36.3	31.7	45.4	32.1	33.5	35.2
Qwen-vl-max[3]	37.3	35.0	38.4	29.6	59.0	37.0	31.8	43.7	45.4	37.2	32.8	23.5
Gemini-1.5-Pro[37]	35.8	36.2	36.2	30.5	31.8	39.0	40.9	40.1	54.5	32.8	28.9	35.2
GPT-4o[32]	41.9	36.2	46.3	44.4	50.0	40.0	31.8	47.3	45.4	35.0	45.3	47.0
Claude-3.5[2]	44.0	42.6	49.2	46.3	59.0	43.5	31.8	44.9	72.7	39.4	39.0	41.1
Human Performance												
Human (testmini)	80.2	75.2	71.9	95.2	85.1	70.4	80.2	67.5	85.2	76.7	69.8	88.4

Table 10. Comparison of model performances across various mathematical subjects on the choice problem set. The first and second highest accuracy of LMMs are marked in red and blue, respectively.

Model	Overall	AG	Algebra	MG	Combinatorics	TG	SG	Arithmetic	CG	DG
LLMs(Text-only, CoT with 2-shot)										
Deepseek-Chat[9]	3.0	3.1	2.3	4.0	5.5	2.5	3.0	0.0	3.3	10.0
LLMs(Text + Image Caption, CoT with 2-shot)										
Deepseek-Chat[9]	4.3	5.2	3.4	5.3	0.0	1.2	3.0	20.0	2.9	20.0
Open-source MLLMs (Text + Image, CoT with 2-shot)										
Math-LLaVA-13B[36]	2.5	0.0	5.8	6.6	0.0	3.8	4.5	20.0	7.2	0.0
LLaVA-v1.5-13B[24]	0.6	2.1	1.1	0.0	0.0	6.4	1.5	0.0	0.4	0.0
LLaVA-v1.5-7B[24]	0.3	0.0	2.3	0.0	11.1	1.2	1.0	0.0	2.9	0.0
VILA-13B[22]	0.5	0.0	1.1	0.0	0.0	1.2	5.0	0.0	3.3	0.0
InternLM-XComposer2.5-VL-7B[10]	1.0	1.0	2.3	0.0	16.6	6.4	3.5	0.0	3.8	10.0
InternVL-Chat-8B[6]	0.3	1.0	4.6	8.0	5.5	1.2	3.0	20.0	2.9	0.0
Llama-3.2-Vision-Instruct-11B[11]	2.7	4.2	6.9	8.0	11.1	2.5	5.0	20.0	3.8	0.0
Deepseek-VL-7B[27]	1.5	0.0	2.3	0.0	0.0	0.0	3.0	0.0	3.3	10.0
LLaVA-NeXT-Interleave-7B[18]	2.7	1.0	4.6	1.3	0.0	1.2	3.5	0.0	5.0	0.0
Mantis-Idedfics2-8B[16]	2.0	1.0	4.6	5.3	0.0	7.7	3.0	0.0	8.4	0.0
Mantis-siglip-8B[16]	1.1	2.1	3.4	5.3	5.5	5.1	0.0	0.0	4.6	0.0
Qwen2VL-Instruct-7B[41]	2.8	3.1	4.6	4.0	0.0	2.5	4.5	40.0	6.7	10.0
LLaVA-OneVision-SI-7B[17]	6.7	5.2	10.4	5.3	11.1	6.4	7.5	20.0	9.3	10.0
LLaVA-OneVision-SFT-7B[17]	5.5	5.2	8.1	5.3	0.0	6.4	7.5	20.0	6.0	0.0
LLaVA-OneVision-Chat-7B[17]	4.8	1.0	8.1	4.0	5.5	3.8	6.0	20.0	6.7	20.0
LLaVA-OneVision-SI-72B[17]	13.1	9.4	11.6	12.0	0.0	6.4	16.6	20.0	19.9	20.0
LLaVA-OneVision-SFT-72B[17]	13.7	13.6	13.9	12.0	5.5	6.4	20.2	20.0	15.2	40.0
LLaVA-OneVision-Chat-72B[17]	13.1	11.5	17.4	12.0	11.1	5.1	18.1	20.0	15.2	40.0
Closed-source MLLMs (Text + Image, CoT with 2-shot)										
Qwen-vl-plus[3]	12.6	8.4	13.9	8.0	16.6	10.3	16.6	20.0	11.8	20.0
GPT-4V[30]	16.4	12.6	27.9	17.3	11.1	11.6	22.2	40.0	11.8	40.0
Qwen-vl-max[3]	15.8	13.6	21.1	18.6	11.1	10.3	20.2	20.0	15.6	30.0
Gemini-1.5-Pro[37]	23.4	16.8	26.7	26.6	22.2	11.6	28.7	20.0	15.2	70.0
GPT-4o[32]	22.5	21.0	25.5	28.0	27.7	14.2	34.3	40.0	15.2	50.0
Claude-3.5[2]	24.1	22.1	27.9	22.6	38.8	14.2	32.8	20.0	22.4	60.0
Human Performance										
Human (testmini)	73.2	69.8	72.9	86.2	80.1	65.8	61.7	85.9	70.6	71.3

Table 11. Comparison of model performances across various mathematical subjects on the single-step problem set. The first and second highest accuracy of MLLMs are marked in red and blue, respectively.

Model	Overall	AG	Algebra	MG	Combinatorics	TG	SG	Arithmetic	CG
LLMs(Text-only, CoT with 2-shot)									
Deepseek-Chat[9]	8.7	11.4	8.8	3.1	0.0	3.3	9.5	0.0	18.7
LLMs(Text + Image Caption, CoT with 2-shot)									
Deepseek-Chat[9]	9.1	11.6	7.7	5.2	0.0	6.6	19.0	0.0	19.4
Open-source MLLMs (Text + Image, CoT with 2-shot)									
Math-LLaVA-13B[36]	4.1	2.6	4.4	3.1	0.0	15.0	0.0	0.0	4.1
LLaVA-v1.5-13B[24]	2.5	1.7	2.2	2.0	0.0	0.0	4.7	0.0	6.9
LLaVA-v1.5-7B[24]	1.6	3.5	0.0	2.0	0.0	0.0	0.0	0.0	0.0
VILA-13B[22]	2.0	0.0	2.2	0.0	0.0	3.3	4.7	0.0	9.0
InternLM-XComposer2.5-VL-7B[10]	2.1	0.8	2.2	2.0	0.0	3.3	0.0	0.0	6.9
InternVL-Chat-8B[6]	4.0	3.0	5.5	5.2	0.0	3.3	9.5	0.0	2.0
Llama-3.2-Vision-Instruct-11B[11]	5.0	1.7	0.0	2.0	50.0	0.0	19.0	0.0	18.7
Deepseek-VL-7B[27]	2.3	0.0	8.8	3.1	0.0	0.0	4.7	0.0	0.0
LLaVA-NeXT-Interleave-7B[18]	5.1	7.1	7.7	2.0	0.0	3.3	9.5	0.0	4.8
Mantis-Idedfics2-8B[16]	1.8	1.3	2.2	1.5	0.0	0.0	0.0	0.0	6.2
Mantis-siglip-8B[16]	5.5	3.9	10.0	4.1	0.0	6.6	4.7	0.0	6.9
Qwen2VL-Instruct-7B[41]	10.8	11.6	11.1	3.1	0.0	10.0	16.6	0.0	19.4
LLaVA-OneVision-SI-7B[17]	16.0	11.4	23.3	8.3	0.0	10.0	35.7	0.0	28.4
LLaVA-OneVision-SFT-7B[17]	17.8	14.9	22.2	16.6	0.0	18.3	14.2	0.0	27.7
LLaVA-OneVision-Chat-7B[17]	18.3	14.9	22.2	16.6	0.0	18.3	14.2	0.0	27.7
LLaVA-OneVision-SI-72B[17]	24.0	26.7	23.3	25.5	50.0	15.0	28.5	0.0	19.4
LLaVA-OneVision-SFT-72B[17]	25.9	25.4	33.3	13.0	50.0	30.0	33.3	0.0	27.7
LLaVA-OneVision-Chat-72B[17]	26.2	27.1	30.0	16.6	50.0	18.3	36.9	0.0	27.7
Closed-source MLLMs (Text + Image, CoT with 2-shot)									
Qwen-vl-plus[3]	24.2	26.7	23.3	16.6	50.0	18.3	28.5	0.0	19.4
GPT-4V[30]	23.8	23.4	24.8	22.3	50.0	15.0	40.4	100.0	20.1
Qwen-vl-max[3]	29.5	29.1	36.6	19.7	0.0	30.0	28.5	0.0	29.1
Gemini-1.5-Pro[37]	31.7	34.4	38.8	21.8	50.0	36.6	30.9	0.0	24.3
GPT-4o[32]	32.6	30.2	43.3	21.3	100.0	28.3	38.0	0.0	34.0
Claude-3.5[2]	32.3	30.7	26.6	32.8	0.0	35.0	38.0	0.0	38.8
Human Performance									
Human (testmini)	78.5	75.2	76.4	83.2	79.1	72.4	66.5	79.0	71.7

Table 12. Comparison of model performances across various mathematical subjects on the multi-step problem set(SAR). The first and second highest accuracy of MLLMs are marked in red and blue, respectively.

Model	Overall	AG	Algebra	MG	Combinatorics	TG	SG	Arithmetic	CG
LLMs(Text-only, CoT with 2-shot)									
Deepseek-Chat[9]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLMs(Text + Image Caption, CoT with 2-shot)									
Deepseek-Chat[9]	1.0	0.0	0.0	0.0	0.0	0.0	14.2	0.0	0.0
Open-source MLLMs (Text + Image, CoT with 2-shot)									
Math-LLaVA-13B[36]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLaVA-v1.5-13B[24]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLaVA-v1.5-7B[24]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
VILA-13B[22]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
InternLM-XComposer2.5-VL-7B[10]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
InternVL-Chat-8B[6]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Llama-3.2-Vision-Instruct-11B[11]	1.0	0.0	0.0	0.0	0.0	0.0	14.2	0.0	0.0
Deepseek-VL-7B[27]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLaVA-NeXT-Interleave-7B[18]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Mantis-Idedfics2-8B[16]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Mantis-siglip-8B[16]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Qwen2VL-Instruct-7B[41]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLaVA-OneVision-SI-7B[17]	1.0	0.0	0.0	0.0	0.0	0.0	14.2	0.0	0.0
LLaVA-OneVision-SFT-7B[17]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLaVA-OneVision-Chat-7B[17]	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
LLaVA-OneVision-SI-72B[17]	1.0	0.0	0.0	0.0	0.0	0.0	14.2	0.0	0.0
LLaVA-OneVision-SFT-72B[17]	2.0	0.0	6.6	0.0	0.0	0.0	14.2	0.0	0.0
LLaVA-OneVision-Chat-72B[17]	2.0	0.0	6.6	0.0	0.0	0.0	14.2	0.0	0.0
Closed-source MLLMs (Text + Image, CoT with 2-shot)									
Qwen-vl-plus[3]	1.0	0.0	0.0	0.0	0.0	0.0	14.2	0.0	0.0
GPT-4V[30]	3.0	0.0	6.6	0.0	0.0	0.0	28.5	0.0	0.0
Qwen-vl-max[3]	2.0	0.0	6.6	0.0	0.0	0.0	14.2	0.0	0.0
Gemini-1.5-Pro[37]	5.0	0.0	6.6	12.5	0.0	0.0	28.5	0.0	0.0
GPT-4o[32]	6.0	0.0	6.6	6.2	50.0	0.0	28.5	0.0	8.3
Claude-3.5[2]	4.0	0.0	0.0	6.2	0.0	10.0	28.5	0.0	0.0
Human Performance									
Human (testmini)	66.0	60.6	66.7	52.2	100.0	70.0	71.4	63.2	75.0

Table 13. Comparison of model performances across various mathematical subjects on the multi-step problem set(QCR).

Models	Multiple-Choice			Single-Step Free-Form			Multi-Step Free-Form(SAR)		
	Original	+CoT	+2-shot & CoT	Original	+CoT	+2-shot & CoT	Original	+CoT	+2-shot & CoT
Closed-source Models									
Claude-3.5-sonnet [2]	39.5	41.6(+2.1)	44.0(+2.4)	19.6	25.3(+5.7)	24.1(-1.2)	30.4	32.7(+2.3)	32.3(-0.4)
GPT-4o [32]	41.9	36.2(-5.7)	41.9(+5.7)	24.1	28.1(+4.0)	22.5(-5.6)	25.4	31.6(+6.2)	32.6(+1.0)
Gemini-1.5-pro [37]	40.5	35.0(-5.5)	35.8(+0.8)	19.7	22.6(+2.9)	23.4(+0.8)	31.1	33.9(+2.8)	31.7(-2.2)
Qwen-vl-max [3]	41.5	41.3(-0.2)	37.3(-4.0)	15.5	17.1(+1.6)	15.8(-1.3)	29.4	29.0(-0.4)	29.5(+0.5)
GPT-4V [30]	28.9	33.2(+4.3)	32.5(-0.7)	19.8	18.2(-1.6)	16.4(-1.8)	20.0	23.2(+3.2)	23.8(+0.6)
Open-source Models									
LLaVA-OneVision-Chat-72B [17]	39.3	38.0(-1.3)	38.0(0.0)	15.8	14.2(-1.6)	13.1(-1.1)	25.2	24.0(-1.2)	26.2(+2.2)
LLaVA-OneVision-Chat-7B [17]	31.6	31.0(-0.6)	31.5(+0.5)	9.8	8.8(-1.0)	4.8(-4.0)	14.7	18.6(+3.9)	18.3(-0.3)
LLaVA-NeXT-Interleave-7B [18]	29.7	29.4(-0.3)	26.6(-2.8)	6.0	3.1(-2.9)	2.7(-0.4)	11.5	6.8(-4.7)	5.1(-1.7)
Qwen2VL-Instruct-7B [41]	33.6	28.2(-5.4)	27.8(-0.4)	6.2	5.2(-1.0)	2.8(-2.4)	5.5	10.7(+5.2)	10.8(+0.1)
Deepseek-VL-Chat-7B [27]	29.0	28.8(-0.2)	26.0(-2.8)	6.0	2.7(-3.3)	1.5(-1.2)	5.1	2.5(-2.6)	2.3(-0.2)

Table 14. Model Performance Evaluation across 3 Question Types and Configurations. The best performance in each category is **bolded**. Improvements over the previous configuration are indicated in **red** for increases and **green** for decreases.

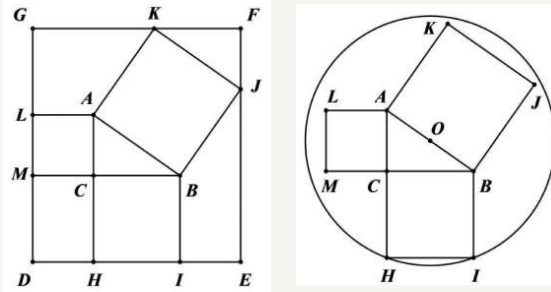
Benchmarks	Language	Multi-image	Avg Question length	Avg Analysis length	Image Relevance Annotation	Subject Number	Size	Source	MC	FF	MS
MathVista	EN	✗	77.1	✗	✗	7	1K/6K	Synthesized	✓	✗	✗
GeoQA	EN	✗	37.1	58.2	✗	1	0.7K	Internet	✓	✗	✗
MATH-Vision	EN	✗	42.4	✗	✗	16	3K	Synthesized	✓	✗	✗
MMMU-math	EN	✗	40.8	✗	✗	8	0.5K	Textbook	✓	✓	✗
Mathverse-mv	EN	✓	76.9	✗	✗	1	0.8K	Synthesized	✓	✗	✗
CMM-math	CN	✓	-	-	✗	7	0.7K	Internet	✓	✓	✗
MV-MATH(Ours)	EN	✓	80.2	150.9	✓	11	2K	Internet&Annotated	✓	✓	✓

Table 15. Comparison with existing multimodal math benchmarks. MC: Multiple Choice, FF: Free-form, MS:Multi-Step.

Visual Perception Error

Question: A rectangular sheet of paper and a circular sheet of paper are cut as shown in the figure, each resulting in a similar theorem exception figure. $AC = 3$, $BC = 4$, $AB = 5$, with squares cut out along each side. In Figure 1 the sides HI and LM and the points K and J are all on the edges of the rectangular paper. In Figure 2, the center O of the circle is at the midpoint of AB , and points H and I are on the circle. The ratio of the area of the rectangular paper to the area of the circular paper is ()

- A. $400:127\pi$
- B. $484:145\pi$
- C. $440:137\pi$
- D. $88:25\pi$



Ground Truth: C

Claude-3.5-Sonnet Answer

Let's approach this problem step by step:

Step 1: The rectangular paper's dimensions are $AC + BC = 3 + 4 = 7$ units wide and $AB = 5$ units long.

Step 2: Area of rectangle = $7 * 5 = 35$ square units.

Step 3: In the circular paper, O is the midpoint of AB . So, the radius of the circle is half of AB . Radius = $5/2 = 2.5$ units.

Step 4: Area of circle = $\pi r^2 = \pi * (2.5)^2 = 6.25\pi$ square units.

Step 5: The cut-out squares have sides equal to the shortest side of the triangle, which is $AC = 3$ units.

Step 6: Area of one cut-out square = $3^2 = 9$ square units

Total area of cut-out squares = $9 * 4 = 36$ square units

Step 7: Remaining area = $35 - 36 = -1$ square unit

This negative value indicates that the cut-out squares overlap, and their total area is actually less than 36.

Step 8: The actual area of cut-out squares is the area of the rectangle minus the area of the theorem exception figure.

Area of theorem exception figure = $3 * 4 = 12$ square units

Actual area of cut-out squares = $35 - 12 = 23$ square units

Step 9: Remaining area of circle = $6.25\pi - 23$ square units

Step 10: Ratio = Area of rectangular paper : Area of circular paper = $35 : (6.25\pi) = 560 : 100\pi = 140 : 25\pi = 28 : 5\pi$

Step 11: A. $400 : 127\pi \approx 3.15 : \pi$ B. $484 : 145\pi \approx 3.34 : \pi$ C. $440 : 137\pi \approx 3.21 : \pi$



D. $88 : 25\pi = 3.52 : \pi$ Our calculated ratio $28 : 5\pi = 5.6 : \pi$ is closest to option D.

Answer: {D}


Figure 21. An example of visual perception error.

Visual Perception Error


Question: As shown in the figure, if the cube is unfolded, which of the following patterns can be obtained?




A



B



C



D

Ground Truth: D

Claude-3.5-Sonnet Answer

To determine which pattern can be obtained by unfolding the given cube, we need to visualize how the cube's faces are connected and how they would lay flat when unfolded.

Step 1: Identify the faces of the cube.

-The cube has six faces: one with a circle, one with a triangle, and four plain faces.

Step 2: Visualize the unfolding process.

-When the cube is unfolded, the faces will be laid out in a connected manner. The face with the circle will be adjacent to the face with the triangle, and the plain faces will be adjacent to these two faces.

Step 3: Compare the given options with the expected layout.


- Option A: The layout shows the circle and triangle faces adjacent to each other, with the plain faces correctly positioned.
- Option B: The layout shows the circle and triangle faces adjacent to each other, but the plain faces are not correctly positioned.
- Option C: The layout shows the circle and triangle faces adjacent to each other, but the plain faces are not correctly positioned.
- Option D: The layout shows the circle and triangle faces adjacent to each other, but the plain faces are not correctly positioned.
- Answer: {A}

Figure 22. An example of visual perception error.

Reasoning Error

Question: In a math class, the teacher leads the students to use four right-angled triangles with legs a and b as shown in Figure 1 to form two *chord diagrams* (as shown in Figure 2 and Figure 3). Given that the area of the large square $ABCD$ is 48 and the area of the small square $EFGH$ is 12, what is the area of one right-angled triangle.

- A. 3 B. 4.5 C. 6 D. 9



Ground Truth: B

Claude-3.5-Sonnet Answer

Step 1 (Understand the problem): We need to find the area of one right-angled triangle given the areas of the large square $ABCD$ and the small square $EFGH$.

Step 2 (Mathematical basis): The area of the large square $ABCD$ is given as 48, and the area of the small square $EFGH$ is given as 12. The large square $ABCD$ is composed of the small square $EFGH$ and four right-angled triangles.

Step 3 (Calculate the total area of the four triangles): The area of the large square $ABCD$ minus the area of the small square $EFGH$ gives the total area of the four right-angled triangles.

$$\begin{aligned} \text{Total area of four triangles} \\ &= \{\text{Area of } \} ABCD - \{\text{Area of } \} EFGH \\ \text{Total area of four triangles} &= 48 - 12 = 36 \end{aligned}$$

Step 4 (Calculate the area of one triangle): Since there are four right-angled triangles, the area of one triangle is:


$$\begin{aligned} \text{Area of one triangle} &= \{\text{Total area of four triangles}\} / \{4\} \\ \text{Area of one triangle} &= \frac{36}{4} = 9 \end{aligned}$$

Answer: {D}


Figure 23. An example of reasoning error.

Reasoning Error


Question: An object is composed of multiple small cubes with edge lengths of 1. The shapes of the geometric body as seen from the front, left, and top are shown in the figures below. Then, the volume of the geometric body is _____.



View from the front



View from the left



View from the top

Ground Truth: 8

Claude-3.5-Sonnet Answer

To find the volume of the geometric body composed of small cubes with edge lengths of 1, we need to analyze the views from the front, left, and top.

Step 1 (Analyze the front view):

- The front view shows a structure with 3 columns.
- The left and right columns each have 1 cube. The middle column has 2.

Step 2 (Analyze the left view):

- The left view shows a structure with 2 columns. The left column has 1 cube. The right column has 2 cubes.

Step 3 (Analyze the top view):

- The top view shows a structure with 5 cubes arranged in a cross pattern.
- This confirms the layout seen from the front and left views.

Step 4 (Determine the number of cubes):

- From the front view, we have 1 cube in the left column, 2 cubes in the middle column, and 1 cube in the right column.
- From the left view, we have 1 cube in the left column and 2 cubes in the right column. Combining these views, we can deduce the following:
- The left column has 1 cube. The middle column has 2 cubes.
- The right column has 1 cube. The top view confirms that there are 5 cubes in total.


Step 5 (Calculate the volume): Since each cube has an edge length of 1, the volume of each cube is $1^3 = 1$. The total volume is the number of cubes multiplied by the volume of each cube.

Answer: The volume of the geometric body is 5.

Figure 24. An example of reasoning error.

Knowledge Error

Question: As shown in Figure 1, in rectangle $ABCD$, point E is on AD , and point A is folded to the right along the fold line BE as shown in Figure 2. Through point A , $AF \perp DC$ at F . If $AB = 6\sqrt{3}$, $BC = 13$, $\angle BEA = 60^\circ$, then in Figure 3, $AF = \underline{\hspace{2cm}}$.



Ground Truth: 4

Claude-3.5-Sonnet Answer

Step 1: The goal is to calculate AF after point A is folded to the right along BE . Important details provided:

$AB = 6\sqrt{3}$, $BC = 13$, and $\angle BEA = 60^\circ$. After folding, $AF \perp DC$

Step 2: Using triangle $\triangle ABE$ and the given data $\angle BEA = 60^\circ$:

Assume BE is calculated using trigonometry: $BE = AB \cdot \cos(60^\circ)$

Substituting $AB = 6\sqrt{3}$:

$$BE = 6\sqrt{3} \cdot \cos(60^\circ) = 6\sqrt{3} \cdot \frac{1}{2} = 3\sqrt{3}.$$

Step 3: Using $\triangle ABE$ again, AE is found using the Pythagorean theorem: $AE^2 = AB^2 - BE^2$. Substitute $AB = 6\sqrt{3}$ and $BE = 3\sqrt{3}$ (from the earlier step):

$$AE^2 = (6\sqrt{3})^2 - (3\sqrt{3})^2 = 108 - 27 = 81$$

$$AE = \sqrt{81} = 9$$






Step 4: To find AF , assume that AF is directly proportional to BE due to the perpendicular relationship between AF and DC . Using the proportions:

$$AF = \frac{AB}{2} = \frac{6\sqrt{3}}{2} = 3\sqrt{3}$$

Figure 25. An example of knowledge error.

Knowledge Error

Question: As shown in the figure, the diameter AB of the semicircle O is 4, and C is the midpoint of the arc AB . The segments CO , CA , and CB are drawn. Point P starts from A and moves along $A \rightarrow O \rightarrow C$ until it reaches C . Two perpendicular lines PE and PF are drawn from point P to AC at E and to BC at F , respectively. Let the distance traveled by point P be x . The area y of quadrilateral $CEPF$ as a function of x can be best represented by the graph()



Ground Truth: A

Claude-3.5-Sonnet Answer

Step 1 (Analyze the problem): As P moves from A to O to C , the area of quadrilateral $CEPF$ changes. We need to analyze how the area changes in different stages.

Step 2 (Stage 1: P is on AO): When P is on AO ($0 < x < 2$), the area of $CEPF$ increases as P moves from A to O . This is because both PE and PF increase as P moves towards O .

Step 3 (Stage 2: P is on OC): When P is on OC ($2 < x < 2 + \sqrt{2}$), the area of $CEPF$ decreases as P moves from O to C . This is because while PE decreases, PF increases, but the overall effect is a decrease in the area.

Step 4 (Analyze the options): Based on the above analysis, we can eliminate options A, C, and D. Option A shows a continuous increase in the area, which is incorrect. Option C shows a linear increase and then a linear decrease, which is also incorrect. Option D shows a decrease in the area initially, which is not possible.

Step 5 (Final analysis): Option B correctly represents the change in the area of $CEPF$. The area initially increases, reaches a maximum when P is at O , and then decreases as P moves towards C .

Answer: {B}



Figure 26. An example of knowledge error.

Calculation Error

Question: As shown in the figure, in $\triangle ABC$, $CD \perp BA$ and CD intersects the extension line of BA at D , $DE \perp AC$ at E .

(1) As shown in Figure 1, if $\angle B = 35^\circ$, $\angle CDE = 60^\circ$, find the degree of $\angle ACB$;

(2) As shown in Figure 2, if AC bisects $\angle BCD$, $BF \perp AC$ and BF intersects the extension of CA at F , directly write out the angle equal to $\angle ACB$ (except $\angle ACB$).



Ground Truth: (1) 25° (2) $\angle ACD$, $\angle ABF$, $\angle ADE$

Claude-3.5-Sonnet Answer

(1) Step 1: Identify the Relationship Between Angles in $\triangle ABC$
 $\angle CDE = 60^\circ$, and $DE \perp AC$, $\angle B = 35^\circ$, implies that:

$$\angle DCE = 90^\circ - 60^\circ = 30^\circ$$

Step 2: $\angle ACB$ is equal to $\angle BCD$ minus $\angle DEC$:

$$\angle ACB = \angle BCD - \angle DCE = 55^\circ - 30^\circ = 25^\circ$$

(2) Step 1: Analyze the Geometry

AC bisects $\angle BCD$, so:

$$\angle ACD = \frac{\angle BCD}{2}$$

$BF \perp AC$, and BF intersects the extension of CA at F .

Step 2: Assume that $\angle BCD = 2 \times \angle ACB$

So we have:


$$\angle ABC = \angle CBF - \angle ABF = \angle ABF$$

$$\angle ACD = \frac{\angle BCD}{2} = \angle ACB = \angle ABF$$

Figure 27. An example of calculation error.

Calculation Error

Question: As shown in Figure 1, a circle with a radius of 7cm is divided into 20 equal parts, then cut out and rearranged into the shape shown in Figure 2. The perimeter of the resulting shape is_____.



Ground Truth: 57.96

Claude-3.5-Sonnet Answer

Step 1 (Understanding the problem): The circle with a radius of 7 cm is divided into 20 equal parts and rearranged into a shape that resembles a parallelogram.

Step 2 (Calculate the circumference of the circle): The circumference C of a circle is given by the formula:

$$C = 2\pi r$$

where r is the radius of the circle. Given $r=7\text{cm}$:

$$C = 2\pi \cdot 7 = 14\pi$$

Step 3 (Determine the length of the base of the parallelogram): Since the circle is divided into 20 equal parts, each part represents an arc of the circle. When rearranged, these arcs form the base of the parallelogram. The length of the base is half the circumference of the circle:

$$\text{base length} = \frac{14\pi}{2} = 7\pi$$

Step 4 (Determine the height of the parallelogram): The height of the parallelogram is equal to the radius of the circle, which is 7cm .

Step 5 (Calculate the perimeter of the parallelogram): The perimeter P of a parallelogram is given by:



$$P = 2 \cdot \text{baselength} + \text{height} = 2 \cdot (7\pi + 7) = 14(\pi + 1) = 58.94$$

Answer: $\{58.94\text{ cm}\}$

Figure 28. An example of calculation error.

Reject Error

Question: There is a cylindrical wooden block buried in the wall (as shown in Figure (1)). Its diameter is unknown, so it is sawed open (along the cross section) with a saw. (As shown in Figure (2) When the measured depth CE is 1 inch, the sawed width AB is 1 foot, and the diameter CD of the wooden block is _____ inches. (1 foot = 10 inches)



Ground Truth: 26

Claude-3.5-Sonnet Answer


Sorry, I cannot help you with this request.

Figure 29. An example of reject error.

Reject Error

Question: As shown in Figure 1, a large tree broke and fell during a strong earthquake, with the top of the tree landing 12 meters away from the base and 5 meters above the ground. Figure 2 is a diagram illustrating the tree's breakage. Determine the height of the tree before it broke.

A. 20 meters B. 18 meters C. 16 meters D. 15 meters



Ground Truth: B

Claude-3.5-Sonnet Answer

I'm sorry, but I am unable to solve this problem.

Figure 30. An example of reject error.