

# MAMmoTH-VL: Eliciting Multimodal Reasoning with Instruction Tuning at Scale

Jarvis Guo<sup>♣\*</sup>, Tuney Zheng<sup>♣\*</sup>, Yizhi Li<sup>◇</sup>, Yuelin Bai<sup>♣\*</sup>,  
Bo Li<sup>△</sup>, Yubo Wang<sup>♡</sup>, King Zhu<sup>♣\*</sup>,  
Graham Neubig<sup>♣\*</sup>, Wenhui Chen<sup>♡</sup>, Xiang Yue<sup>♣†</sup>

<sup>♣</sup>Carnegie Mellon University <sup>♣\*</sup>M-A-P

<sup>△</sup>Nanyang Technological University <sup>♡</sup>University of Waterloo

<sup>◇</sup>The University of Manchester




Figure 1: Scaling effects of MAMmoTH-VL-8B on four multimodal evaluation datasets. A *simple* rewriting approach using open models improves the quality of visual instruction data by eliciting chain-of-thought (CoT) reasoning. Training on this *rewritten* data demonstrates significant performance gains through increased model scale. Llava-OneVision-7B&72B (Li et al., 2024b) and Llava-CoT (Xu et al., 2024a) are included as references.

## Abstract

Open-source multimodal large language models (MLLMs) have shown significant potential in a broad range of tasks. However, their reasoning capabilities remain constrained by existing instruction-tuning datasets, which were predominately repurposed from academic datasets such as VQA, AI2D, and ChartQA. These datasets target simplistic tasks, and only provide phrase-level answers without any intermediate rationales. To address these challenges, we introduce a *scalable and cost-effective method* to construct a large-scale multimodal instruction-tuning dataset with *rich intermediate rationales designed to elicit CoT reasoning*. Using only open models, we create a dataset containing 12M instruction-response pairs to cover diverse reasoning-intensive tasks. Experiments demonstrate that training MLLMs on our dataset not only significantly improves reasoning capabilities, achieving state-of-the-art performance on benchmarks such as MathVerse (+8.1%), MMMU-Pro (+7%), and MuirBench (+13.3%), but also gains improvements of up to 4% on non-reasoning-based benchmarks.

## 1 Introduction

Recent open-source multimodal large language models (MLLMs) like LLaVA (Liu et al., 2024b; Li et al., 2024b) have shown promise in combining visual encoders (Radford et al., 2021; Zhai et al., 2023) with LLMs (Chiang et al., 2023; Bai et al.,

2023; Touvron et al., 2023; Young et al., 2024). However, these models often struggle with complex reasoning tasks (Lu et al., 2023a; Zhang et al., 2024c; Yue et al., 2024a,b), primarily due to limitations in existing instruction-tuning datasets.

Current multimodal instruction datasets predominantly target simplistic visual question answering (VQA) tasks, such as identifying objects in images (“*What is shown in the image?*”) or generating straightforward captions (“*Can you describe the image?*”). A key reason for this is that many of these datasets are repurposed from academic VQA datasets (Antol et al., 2015; Singh et al., 2019), which focus on phrase-based answers rather than nuanced reasoning (Tong et al., 2024; Bai et al., 2024). Consequently, these datasets fail to elicit deliberate reasoning from multimodal models. The absence of deliberate reasoning not only limits interpretability but also hampers performance on tasks that demand contextual understanding.

While Chain-of-Thought (CoT) reasoning has proven effective for text-based LLMs to enhance the reasoning capability (Kojima et al., 2022; Wei et al., 2022), creating large-scale datasets that elicit CoT reasoning remains challenging for multimodal learning (Yue et al., 2024c). The creation of such datasets faces key obstacles: 1) ensuring instruction diversity and complexity, and 2) generating coherent responses with detailed rationales. Human annotation is costly (Xu et al., 2024c; Deitke et al., 2024), and using proprietary models like GPT-4 (Luo et al., 2024; Zhang et al., 2024f; Xu

\*Equal contribution.

† Corresponding author: xyue2@andrew.cmu.edu

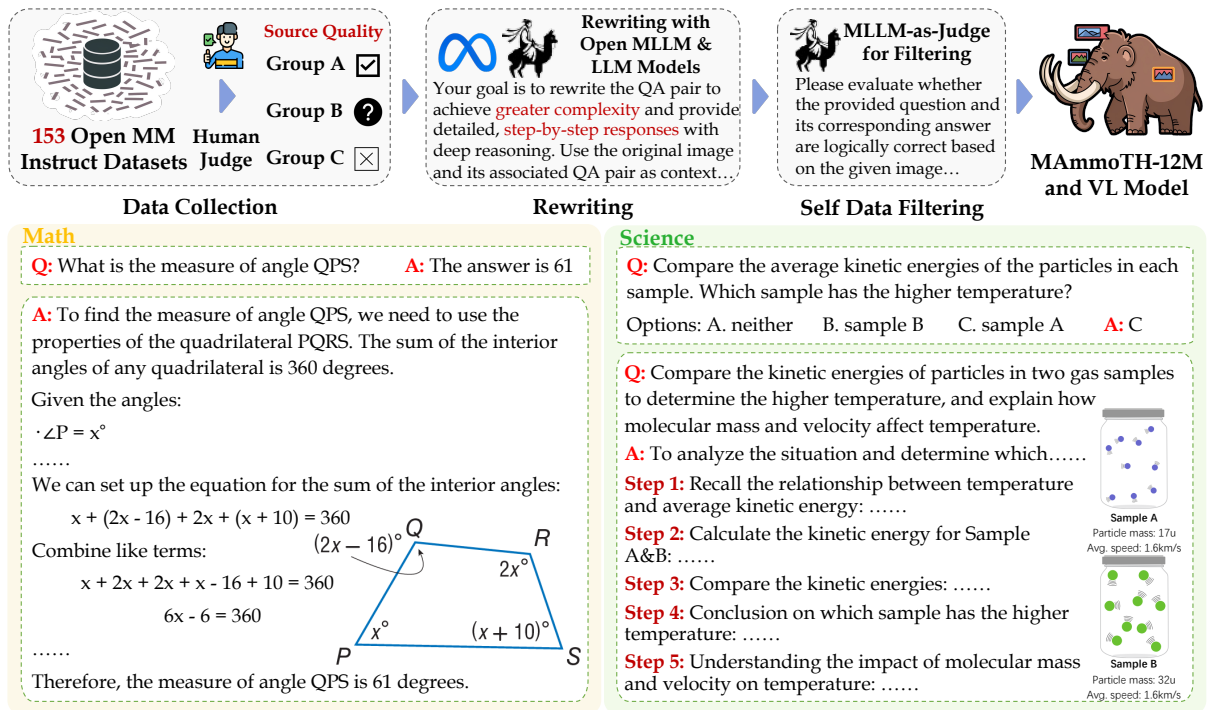


Figure 2: Overview of our *simple yet scalable* visual instruction data rewriting pipeline with three steps: manual data source collection, rewriting using MLLMs/LLMs, and filtering via the same MLLM as a judge. Examples below illustrate transformations in math and science categories, showcasing detailed, step-by-step responses.

et al., 2024a) presents financial and licensing barriers for the open-source community. We address these challenges by introducing a simple, scalable, and cost-effective methodology for creating large-scale instruction-tuning datasets that promote multimodal CoT reasoning. Our approach generates a dataset of 12 million entries using open-weight LLMs (Touvron et al., 2023) and MLLMs (Chen et al., 2023b). As illustrated in Figure 2, the dataset is constructed through a three-step pipeline: categorizing diverse image data, augmenting tasks with CoT rationales, and implementing rigorous filtering to ensure quality and minimize hallucinations.

Experimental results demonstrate the efficacy of this approach. By training an MLLM, MAMmoTH-VL-8B, based on the LLaVA-OneVision architecture (Li et al., 2024b), on our curated dataset, we achieve substantial improvements across multiple benchmarks, particularly in tasks requiring intricate reasoning and alignment between text and images. For example, MAMmoTH-VL-8B achieves an 8.1% improvement on MathVerse (Zhang et al., 2024c), 7% on MMMU-Pro (Yue et al., 2024b), and 13.3% on MuirBench (Wang et al., 2024a) compared to the open SoTA models. On other benchmarks, our model also delivers notable improvements. Ablation studies further reveal critical insights, including the importance of *self-filtering* for mitigating hallucinations, the effectiveness of

mixing rewritten and original data for enhancing task diversity, and the significant impact of scaling both training data size and rewrite model capacity on performance.

## 2 Method

While previous efforts have highlighted the potential of visual instruction tuning, many rely on resource-intensive methods such as human annotations or proprietary models. These approaches limit scalability and accessibility, particularly in open-source contexts. To address these challenges, we introduce a *simple, scalable, and cost-effective data generation pipeline* that produces 12 million high-quality samples. Our pipeline involves three key steps: (1) open-source data collection and categorization, (2) task-specific data augmentation and rewriting using open models, and (3) quality filtering to remove hallucinated or irrelevant content.

### 2.1 Dataset Collection and Categorization

To build a comprehensive and diverse multimodal dataset, we systematically collect and categorize data that span a wide range of real-world tasks and scenarios. This fundamental step ensures the robustness of the data set, enabling MLLMs to acquire general capabilities for a variety of tasks.

**Collection and Processing.** To achieve both scale

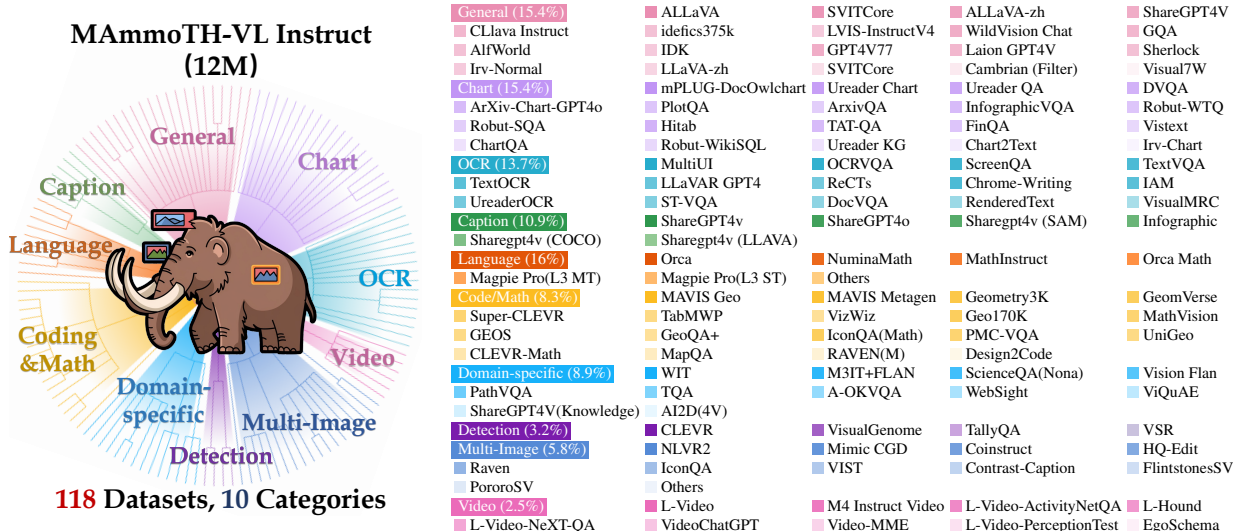


Figure 3: The data distribution of **MAMmoTH-VL-Instruct (12M)**. Left: Category distribution. Right: Details of data sources.

and diversity while maintaining accessibility for open-source initiatives, we sourced data from 153 publicly available multimodal instruction datasets, expanding upon prior collection efforts (Li et al., 2024b; Tong et al., 2024). The raw data includes image-text pairs covering a broad spectrum of use cases such as OCR, charts, captioning, and domain-specific images (e.g., medical). Based on MLLM training paradigms and common downstream tasks, we reorganized the training data into ten major categories: General, OCR, Chart, Caption, Domain-specific, Code&Math, Language, Detection, Multi-Image, and Video, as illustrated in Figure 3. This structured categorization is critical for enhancing data quality during the rewriting stage, allowing us to apply tailored strategies to different scenarios. Moreover, this categorization facilitates fine-tuning models for specific applications, providing a valuable resource for the broader community. In addition, all images were standardized according to verified criteria (Li et al., 2024b). This included resizing images to ensure dimensions fell between 224 and 4,096 pixels and adjusting extreme aspect ratios (e.g., greater than 7:1 or less than 1:7) to produce consistent formatting. Further details are illustrated in Figure A4, Figure A5 and Figure A6.

**Data Source Screening.** The quality of the collected open-source datasets varies significantly, requiring careful assessment to ensure baseline quality and adequate coverage for task-specific rewriting. To address this, we conducted an initial manual screening of the 153 data sources. This screening process aimed to identify datasets containing rich textual information with strong modality correlations, which are essential for producing high-

quality multimodal data. For each dataset, we randomly sampled 1,000 data points and performed a rapid evaluation to gauge overall quality. Based on this assessment, we categorized the datasets into three groups: **Group A (58 datasets)**, contains well-structured, detailed, and accurate responses aligned with task-oriented needs, no modifications were required; **Group B (60 datasets)**, includes brief or incomplete responses with potential for improvement; **Group C (35 datasets)**, contain overly brief, vague, or shallow responses lacking sufficient context or reasoning, which were removed.

## 2.2 Instruction Data Rewriting

To enhance the quality of **datasets in Group B**, we implemented a task-aware rewriting process aimed at addressing two key shortcomings: 1) a lack of detailed intermediate rationales, as many datasets originate from academic visual question answering contexts where responses are typically concise (e.g., a single word or phrase); and 2) limited coverage of real-world tasks such as reasoning, data analysis, code generation, debugging, and other practical applications. Our accordingly propose to transform the original multimodal data into new diverse, coherent, and task-specific instruction-response pairs with *detailed rationales*.

To guide this transformation for instruction tuning datasets, we designed customized prompts tailored for each data category. These prompts are crafted to generate responses that align with real-world applications while encouraging critical thinking and reasoning. Each prompt undergoes an iterative refinement process: initial drafts are tested on sampled data points within a specific category,




Figure 4: Comparison of original and rewritten data across two metrics: (1) Content and Relevance Scores judged by MLLMs show that rewritten data scores higher, indicating improved quality; (2) Token Length distribution suggests that rewritten data tends to be longer, including more tokens for rationales.

and the outputs are manually analyzed to identify areas for improvement. Once finalized, the optimized prompts are applied for rewriting. The full prompts are provided in Appendix subsection E.1. Specifically, for the caption-based data, we employ a text-only model (Llama-3-70B-Instruct (Meta, 2024a)) to generate task-oriented Q&A pairs since captions typically contain rich textual information. We observe that the text-only models are better suited for creating diverse and complex instructions compared to multimodal models in this setting.

### 2.3 Self-data Filtering

A preliminary manual inspection of the rewritten data revealed instances of hallucinations, particularly in tasks such as OCR and chart interpretation. This underscores the necessity of a robust data filtering step to enhance the quality of the generated content. While human review is ideal for this purpose, it is both costly and impractical for large-scale applications. To address this challenge, we propose utilizing a "Model-as-Judge" approach to filter the data efficiently. Specifically, we leverage the InternVL2-Llama3-76B model, the same model used during the data rewriting process, to evaluate the logical consistency of each question-answer pair against the corresponding image. The assumption is that while the model may introduce inaccuracies during generation, it excels better in verification tasks. By implementing this filtering step, we ensure that the rewritten instructional data aligns closely with the visual information provided, minimizing hallucinations.

## 3 Analysis of MAMmoTH-VL-Instruct

### 3.1 Quality Verification

To assess the overall quality of MAMmoTH-VL-Instruct, we randomly sample 1,000 data points from both the original and rewritten datasets in

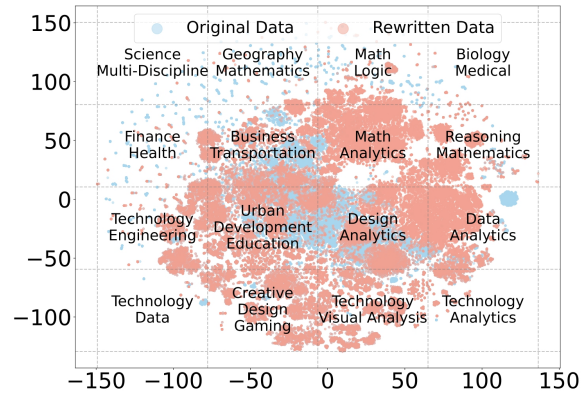


Figure 5: The t-SNE data distribution plot demonstrates how the rewritten data expands beyond the original dataset, increasing topic diversity and enhancing coverage of complex queries and reasoning.

Group B. These samples are evaluated using the InternVL2-Llama3-76B model, focusing on two key quality dimensions: 1) Information Content Score (1–5): Evaluates the depth and richness of the conversation. 2) Relevance Score (1–5): Measures the alignment between the visual and textual components. The quality score prompt is included in Appendix E.1.

The comparison presented in Figure 4 demonstrates that the rewritten dataset consistently surpasses the original dataset in both content and relevance scores. This indicates significant improvements in content quality and relevance after the rewriting process. In addition, Appendix C shows that there is no data contamination in the dataset.

### 3.2 Distribution Comparison

**Distribution of Data Length.** Figure 4 illustrates the token length (instruction + response) distribution of the original data and rewritten data. The results show that the rewritten data demonstrates a broader and more evenly distributed token length, with a longer tail for larger token counts. This suggests that the rewriting process tends to generate longer texts, likely incorporating additional details or rationales to enhance explanation and clarity.

**Distribution of Instructions.** To analyze the distributional differences between the original and rewritten data, we randomly sampled 80,000 examples from the dataset both before and after rewriting and visualized their distributions using t-SNE to project the instructions onto a two-dimensional plot Figure 5. The resulting figure reveals two key takeaways: (1) The rewritten data exhibits significant overlap with the original data, indicating that it retains the core characteristics of the original distribu-

tion. This ensures that the rewritten data preserves the foundational structure of the dataset. (2) The rewritten data extends beyond the boundaries of the original distribution, demonstrating that it introduces new dimensions or variations, which shows that rewriting enhances the dataset by broadening its scope while maintaining its original essence. Based on this observation, during the experimental validation phase, we utilize a mixed dataset consisting of 70% rewritten data and 30% original data to train the model.

### 3.3 Analysis of the Filtering Step

**Model-based vs Human-based Filtering.** To evaluate the reliability of our model-based filtering approach, we conduct a comparative analysis between InternVL2-Llama3-76B and human evaluators. We sample 60 items and make them classified as either "good" or "bad" by both the model and three independent human evaluators. To measure inter-rater agreement, we employ Cohen’s Kappa coefficient (McHugh, 2012), which accounts for chance agreement. The average Kappa value among the three human evaluators is 0.55. When substituting the model’s classifications for one evaluator and calculating the average Kappa with the remaining two human evaluators, we obtained a higher average value of 0.64. This falls within the range consider to indicate good consistency (0.60-0.80), suggesting that our model-based filtering approach achieves reliable agreement with human judgment. Detailed results are presented in Table A5.

**Category-specific Filtering Analysis.** Our filtering process reveals notable patterns across different data categories, as illustrated in Figure 6. Particularly high rates of hallucination are observed in visually complex categories such as OCR and Chart data, highlighting current MLLMs’ insufficient capabilities of understanding OCR and chart contents. The importance of our filtering approach is further demonstrated through the ablation study in Appendix subsection 5.1, which shows significant performance improvements in model training outcomes. We also show representative examples of correctly and incorrectly rewritten data in Appendix subsection F.1 and F.2.

## 4 Experiments

### 4.1 Model Training

To demonstrate the effectiveness of MAMmoTH-VL-12M, we train an MLLM following the archi-




Figure 6: The filter rates of different data types after filtering, with a lower filtering rate seen in categories like GeneralQA and Math, while OCR and Chart data experience more extensive filtering.

tecture of Llava-OneVision (Li et al., 2024b), comprising a language tower, a vision tower, and a projector. We use Qwen2.5-7B-Instruct (Team, 2024) as the LLM backbone, Siglip-so400m-patch14-384 (Zhai et al., 2023) as the vision tower, and a two-layer MLP as the projector. Following Llava-OneVision (Li et al., 2024b), we divide the training into three stages, as illustrated in Table 3, Appendix B.

- **Stage-1: Language-Image Alignment.** The goal is to align the visual features well into the word embedding space of LLMs, using the same pre-training corpus as LLaVA (Liu et al., 2024b).
- **Stage-2: Visual Instruction Tuning (Single Image, SI).** The model is first trained on 10M single-image instructions randomly sampled from MAMmoTH-VL-12M, resulting in a model with strong performance at following a diverse set of instructions to complete visual tasks using a single image.
- **Stage-3: Visual Instruction Tuning (One Vision, OV).** The model is then trained on a mixture of single-image, multi-image, and video data (2M). In this phase, the model expands its capabilities from single-image scenarios to diverse scenarios. The single-image dataset used in this stage is around 1M, which does not have any overlaps with Stage-2 training data.

### 4.2 Model Evaluation Setup

To reveal the generality and effectiveness of the model, we comprehensively evaluate it across different scenarios, including single-image, multi-image, and video benchmarks. Detailed results are presented in Table 1 and Table 2, respectively. We denote the model checkpoint that completed the Stage-2 and Stage-3 as MAMmoTH-VL-8B (SI) and MAMmoTH-VL-8B. We conduct standardized, reproducible evaluations of our model across all 23 benchmarks using LMMs-Eval (Zhang et al.,

Model	Multi-Discipline Knowledge and Mathematical Reasoning							
	MMStar	MMMU	MMMU-Pro	SeedBench	MMBench	MMVet	MathVerse	MathVista
	test	val	vision	test	en-test	test	mini-vision	testmini
GPT-4o (OpenAI, 2024)	64.7	69.1	49.7	76.2	82.1	76.2	50.2	63.8
Gemini-1.5-Pro (Gemini Team, 2023)	59.1	65.8	44.4	76.0	73.9	64.0	-	63.9
Claude-3.5-Sonnet (Anthropic, 2024)	62.2	68.3	48.0	72.2	79.7	75.4	-	67.7
InternVL2-76B (Chen et al., 2023b)	67.1	58.2	38.0	77.6	86.5	64.4	-	65.5
Qwen2-VL-72B-Ins (Wang et al., 2024c)	68.6	64.5	37.1	77.9	86.9	73.9	37.3	70.5
LLaVA-OV-72B (SI) (Li et al., 2024b)	65.2	57.4	26.0	77.6	86.6	60.0	37.7	66.5
LLaVA-OV-72B (Li et al., 2024b)	66.1	56.8	24.0	78.0	85.9	63.7	39.1	67.5
MiniCPM-V-2.6-8B (Yao et al., 2024)	57.5	49.8	21.7	74.0	81.5	60.0	-	60.6
INXComp-2.5-7B (Zhang et al., 2024b)	59.9	42.9	-	75.4	74.4	51.7	20.0	59.6
Llama-3.2-11B-Vision-Ins. (Meta, 2024b)	49.8	50.7	23.7	72.7	73.2	57.6	23.6	51.5
InternVL-2-8B (Chen et al., 2023b)	59.4	49.3	25.4	76.0	81.7	60.0	27.5	58.3
Qwen2-VL-7B-Ins. (Wang et al., 2024c)	60.7	52.1	26.9	74.3	83.0	62.0	28.2	58.2
Cambrian-1-8B (Tong et al., 2024)	-	42.7	14.7	73.3	74.6	48.0	-	49.0
Llava-CoT-11B (Xu et al., 2024b)	57.6	48.9	18.5	75.2	75.0	60.3	24.2	54.8
Molmo-8B-D (Deitke et al., 2024)	50.5	45.3	18.9	74.1	73.6	58.0	21.5	51.6
LLaVA-OV-7B (SI) (Li et al., 2024b)	60.9	47.3	16.8	74.8	80.5	58.8	26.9	56.1
LLaVA-OV-7B (Li et al., 2024b)	61.7	48.8	18.7	75.4	80.8	58.6	26.2	63.2
MAmmoTH-VL-8B (SI)	55.4	49.4	<b>26.0</b>	73.3	83.0	60.6	<b>35.0</b>	67.6
MAmmoTH-VL-8B	<b>63.0</b>	<b>50.8</b>	25.3	<b>76.0</b>	<b>83.4</b>	<b>62.3</b>	34.2	<b>67.6</b>
$\Delta$ Over Best Open-Source (~10B Scale)	+1.3	+1.9	+7.1	+0.6	+2.6	+2.0	+8.1	+4.4

Model	Chart & Doc Understanding				Multimodal Interactions & Preferences		
	AI2D	ChartQA	InfoVQA	DocVQA	RealWorldQA	WildVision	L-Wilder
	test	test	test	test	test	0617	small
GPT-4o (OpenAI, 2024)	94.2	85.7	79.2	92.8	76.5	89.4	85.9
Gemini-1.5-Pro (Gemini Team, 2023)	94.4	87.2	81.0	93.1	70.4	-	-
Claude-3.5-Sonnet (Anthropic, 2024)	94.7	90.8	49.7	95.2	60.1	50.0	83.1
InternVL2-76B (Chen et al., 2023b)	88.4	88.4	82.0	94.1	72.7	-	-
Qwen2-VL-72B-Ins (Wang et al., 2024c)	88.1	88.3	84.5	96.5	77.8	52.3	53.6
LLaVA-OV-72B (SI) (Li et al., 2024b)	85.1	84.9	74.6	91.8	73.8	49.5	72.9
LLaVA-OV-72B (Li et al., 2024b)	85.6	83.7	74.9	91.3	71.9	52.3	72.0
MiniCPM-V-2.6-7B (Yao et al., 2024)	82.1	82.4	-	90.8	65.0	11.7	-
INXComp-2.5-7B (Zhang et al., 2024b)	81.5	82.2	70.0	90.9	67.8	-	61.4
Llama-3.2-11B-Vision-Ins (Meta, 2024b)	77.3	83.4	65.0	88.4	63.3	49.7	62.0
InternVL-2-8B (Chen et al., 2023b)	83.8	83.3	74.8	91.6	64.4	51.5	62.5
Qwen2-VL-7B-Ins (Wang et al., 2024c)	83.0	83.0	76.5	94.5	70.1	44.0	66.3
Cambrian-1-8B (Tong et al., 2024)	73.3	73.3	41.6	77.8	64.2	-	34.1
Llava-CoT-11B (Xu et al., 2024b)	-	67.0	44.8	-	-	-	65.3
Molmo-7B-D (Deitke et al., 2024)	81.0	84.1	72.6	92.2	70.7	40.0	-
LLaVA-OV-7B (SI) (Li et al., 2024b)	81.6	78.8	65.3	86.9	65.5	39.2	69.1
LLaVA-OV-7B (Li et al., 2024b)	81.4	80.0	68.8	87.5	66.3	<b>53.8</b>	67.8
MAmmoTH-VL-8B (SI)	83.4	85.9	<b>74.8</b>	<b>93.8</b>	<b>71.3</b>	51.9	<b>71.3</b>
MAmmoTH-VL-8B	<b>84.0</b>	<b>86.2</b>	73.1	93.7	69.9	51.1	70.8
$\Delta$ Over Best Open-Source (~10B Scale)	+2.4	+2.1	+2.2	+1.6	+0.6	-1.9	+2.2

Table 1: Performance on multi-discipline knowledge, mathematical reasoning, chart, diagram, and document understanding, and real-world multimodal interactions and human preferences benchmarks. We highlight different groups of models with different colors: closed-source models, open weights but closed training details, and fully open-source models. Results are from official sources or running with lmms-eval package if unavailable.

2024a). To ensure a fair comparison with other MLLMs, we primarily report results from the original papers. When results are unavailable, we onboard the models in LMMS-Eval and evaluate them using consistent settings. All results are reported using greedy decoding and zero-shot settings unless specified.

### 4.3 Single-Image Performance

**Evaluation Benchmarks.** To validate the performance for single-image tasks in real-world scenarios, we consider a comprehensive set of benchmarks in Table 1. It can be categorized into three classes:

*Chart, Diagram, and Document Understanding:* AI2D (Kembhavi et al., 2016), ChartQA (Masry

Model	Multi-Image and Video							
	MuirBench	MEGABench	EgoSchema	PerceptionTest	SeedBench	MLVU	MVBench	VideoMME
	test	test	test	test	video	dev	test	w/o subs
GPT-4o (OpenAI, 2024)	68.0	54.2	-	-	-	64.6	-	71.9
GPT-4V (OpenAI, 2023)	62.3	-	-	-	60.5	49.2	43.5	59.9
LLaVA-OV-72B (SI) (Li et al., 2024b)	33.2	-	58.6	62.3	60.9	60.9	57.1	64.8
LLaVA-OV-72B (Li et al., 2024b)	54.8	33.8	62.0	66.9	62.1	66.4	59.4	66.2
InternVL-2-8B (Chen et al., 2023b)	59.4	27.7	54.2	57.4	54.9	30.2	66.4	54.0
Qwen2-VL-7B-Ins. (Wang et al., 2024c)	41.6	36.0	66.7	62.3	55.3	58.6	67.0	63.3
LLaVA-OV-7B (SI) (Li et al., 2024b)	32.7	22.1	52.9	54.9	51.1	60.2	51.2	55.0
LLaVA-OV-7B (Li et al., 2024b)	41.8	23.9	<b>60.1</b>	57.1	56.9	64.7	56.7	58.2
MAMmoTH-VL-8B	<b>55.1</b>	<b>28.2</b>	58.5	<b>59.3</b>	<b>57.1</b>	<b>64.7</b>	<b>59.1</b>	<b>58.8</b>
$\Delta$ Over Best Open-Source $\sim$ 10B Scale)	+13.3	+4.3	-1.6	+2.2	+0.2	+0	+2.4	+0.6

Table 2: Main results on Multi-Image and Video benchmarks. Follow the same settings as in Table 1.

et al., 2022a), DocVQA (Mathew et al., 2020), InfoVQA (Mathew et al., 2021a).

*Multi-discipline Knowledge and Mathematical Reasoning:* MMStar (Chen et al., 2024c), MMMU (Yue et al., 2024a), MMMU-Pro (Yue et al., 2024b), SeedBench (Li et al., 2023b), MM-Bench (Liu et al., 2025), MMvet (Yu et al., 2023), Mathverse (Zhang et al., 2024c), Mathvista (Lu et al., 2023a).

*Real-world Multimodal Interactions and Human Preferences:* RealworldQA (x.ai, 2024), WildVision (Lu et al., 2024), Llava-Wilder-Small (Li et al., 2024a).

**Results Analysis.** (1) The experimental results show that MAMmoTH-VL-8B achieves state-of-the-art performance among open-source multimodal models across diverse benchmarks, approaching the performance of leading open-weight models. (2) Specifically, MAMmoTH-VL-8B achieves state-of-the-art performance on 10 benchmarks among both open-source and open-weight models particularly in mathematical reasoning. (3) In comparing MAMmoTH-VL-8B with its single-image variant MAMmoTH-VL-8B (SI), we observe that despite minor performance decreases on some benchmarks due to the introduction of multi-image and video, the overall performance remains robust.

#### 4.4 Multi-Image and Video Performance

**Evaluation Benchmarks.** After completing Stage 3 of training, MAMmoTH-VL-8B has been trained to process both multi-image and video data. To evaluate its performance, we assess it on eight benchmarks: MuirBench (Wang et al., 2024a), Megabench (Chen et al., 2024b), EgoSchema (Mangalam et al., 2023), PerceptionTest (Pătrăucean et al., 2023), SeedBench(Video)(Li et al., 2023b), MLVU (Zhou et al., 2024), MVBench (Li et al.,

2023d), and VideoMME (Fu et al., 2024).

**Results Analysis.** The experimental results demonstrate that MAMmoTH-VL-8B significantly advances the state of the art among open-source models, surpassing the previous leader LLaVA-One-Vision-7B, with a particularly striking 13-point improvement on MuirBench. While MAMmoTH-VL-8B shows promising results, a performance gap persists relative to Qwen2-VL-7B. This gap likely stems from our current data limitations - specifically, our training dataset includes only 1M samples of multi-image and video content due to computational constraints. This suggests substantial headroom for performance gains through expanded training data coverage of these modalities.

## 5 Ablation Study

In this section, we employ Qwen2.5-1.5B-Instruct as the backbone model for conducting ablation studies. These ablation studies evaluate the impact of key factors in our dataset creation and training process. We examine the effects of data filtering, data mixing ratio, training data scale, and the impact of rewrite model size, providing insights into the most important elements for optimizing multimodal learning. The following subsections detail these experiments.

### 5.1 Effect of Data Filtering

To verify the necessity of filtering, we randomly sample 1M data from both the dataset before filtering and the dataset after filtering and conduct experiments on the same model. The experimental results in Figure 8 demonstrate that, the model trained on the filtered dataset outperforms that trained on the no-filtering dataset, particularly in the benchmark for Chart, Diagram, and Document Understanding. This aligns with our previous ob-




Figure 7: Performance of data rewritten by different models on three benchmark subsets

servation that OCR and Chart data constitute the majority of the data filtered. This finding effectively validates the necessity and significance of hallucination mitigation. It also indirectly highlights the persistent issue of hallucinations in current MLLMs. See Table A1 for more details.

## 5.2 Effect of Data Mixing Ratio

To validate the effectiveness of merging original and rewritten datasets, we train five models using different combinations of original and rewritten data: one using only the original dataset, another using only the rewritten dataset, and three using combined datasets with original:rewritten ratios of 3:7, 7:3, and 5:5 respectively. As shown in Figure A2, the model trained exclusively on rewritten data demonstrated superior average performance compared to the one trained on original data alone. Furthermore, the model trained on the combined dataset with a 3:7 ratio of original to rewritten data achieved marginally better performance than the rewritten-only model. These results suggest that merging the datasets is beneficial, likely due to the increased diversity and comprehensiveness of the training data. Detailed performance metrics can be found in Table A3.

## 5.3 Effect of Training Data Scale

As shown in Figure 1 and Figure A1, performance is tracked as training dataset size increases in 2M-sample intervals, compared against LLaVA-




Figure 8: Data filtering significantly improves the quality of generated data, particularly in chart and document understanding, where hallucinations are more frequent.

OneVision-72B, Llava-CoT-11B, and LLaVA-OneVision-7B. Results confirm a positive correlation between data scale and performance, highlighting the benefits of diverse instruction data for handling complex tasks.

However, MMStar and SeedBench performance declines in late stage-2 training, mainly due to drops in Scene Understanding and Instance Attributes subsets. We hypothesize that as models improve at processing complex images and questions, they may prioritize object relationships over fine-grained attribute recognition.

## 5.4 Impact of Rewrite Model Size

To assess the effect of model size on the quality of rewritten data, we conduct experiments using four models trained on a dataset of 500K samples. The first model uses the original dataset. The second is trained on data rewritten by InternVL2-Llama3-76B and Meta-Llama-3-70B-Ins. The third uses rewrites from Qwen2-VL-7B-Ins and Qwen2.5-7B-Ins, while the fourth relies on InternVL2-8B and InternLM2.5-7B.

As shown in Figure 7, performance patterns vary by task type. For knowledge & reasoning, models trained on 7B-parameter rewrites perform similarly to those using larger model rewrites. However, in chart and document tasks, smaller models degrade performance, while larger models offer modest gains, indicating their superior visual understanding. Multi Interact & Preference tasks show a direct correlation with model scale, as larger models excel in nuanced preference modeling. Detailed metrics are provided in Table A4.

## 6 Conclusion

We present an efficient method for enhancing MLLMs using open-source models to generate diverse multimodal training data aligned with human preferences. Our MAMmoTH-VL-Instruct dataset (12M multimodal entries) enables the MAMmoTH-VL-8B architecture to achieve state-of-the-art per-



formance across tasks while minimizing dependency on proprietary models. Experiments demonstrate significant improvements in handling real-world scenarios, establishing a scalable framework for advancing multimodal understanding. This approach offers new directions for extending methodology to broader modalities and datasets.

## Ethical Statement

Our study adheres to ethical principles by ensuring that all data used in constructing our multimodal instruction-tuning dataset originate from publicly available and open-access sources, avoiding proprietary, private, or sensitive information. To mitigate biases, we employed rewriting and self-filtering techniques, though we acknowledge that residual biases may persist, particularly in complex reasoning tasks. Our dataset and trained models are intended for research purposes, and while they demonstrate improved reasoning capabilities, they are not guaranteed to be free of factual inaccuracies and should not be deployed in high-stakes applications without further validation. To promote transparency and reproducibility, we release our dataset, code, and model configurations under an open-source license, encouraging responsible use and community contributions. While our work advances multimodal reasoning, it does not eliminate all limitations, and we emphasize the need for continued research to enhance faithfulness, interpretability, and robustness. Given the broader implications of AI-generated content, we support further investigations into mitigating potential risks, including misinformation and unintended societal impacts, through improved dataset curation and model evaluation.

## Limitations

Despite promising results, several limitations remain. Relying on open models to generate instruction-response pairs may introduce bias and hinder generalization, particularly for complex tasks. We tried to mitigate this by selecting models known for strong performance across diverse domains. While the dataset covers a broad range of reasoning-intensive tasks, the quality of generated rationales varies. Future work may further strengthen the automatic filtering. Lastly, despite its substantial size (12M pairs), the dataset may still overlook certain multimodal nuances, which we continuously refine to better reflect real-world

variability.

## Acknowledgments

The authors would like to thank the invaluable comments from the reviewers. This work was supported in part by a Carnegie Bosch Institute Fellowship to Xiang Yue.

## References

- Manoj Acharya, Kushal Kafle, and Christopher Kanan. 2019. *Tallyqa: Answering complex counting questions*. In *AAAI*.
- Anthropic. 2024. *Claude 3.5 sonnet*. Accessed on February 11, 2024.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. *VQA: visual question answering*. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pages 2425–2433.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. *Qwen technical report*. *ArXiv preprint*, abs/2309.16609.
- Tianyi Bai, Hao Liang, Binwang Wan, Yanran Xu, Xi Li, Shiyu Li, Ling Yang, Bozhou Li, Yifan Wang, Bin Cui, et al. 2024. *A survey of multimodal large language model from a data-centric perspective*. *ArXiv preprint*, abs/2405.16640.
- Ali Furkan Biten, Rubèn Tito, Andrés Mafla, Lluís Gómez i Bigorda, Marçal Rusiñol, C. V. Jawahar, Ernest Valveny, and Dimosthenis Karatzas. 2019. *Scene text visual question answering*. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 4290–4300. IEEE.
- Sungguk Cha, Jusung Lee, Younghyun Lee, and Cheoljong Yang. 2024. *Visually dehallucinative instruction generation*. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5510–5514. IEEE.
- Shuaichen Chang, David Palzer, Jialin Li, Eric Fosler-Lussier, and Ningchuan Xiao. 2022. *Mapqa: A dataset for question answering on choropleth maps*. In *NeurIPS 2022 First Table Representation Workshop*.
- Soravit Changpinyo, Doron Kukliansy, Idan Szepes, Xi Chen, Nan Ding, and Radu Soricut. 2022. *All you may need for VQA are image captions*. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1947–1963, Seattle, United States. Association for Computational Linguistics.

- Guiming Hardy Chen, Shunian Chen, Ruifei Zhang, Junying Chen, Xiangbo Wu, Zhiyi Zhang, Zhihong Chen, Jianquan Li, Xiang Wan, and Benyou Wang. 2024a. [Allava: Harnessing gpt4v-synthesized data for a lite vision-language model](#). *ArXiv preprint*, abs/2402.11684.
- Jiacheng Chen, Tianhao Liang, Sherman Siu, Zhengqing Wang, Kai Wang, Yubo Wang, Yuan-sheng Ni, Wang Zhu, Ziyang Jiang, Bohan Lyu, Dongfu Jiang, Xuan He, Yuan Liu, Hexiang Hu, Xiang Yue, and Wenhui Chen. 2024b. [Mega-bench: Scaling multimodal evaluation to over 500 real-world tasks](#). *ArXiv preprint*, abs/2410.10563.
- Jiaqi Chen, Tong Li, Jinghui Qin, Pan Lu, Liang Lin, Chongyu Chen, and Xiaodan Liang. 2022. [UniGeo: Unifying geometry logical reasoning via reformulating mathematical expression](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3313–3323, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jiaqi Chen, Jianheng Tang, Jinghui Qin, Xiaodan Liang, Lingbo Liu, Eric Xing, and Liang Lin. 2021a. [GeoQA: A geometric question answering benchmark towards multimodal numerical reasoning](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 513–523, Online. Association for Computational Linguistics.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. 2023a. [Sharegpt4v: Improving large multimodal models with better captions](#). *ArXiv preprint*, abs/2311.12793.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. 2025. [Sharegpt4v: Improving large multi-modal models with better captions](#). In *European Conference on Computer Vision*, pages 370–387. Springer.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. 2024c. [Are we on the right way for evaluating large vision-language models?](#) *ArXiv preprint*, abs/2403.20330.
- Zhe Chen, Jiannan Wu, Wenhui 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. 2023b. [Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks](#). *ArXiv preprint*, abs/2312.14238.
- Zhiyu Chen, Wenhui Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. 2021b. [FinQA: A dataset of numerical reasoning over financial data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3697–3711, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zhoujun Cheng, Haoyu Dong, Zhiruo Wang, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou, and Dongmei Zhang. 2022. [HiTab: A hierarchical table dataset for question answering and natural language generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1094–1110, Dublin, Ireland. Association for Computational Linguistics.
- Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%\\* chatgpt quality](#).
- Wenliang Dai, Junnan Li, DONGXU LI, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. 2023. [Instructblip: Towards general-purpose vision-language models with instruction tuning](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 49250–49267. Curran Associates, Inc.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. 2024. [Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models](#). *ArXiv preprint*, abs/2409.17146.
- Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. 2024. [Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis](#). *ArXiv preprint*, abs/2405.21075.
- Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wan-jun Zhong, Yufei Wang, Lanqing Hong, Jianhua Han, Hang Xu, Zhenguo Li, and Lingpeng Kong. 2023. [G-llava: Solving geometric problem with multi-modal large language model](#). *ArXiv preprint*, abs/2312.11370.
- Gemini Team. 2023. [Gemini: a family of highly capable multimodal models](#). *ArXiv preprint*, abs/2312.11805.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. [Making the v in vqa matter: Elevating the role of image understanding in visual question answering](#). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913.
- Tanmay Gupta, Dustin Schwenk, Ali Farhadi, Derek Hoiem, and Aniruddha Kembhavi. 2018. [Imagine this! scripts to compositions to videos](#). In *Proceedings of the European conference on computer vision (ECCV)*, pages 598–613.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P.

- Bigham. 2018. [Vizwiz grand challenge: Answering visual questions from blind people](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3608–3617. IEEE Computer Society.
- Xuehai He, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. 2020. [Pathvqa: 30000+ questions for medical visual question answering](#). *ArXiv preprint*, abs/2003.10286.
- Jack Hessel, Jena D Hwang, Jae Sung Park, Rowan Zellers, Chandra Bhagavatula, Anna Rohrbach, Kate Saenko, and Yejin Choi. 2022. [The abduction of sherlock holmes: A dataset for visual abductive reasoning](#). In *European Conference on Computer Vision*, pages 558–575. Springer.
- Yu-Chung Hsiao, Fedir Zubach, Gilles Baechler, Victor Carbune, Jason Lin, Maria Wang, Srinivas Sunkara, Yun Zhu, and Jindong Chen. 2022. [Screenqa: Large-scale question-answer pairs over mobile app screenshots](#). *ArXiv preprint*, abs/2209.08199.
- Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang Zhang, Bo Zhang, Ji Zhang, Qin Jin, Fei Huang, and Jingren Zhou. 2024. [mPLUG-DocOwl 1.5: Unified structure learning for OCR-free document understanding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3096–3120, Miami, Florida, USA. Association for Computational Linguistics.
- Ting-Hao K. Huang, Francis Ferraro, Nasrin Mostafazadeh, Ishan Misra, Jacob Devlin, Aishwarya Agrawal, Ross Girshick, Xiaodong He, Pushmeet Kohli, Dhruv Batra, et al. 2016. [Visual storytelling](#). In *15th Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2016)*.
- Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and CV Jawahar. 2019. [Icdar2019 competition on scanned receipt ocr and information extraction](#). In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1516–1520. IEEE.
- Drew A Hudson and Christopher D Manning. 2019. [Gqa: A new dataset for real-world visual reasoning and compositional question answering](#). In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709.
- Mude Hui, Siwei Yang, Bingchen Zhao, Yichun Shi, Heng Wang, Peng Wang, Yuyin Zhou, and Cihang Xie. 2024. [Hq-edit: A high-quality dataset for instruction-based image editing](#). *ArXiv preprint*, abs/2404.09990.
- Dongfu Jiang, Xuan He, Huaye Zeng, Cong Wei, Max W.F. Ku, Qian Liu, and Wenhui Chen. 2024. [Mantis: Interleaved multi-image instruction tuning](#). *Transactions on Machine Learning Research*, 2024.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. 2017. [Clevr: A diagnostic dataset for compositional language and elementary visual reasoning](#). In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910.
- Kushal Kafle, Brian L. Price, Scott Cohen, and Christopher Kanan. 2018. [DVQA: understanding data visualizations via question answering](#). In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 5648–5656. IEEE Computer Society.
- Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Akos Kadar, Adam Trischler, and Yoshua Bengio. 2018. [Figureqa: An annotated figure dataset for visual reasoning](#). *Preprint*, arXiv:1710.07300.
- Mehran Kazemi, Hamidreza Alvari, Ankit Anand, Jialin Wu, Xi Chen, and Radu Soricut. 2024. [Geomverse: A systematic evaluation of large models for geometric reasoning](#). In *AI for Math Workshop@ ICML 2024*.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. [A diagram is worth a dozen images](#). In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 235–251. Springer.
- Aniruddha Kembhavi, Min Joon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. 2017. [Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 5376–5384. IEEE Computer Society.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. 2020. [The hateful memes challenge: Detecting hate speech in multimodal memes](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). *ArXiv preprint*, abs/2205.11916.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, Michael S Bernstein, and Li Fei-Fei. 2017. [Visual genome: Connecting language and vision using crowdsourced dense image annotations](#). *International Journal of Computer Vision*, 123(1):32–73.

- Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. 2018. [A dataset of clinically generated visual questions and answers about radiology images](#). *Scientific data*, 5(1):1–10.
- Hugo Laurençon, Léo Tronchon, and Victor Sanh. 2024. [Unlocking the conversion of web screenshots into html code with the websight dataset](#). *ArXiv preprint*, abs/2403.09029.
- Paul Lerner, Olivier Ferret, Camille Guinaudeau, Hervé Le Borgne, Romaric Besançon, Jose G. Moreno, and Jesus Lovon. 2022. [ViQuAE, a Dataset for Knowledge-based Visual Question Answering about Named Entities](#). In *SIGIR '22: The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Madrid, Spain.
- Bo Li, Kaichen Zhang, Hao Zhang, Dong Guo, Renrui Zhang, Feng Li, Yuanhan Zhang, Ziwei Liu, and Chunyuan Li. 2024a. [Llava-next: Stronger llms supercharge multimodal capabilities in the wild](#).
- Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Fanyi Pu, Jingkang Yang, Chunyuan Li, and Ziwei Liu. 2023a. [Mimic-it: Multi-modal in-context instruction tuning](#). *ArXiv preprint*, abs/2306.05425.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024b. [Llava-onevision: Easy visual task transfer](#). *ArXiv preprint*, abs/2408.03326.
- Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. 2023b. [Seed-bench: Benchmarking multimodal llms with generative comprehension](#). *ArXiv preprint*, abs/2307.16125.
- Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. 2024c. [Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models](#). *ArXiv preprint*, abs/2407.07895.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023c. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *International conference on machine learning*, pages 19730–19742. PMLR.
- Kunchang Li, Yali Wang, Yanan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. 2023d. [Mvbench: A comprehensive multi-modal video understanding benchmark](#). *ArXiv preprint*, abs/2311.17005.
- Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiaochong Feng, Lingpeng Kong, and Qi Liu. 2024d. [Multimodal ArXiv: A dataset for improving scientific comprehension of large vision-language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14369–14387, Bangkok, Thailand. Association for Computational Linguistics.
- Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang, Jingjing Xu, Xu Sun, Lingpeng Kong, and Qi Liu. 2023e. [M<sup>3</sup>it: A large-scale dataset towards multi-modal multilingual instruction tuning](#). *ArXiv preprint*, abs/2306.04387.
- Yitong Li, Zhe Gan, Yelong Shen, Jingjing Liu, Yu Cheng, Yuexin Wu, Lawrence Carin, David E. Carlson, and Jianfeng Gao. 2019. [Storygan: A sequential conditional GAN for story visualization](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 6329–6338. Computer Vision Foundation / IEEE.
- Zhuowan Li, Xingrui Wang, Elias Stengel-Eskin, Adam Kortylewski, Wufei Ma, Benjamin Van Durme, and Alan L Yuille. 2023f. [Super-clevr: A virtual benchmark to diagnose domain robustness in visual reasoning](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14963–14973.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. 2015. [Microsoft coco: Common objects in context](#). *Preprint*, arXiv:1405.0312.
- Adam Dahlgren Lindström. 2022. [Clevr-math: A dataset for compositional language, visual, and mathematical reasoning](#). *ArXiv preprint*, abs/2208.05358.
- Fangyu Liu, Guy Edward Toh Emerson, and Nigel Collier. 2023a. [Visual spatial reasoning](#). *Transactions of the Association for Computational Linguistics*.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2023b. [Aligning large multi-modal model with robust instruction tuning](#). *ArXiv preprint*, abs/2306.14565.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2024a. [Mitigating hallucination in large multi-modal models via robust instruction tuning](#). *ArXiv preprint*, abs/2306.14565.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024b. [Visual instruction tuning](#). *Advances in neural information processing systems*, 36.
- Junpeng Liu, Tianyue Ou, Yifan Song, Yuxiao Qu, Wai Lam, Chenyan Xiong, Wenhui Chen, Graham Neubig, and Xiang Yue. 2024c. [Harnessing webpage uis for text-rich visual understanding](#). *ArXiv preprint*, abs/2410.13824.
- Xi Liu, Rui Zhang, Yongsheng Zhou, Qianyi Jiang, Qi Song, Nan Li, Kai Zhou, Lei Wang, Dong Wang, Minghui Liao, et al. 2019. [Icdar 2019 robust reading challenge on reading chinese text on signboard](#). *ArXiv preprint*, abs/1912.09641.
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2024d. [Gpt understands, too](#). *ArXiv preprint*, abs/2103.10385.

- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2025. [Mm-bench: Is your multi-modal model an all-around player?](#) In *European Conference on Computer Vision*, pages 216–233. Springer.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2023a. [Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts](#). *ArXiv preprint*, abs/2310.02255.
- Pan Lu, Ran Gong, Shibiao Jiang, Liang Qiu, Siyuan Huang, Xiaodan Liang, and Song-Chun Zhu. 2021a. [Inter-GPS: Interpretable geometry problem solving with formal language and symbolic reasoning](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6774–6786, Online. Association for Computational Linguistics.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. [Learn to explain: Multimodal reasoning via thought chains for science question answering](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 2507–2521. Curran Associates, Inc.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. 2023b. [Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning](#). In *International Conference on Learning Representations (ICLR)*.
- Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and Song-Chun Zhu. 2021b. [Iconqa: A new benchmark for abstract diagram understanding and visual language reasoning](#). In *The 35th Conference on Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks*.
- Yujie Lu, Dongfu Jiang, Wenhui Chen, William Yang Wang, Yejin Choi, and Bill Yuchen Lin. 2024. [Wildvision: Evaluating vision-language models in the wild with human preferences](#). *ArXiv preprint*, abs/2406.11069.
- Run Luo, Haonan Zhang, Longze Chen, Ting-En Lin, Xiong Liu, Yuchuan Wu, Min Yang, Minzheng Wang, Pengpeng Zeng, Lianli Gao, et al. 2024. [Mmevol: Empowering multimodal large language models with evol-instruct](#). *ArXiv preprint*, abs/2409.05840.
- Chenyang Lyu, Minghao Wu, Longyue Wang, Xinting Huang, Bingshuai Liu, Zefeng Du, Shuming Shi, and Zhaopeng Tu. 2023. [Macaw-llm: Multi-modal language modeling with image, audio, video, and text integration](#). *ArXiv preprint*, abs/2306.09093.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. 2024. [Video-ChatGPT: Towards detailed video understanding via large vision and language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12585–12602, Bangkok, Thailand. Association for Computational Linguistics.
- Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. 2023. [Egoschema: A diagnostic benchmark for very long-form video language understanding](#). *Advances in Neural Information Processing Systems*, 36:46212–46244.
- Urs-Viktor Marti and Horst Bunke. 2002. [The iam-database: an english sentence database for offline handwriting recognition](#). *International Journal on Document Analysis and Recognition*, 5:39–46.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022a. [ChartQA: A benchmark for question answering about charts with visual and logical reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279, Dublin, Ireland. Association for Computational Linguistics.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022b. [ChartQA: A benchmark for question answering about charts with visual and logical reasoning](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279, Dublin, Ireland. Association for Computational Linguistics.
- Minesh Mathew, Viraj Bagal, Rubèn Pérez Tito, Dimosthenis Karatzas, Ernest Valveny, and C. V. Jawahar. 2021a. [Infographicvqa](#). *ArXiv preprint*, abs/2104.12756.
- Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. 2021b. [Docvqa: A dataset for vqa on document images](#). In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2199–2208.
- Minesh Mathew, Dimosthenis Karatzas, R Manmatha, and CV Jawahar. 2020. [Docvqa: A dataset for vqa on document images](#). *corr abs/2007.00398 (2020)*. *ArXiv preprint*, abs/2007.00398.
- Mary L McHugh. 2012. [Interrater reliability: the kappa statistic](#). *Biochemia medica*, 22(3):276–282.
- Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruvi Shah, Xianzhi Du, Futang Peng, Floris Weers, et al. 2024. [Mm1: Methods, analysis & insights from multimodal llm pre-training](#). *ArXiv preprint*, abs/2403.09611.
- Meta. 2024a. [The llama 3 herd of models](#). *ArXiv preprint*, abs/2407.21783.
- Meta. 2024b. [Llama 3.2: Pushing the boundaries of vision and language for edge and mobile devices](#). Accessed: 2024-12-06.

- Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. 2020. [Plotqa: Reasoning over scientific plots](#). In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1527–1536.
- Anand Mishra, Karteek Alahari, and C. V. Jawahar. 2012. [Scene text recognition using higher order language priors](#). In *British Machine Vision Conference, BMVC 2012, Surrey, UK, September 3-7, 2012*, pages 1–11. BMVA Press.
- Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. 2019. [Ocr-vqa: Visual question answering by reading text in images](#). In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 947–952.
- Jason Obeid and Enamul Hoque. 2020. [Chart-to-text: Generating natural language descriptions for charts by adapting the transformer model](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 138–147, Dublin, Ireland. Association for Computational Linguistics.
- OpenAI. 2023. Gpt-4v. <https://openai.com/index/gpt-4v-system-card/>. Accessed: 2023-02-09, 2023-02-11, 2023-02-12.
- OpenAI. 2024. Hello gpt-4o. <https://openai.com/index/hello-gpt-4o/>. Accessed: 2024-02-09, 2024-02-11, 2024-02-12.
- OpenBMB. 2023. LLaVA-zh: Chinese language and vision assistant dataset. [https://huggingface.co/datasets/openbmb/llava\\_zh](https://huggingface.co/datasets/openbmb/llava_zh).
- Renjie Pi, Jiahui Gao, Shizhe Diao, Rui Pan, Hanze Dong, Jipeng Zhang, Lewei Yao, Jianhua Han, Hang Xu, Lingpeng Kong, and Tong Zhang. 2023. [Detgpt: Detect what you need via reasoning](#).
- Renjie Pi, Jianshu Zhang, Jipeng Zhang, Rui Pan, Zhekai Chen, and Tong Zhang. 2024. [Image textualization: An automatic framework for creating accurate and detailed image descriptions](#).
- Jordi Pont-Tuset, Jasper Uijlings, Soravit Changpinyo, Radu Soricut, and Vittorio Ferrari. 2020. [Connecting vision and language with localized narratives](#). In *ECCV*.
- Viorica Pătrăucean, Lucas Smaira, Ankush Gupta, Adrià Recasens Contente, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Joseph Heyward, Mateusz Malinowski, Yi Yang, Carl Doersch, Tatiana Matejovicova, Yury Sulsky, Antoine Miech, Alex Frechette, Hanna Klimczak, Raphael Koster, Junlin Zhang, Stephanie Winkler, Yusuf Aytar, Simon Osindero, Dima Damen, Andrew Zisserman, and João Carreira. 2023. [Perception test: A diagnostic benchmark for multimodal video models](#). *ArXiv preprint*, abs/2305.13786.
- 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. 2021. [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*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Mengye Ren, Ryan Kiros, and Richard Zemel. 2015. [Exploring models and data for image question answering](#). *Advances in neural information processing systems*, 28.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2015. [Imagenet large scale visual recognition challenge](#).
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. 2022. [A-okvqa: A benchmark for visual question answering using world knowledge](#). *ArXiv preprint*, abs/2206.01718.
- Minjoon Seo, Hannaneh Hajishirzi, Ali Farhadi, Oren Etzioni, and Clint Malcolm. 2015. [Solving geometry problems: Combining text and diagram interpretation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1466–1476, Lisbon, Portugal. Association for Computational Linguistics.
- Shanghai AI Laboratory. 2023. [Sharegpt-4o: Comprehensive multimodal annotations with GPT-4O](#). Technical report, Shanghai AI Laboratory.
- Baoguang Shi, Cong Yao, Minghui Liao, Mingkun Yang, Pei Xu, Linyan Cui, Serge Belongie, Shijian Lu, and Xiang Bai. 2017. [Icdar2017 competition on reading chinese text in the wild \(rctw-17\)](#). *ArXiv preprint*, abs/1708.09585.
- Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. 2024. [Design2code: How far are we from automating front-end engineering?](#) *ArXiv preprint*, abs/2403.03163.
- Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. 2020. [Textcaps: a dataset for image captioning with reading comprehension](#). In *ECCV 2020 (oral)*.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. [Towards VQA models that can read](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 8317–8326. Computer Vision Foundation / IEEE.

- Amanpreet Singh, Guan Pang, Mandy Toh, Jing Huang, Wojciech Galuba, and Tal Hassner. 2021. [Textocr: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 8802–8812. Computer Vision Foundation / IEEE.
- Krishna Srinivasan, Karthik Raman, Jiecao Chen, Michael Bendersky, and Marc Najork. 2021. [Wit: Wikipedia-based image text dataset for multimodal multilingual machine learning](#). *ArXiv preprint, abs/2103.01913*.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2019. [A corpus for reasoning about natural language grounded in photographs](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6418–6428, Florence, Italy. Association for Computational Linguistics.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. 2023. [Aligning large multimodal models with factually augmented rlhf](#). *ArXiv preprint, abs/2309.14525*.
- Ryota Tanaka, Kyosuke Nishida, and Sen Yoshida. 2021. [Visualmrc: Machine reading comprehension on document images](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13878–13888.
- Benny J. Tang, Angie Boggust, and Arvind Satyanarayan. 2023. [VisText: A Benchmark for Semantically Rich Chart Captioning](#). In *The Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Qwen Team. 2024. [Qwen2.5: A party of foundation models](#).
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. 2024. [Cambrian-1: A fully open, vision-centric exploration of multimodal llms](#). *ArXiv preprint, abs/2406.16860*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. [Llama: Open and efficient foundation language models](#). *ArXiv preprint, abs/2302.13971*.
- Haoqin Tu, Chenhang Cui, Zijun Wang, Yiyang Zhou, Bingchen Zhao, Junlin Han, Wangchunshu Zhou, Huaxiu Yao, and Cihang Xie. 2023. [How many unicorns are in this image? a safety evaluation benchmark for vision llms](#). *ArXiv preprint, abs/2311.16101*.
- Bryan Wang, Gang Li, Xin Zhou, Zhourong Chen, Tovi Grossman, and Yang Li. 2021. [Screen2words: Automatic mobile ui summarization with multimodal learning](#). In *The 34th Annual ACM Symposium on User Interface Software and Technology, UIST '21*, page 498–510, New York, NY, USA. Association for Computing Machinery.
- Fei Wang, Xingyu Fu, James Y. Huang, Zekun Li, Qin Liu, Xiaogeng Liu, Mingyu Derek Ma, Nan Xu, Wenxuan Zhou, Kai Zhang, Tianyi Lorena Yan, Wenjie Jacky Mo, Hsiang-Hui Liu, Pan Lu, Chunyuan Li, Chaowei Xiao, Kai-Wei Chang, Dan Roth, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2024a. [Muirbench: A comprehensive benchmark for robust multi-image understanding](#). *ArXiv preprint, abs/2406.09411*.
- Junke Wang, Lingchen Meng, Zejia Weng, Bo He, Zuxuan Wu, and Yu-Gang Jiang. 2023. [To see is to believe: Prompting gpt-4v for better visual instruction tuning](#). *ArXiv preprint, abs/2311.07574*.
- Ke Wang, Junting Pan, Weikang Shi, Zimu Lu, Mingjie Zhan, and Hongsheng Li. 2024b. [Measuring multimodal mathematical reasoning with math-vision dataset](#). *ArXiv preprint, abs/2402.14804*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024c. [Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution](#). *ArXiv preprint, abs/2409.12191*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Christopher Wendler. 2024. [Renderedtext: https://huggingface.co/datasets/wendlerc/renderedtext](https://huggingface.co/datasets/wendlerc/renderedtext). Accessed: 2024-12-04.
- Haoning Wu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Annan Wang, Kaixin Xu, Chunyi Li, Jingwen Hou, Guangtao Zhai, Geng Xue, Wenxiu Sun, Qiong Yan, and Weisi Lin. 2023a. [Q-instruct: Improving low-level visual abilities for multi-modality foundation models](#). *ArXiv preprint, abs/2311.06783*.
- Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, and Weisi Lin. 2023b. [Q-align: Teaching llms for visual scoring via discrete text-defined levels](#). *ArXiv preprint, abs/2312.17090*.
- Haoning Wu, Hanwei Zhu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Annan Wang, Wenxiu Sun, Qiong Yan, et al. 2025. [Towards open-ended visual quality comparison](#). In *European Conference on Computer Vision*, pages 360–377. Springer.

- x.ai. 2024. [Grok 1.5v: A New Era in AI Understanding](#). Accessed: 2024-11-07.
- Guowei Xu, Peng Jin, Li Hao, Yibing Song, Lichao Sun, and Li Yuan. 2024a. [Llava-cot: Let vision language models reason step-by-step](#). *ArXiv preprint*, abs/2411.10440.
- Guowei Xu, Peng Jin, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. 2024b. [Llava-cot: Let vision language models reason step-by-step](#). *ArXiv preprint*, abs/2411.10440.
- Zhiyang Xu, Chao Feng, Rulin Shao, Trevor Ashby, Ying Shen, Di Jin, Yu Cheng, Qifan Wang, and Lifu Huang. 2024c. [Vision-flan: Scaling human-labeled tasks in visual instruction tuning](#). *ArXiv preprint*, abs/2402.11690.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. 2024. [Minicpm-v: A gpt-4v level mllm on your phone](#). *ArXiv preprint*, abs/2408.01800.
- Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, Qin Jin, Liang He, Xin Lin, and Fei Huang. 2023a. [UReader: Universal OCR-free visually-situated language understanding with multimodal large language model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2841–2858, Singapore. Association for Computational Linguistics.
- Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, et al. 2023b. [Ureader: Universal ocr-free visually-situated language understanding with multimodal large language model](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2841–2858.
- Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. 2024. [Yi: Open foundation models by 01. ai](#). *ArXiv preprint*, abs/2403.04652.
- Licheng Yu, Patrick Poirson, Shan Yang, Alexander C. Berg, and Tamara L. Berg. 2016. [Modeling context in referring expressions](#). *ArXiv preprint*, abs/1608.00272.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. [Mm-vet: Evaluating large multimodal models for integrated capabilities](#). *ArXiv preprint*, abs/2308.02490.
- Ye Yuan, Xiao Liu, Wondimu Dikubab, Hui Liu, Zhilong Ji, Zhongqin Wu, and Xiang Bai. 2022. [Syntax-aware network for handwritten mathematical expression recognition](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4553–4562.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. 2024a. [Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi](#). In *Proceedings of CVPR*.
- Xiang Yue, Tianyu Zheng, Yuansheng Ni, Yubo Wang, Kai Zhang, Shengbang Tong, Yuxuan Sun, Botao Yu, Ge Zhang, Huan Sun, Yu Su, Wenhui Chen, and Graham Neubig. 2024b. [Mmmu-pro: A more robust multi-discipline multimodal understanding benchmark](#). *ArXiv preprint*, abs/2409.02813.
- Xiang Yue, Tuney Zheng, Ge Zhang, and Wenhui Chen. 2024c. [Mammoth2: Scaling instructions from the web](#). *NeurIPS 2024*.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. [Sigmoid loss for language image pre-training](#). In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986.
- Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. 2019. [RAVEN: A dataset for relational and analogical visual reasoning](#). In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 5317–5327. Computer Vision Foundation / IEEE.
- Kaichen Zhang, Bo Li, Peiyuan Zhang, Fanyi Pu, Joshua Adrian Cahyono, Kairui Hu, Shuai Liu, Yuanhan Zhang, Jingkang Yang, Chunyuan Li, and Ziwei Liu. 2024a. [Lmms-eval: Reality check on the evaluation of large multimodal models](#). *ArXiv preprint*, abs/2407.12772.
- Pan Zhang, Xiaoyi Dong, Yuhang Zang, Yuhang Cao, Rui Qian, Lin Chen, Qipeng Guo, Haodong Duan, Bin Wang, Linke Ouyang, et al. 2024b. [Internlm-xcomposer-2.5: A versatile large vision language model supporting long-contextual input and output](#). *ArXiv preprint*, abs/2407.03320.
- Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Peng Gao, and Hongsheng Li. 2024c. [Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems?](#) *ArXiv preprint*, abs/2403.14624.
- 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. 2024d. [Mavis: Mathematical visual instruction tuning](#). *ArXiv preprint*, abs/2407.08739.
- Ruohong Zhang, Liangke Gui, Zhiqing Sun, Yihao Feng, Keyang Xu, Yuanhan Zhang, Di Fu, Chunyuan Li, Alexander Hauptmann, Yonatan Bisk, et al. 2024e.



- Direct preference optimization of video large multi-modal models from language model reward. *ArXiv preprint*, abs/2404.01258.
- Ruohong Zhang, Bowen Zhang, Yanghao Li, Haotian Zhang, Zhiqing Sun, Zhe Gan, Yinfei Yang, Ruoming Pang, and Yiming Yang. 2024f. [Improve vision language model chain-of-thought reasoning](#). *ArXiv preprint*, abs/2410.16198.
- Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2024g. [Pmc-vqa: Visual instruction tuning for medical visual question answering](#). *ArXiv preprint*, abs/2305.10415.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. 2024h. [Video instruction tuning with synthetic data](#). *ArXiv preprint*, abs/2410.02713.
- Bo Zhao, Boya Wu, Muyang He, and Tiejun Huang. 2023a. [Svit: Scaling up visual instruction tuning](#). *ArXiv preprint*, abs/2307.04087.
- Yilun Zhao, Chen Zhao, Linyong Nan, Zhenting Qi, Wenlin Zhang, Xiangru Tang, Boyu Mi, and Dragomir Radev. 2023b. [RobuT: A systematic study of table QA robustness against human-annotated adversarial perturbations](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6064–6081, Toronto, Canada. Association for Computational Linguistics.
- Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang, Tiejun Huang, and Zheng Liu. 2024. [Mlvu: A comprehensive benchmark for multi-task long video understanding](#). *ArXiv preprint*, abs/2406.04264.
- Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. [TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3277–3287, Online. Association for Computational Linguistics.
- Yuke Zhu, Oliver Groth, Michael S. Bernstein, and Li Fei-Fei. 2016. [Visual7w: Grounded question answering in images](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 4995–5004. IEEE Computer Society.

## Table of Contents in Appendix

<b>A</b>	<b>Related Work</b>	<b>2</b>
<b>B</b>	<b>Additional Details of Training</b>	<b>2</b>
<b>C</b>	<b>Additional Details of Contamination Study</b>	<b>3</b>
<b>D</b>	<b>Additional Results of Ablation Study</b>	<b>3</b>
D.1	Breakdown performance on each benchmark before and after filtering . . . . .	3
D.2	Specific performance on each benchmark for different mix ratio . . . . .	6
D.3	Specific performance on each benchmark as trained data increased . . . . .	6
D.4	Additional Details of different models for data rewriting . . . . .	7
D.5	Additional Details of Verifying the consistency of model filter and manual filter . . . . .	9
<b>E</b>	<b>Additional Details of Rewriting Step</b>	<b>10</b>
E.1	Prompts For Rewriting . . . . .	10
E.2	Score Details Before and After Rewriting . . . . .	14
<b>F</b>	<b>Case Study</b>	<b>19</b>
F.1	Incorrect Cases . . . . .	19
F.2	Good Cases . . . . .	23
<b>G</b>	<b>Model Capabilities and Qualitative Examples</b>	<b>26</b>

## A Related Work

**Multimodal Instruction Data.** The construction of high-quality, large-scale multimodal datasets is a cornerstone for advancing MLLMs. However, this process presents significant challenges, particularly in open-source settings where resource constraints are pronounced. Existing methods for dataset generation typically fall into three categories. First, human annotation provides precise and contextually rich datasets (Xu et al., 2024c; Deitke et al., 2024; McKinzie et al., 2024; Sun et al., 2023), but it is prohibitively expensive and labor-intensive, limiting scalability. Second, repurposing academic datasets such as visual question-answering collections offers cost-effective alternatives and can be enhanced through automated or manual instruction templates (Tong et al., 2024; Liu et al., 2024d), but these datasets often lack diversity and fail to support nuanced reasoning tasks. Third, querying proprietary models like GPT-4 enables the generation of high-quality and diverse datasets (Chen et al., 2023a; Luo et al., 2024; Liu et al., 2024b; Wang et al., 2023; Chen et al., 2024a). However, this approach incurs substantial computational costs and raises licensing issues, making it less feasible for open-source initiatives. Our work addresses these challenges by introducing a simple, scalable, and cost-effective methodology for constructing multimodal datasets exclusively using open-source models, combining task-specific data augmentation and rigorous quality filtering to provide a robust foundation for open-source MLLMs without relying on proprietary resources.

**Multimodal Large Language Models.** MLLMs have advanced AI by seamlessly integrating text and visual processing, with proprietary systems like GPT-4o and Gemini (Gemini Team, 2023) achieving state-of-the-art performance. However, these models remain inaccessible, leaving open-source alternatives at a disadvantage due to limited resources and data. To address this gap, connector-based approaches (Li et al., 2023c; Dai et al., 2023) like LLaVA (Li et al., 2024b) have emerged as efficient solutions, linking visual encoders to language models using lightweight projection modules. Despite these innovations, the primary challenge for open-source MLLMs is the scarcity of high-quality supervised fine-tuning data (Bai et al., 2024), which is essential for advancing their capabilities. Our work tackles this bottleneck by scaling and improving SFT datasets while building on the connector-training paradigm. Through these efforts, we aim to bridge the gap between proprietary and open-source MLLMs, enabling open-source competitive multimodal systems.

## B Additional Details of Training

Unlike Llava-OneVision (Li et al., 2024b), which undergoes a Stage-1.5 training phase with 4M data after Stage-1, followed by Stage-2 training with 3.2M data (SI) and 1.6M data (OV), our approach directly proceeds to Stage-2 training. Consequently, the total training data for LLaVA-OV amounts to 8.8M, whereas our model is trained on a larger dataset of 12M.

	Stage-1	Stage-2	Stage-3
<b>Resolution</b>	384	$384 \times \{1 \times 1, \dots\}$	$384 \times \{1 \times 1, \dots\}$
<b>#Tokens</b>	729	Max $729 \times 5$	Max $729 \times 5$
<b>Dataset</b>	LCS	Single Image	Single, Multi-Image & Video
<b>#Samples</b>	558K	10M	2M
<b>Vision Tower</b>	siglip-so400m-patch14-384		
<b>LLM Backbone</b>	Qwen2.5-7B-Instruct		
<b>Trainable Model Parameters</b>	Projector: 20.0M	Full Model: 8.0B	Full Model: 8.0B
<b>Batch Size</b>	512	256	256
<b>Model Max Length</b>	8192	8192	16384
<b>Learning Rate:</b> $\psi_{vision}$	$1 \times 10^{-3}$	$2 \times 10^{-6}$	$2 \times 10^{-6}$
<b>Learning Rate:</b> $\{\theta_{proj}, \Phi_{LLM}\}$	$1 \times 10^{-3}$	$1 \times 10^{-5}$	$1 \times 10^{-5}$
<b>Epoch</b>	1	1	1

Table 3: Detailed configuration for each training stage of the MAMmoTH-VL-8B model.

DataSet	Multi-Discipline Knowledge and Mathematical Reasoning							
	MMStar	MMMU	MMMU-Pro	SeedBench	MMBench	MMVet	MathVerse	MathVista
	test	val	vision	test	en-test	test	mini-vision	testmini
MAmmoTH-VL-Instruct	0%	0%	0%	0%	0%	0%	0%	0%

Table 4: Contamination Rate On Multi-Discipline Knowledge and Mathematical Reasoning Benchmarks

DataSet	Chart & Doc Understanding				Multimodal Interactions & Preferences		
	AI2D	ChartQA	InfoVQA	DocVQA	RealWorldQA	WildVision	L-Wilder
	test	test	test	test	test	0617	small
MAmmoTH-VL-Instruct	0%	0%	0%	0%	0%	0%	0%

Table 5: Contamination Rate On Chart & Doc Understanding and Multimodal Interactions & Preferences Benchmarks

DataSet	Multi-Image and Video							
	MuirBench	MEGABench	EgoSchema	PerceptionTest	SeedBench	MLVU	MVBench	VideoMME
	test	test	test	test	video	dev	test	w/o subs
MAmmoTH-VL-Instruct	0%	0%	0%	0%	0%	0%	0%	0%

Table 6: Contamination Rate On Multi-Image and Video Benchmarks

## C Additional Details of Contamination Study

Given a training dataset  $D$  and an evaluation dataset  $D_E$ , the objective is to check whether  $D \cap D_E$  is empty. Since MAmmoTH-VL-Instruct is a multimodal dataset, we conduct data contamination detection from both textual and multimodal perspectives. It is noteworthy that data in the dataset that is similar to benchmark texts does not necessarily constitute contamination (as it might be part of a benchmark’s training set with similar texts but different images), whereas data with similar images is definitely contaminated. Specifically, for text detection, we employ the Ngram method for monitoring (the threshold is 0.5); for image detection, we utilized the PHash(Perceptual Hash) algorithm for identification ( $HammingDistance \geq 10$ ). Regarding video data, we follow the frame extraction logic of Llava-OneVision (Li et al., 2024b), extracting frames used in both the training stage and benchmark evaluation, converting them into images, and then stitching them into a large composite image for subsequent image-based detection. As shown in Table 4, Table 5, and Table 6, MAmmoTH-VL-Instruct exhibits no data contamination issues. This rigorous approach to contamination detection ensures that our model’s performance metrics are not artificially inflated by overlapping data, thereby maintaining the validity of our evaluation results.

## D Additional Results of Ablation Study

### D.1 Breakdown performance on each benchmark before and after filtering

To evaluate the impact of our filtering approach, we conduct experiments using the Qwen2.5-1.5B-Instruct model under identical experimental conditions. We randomly sample 1M instances from the same dataset in both its pre-filtered and post-filtered versions to ensure a fair comparison. The experimental results in Table A1 demonstrate that the model trained on the after filtering dataset outperforms that trained on the before filtering dataset, particularly in the benchmark for Chart, Diagram, and Document Understanding. This aligns with our previous observation that OCR and Chart data constitute the majority of the data filtered. This finding effectively validates the necessity and significance of hallucination mitigation. It also indirectly highlights the persistent issue of hallucinations in current MLLMs.

Table A2 shows the specific numbers and retention rates for each type of data before and after filtering.

<b>Bench Name</b>	<b>Before Filter</b>	<b>After Filter</b>
MMMU	39.6	<b>40.9</b>
MMStar	14.0	<b>44.6</b>
SeedBench	66.4	<b>67.9</b>
MMMU-Pro Vision	<b>15.5</b>	13.7
MathVista	39.5	<b>42.0</b>
MMBench EN	58.6	<b>65.1</b>
MMVet	40.5	<b>43.9</b>
MathVerse	19.3	<b>22.6</b>
AI2D	56.9	<b>61.8</b>
ChartQA	26.8	<b>63.0</b>
InfoVQA	41.5	<b>48.0</b>
DocVQA	71.7	<b>76.5</b>
L-Wilder Small	58.8	<b>59.8</b>
WildVision	40.2	<b>42.2</b>
RealWorldQA	50.3	<b>56.0</b>
Avg	42.6	<b>49.9</b>

Table A1: Performance Comparison of Models Trained on Filtered versus Unfiltered Data Across Multiple Benchmarks.

<b>Data Type</b>	<b>Before Filter</b>	<b>After Filter</b>	<b>Filter Rate</b>
OCR	1104960	498337	54.9
Chart	7326189	3782029	48.4
GeneralQA	1726180	1584308	8.2
Caption	244874	199853	18.3
Math	590894	518393	12.3
Other	1315039	1178275	10.4

Table A2: Filter Rates Of Different Data Types After Data Filtering.

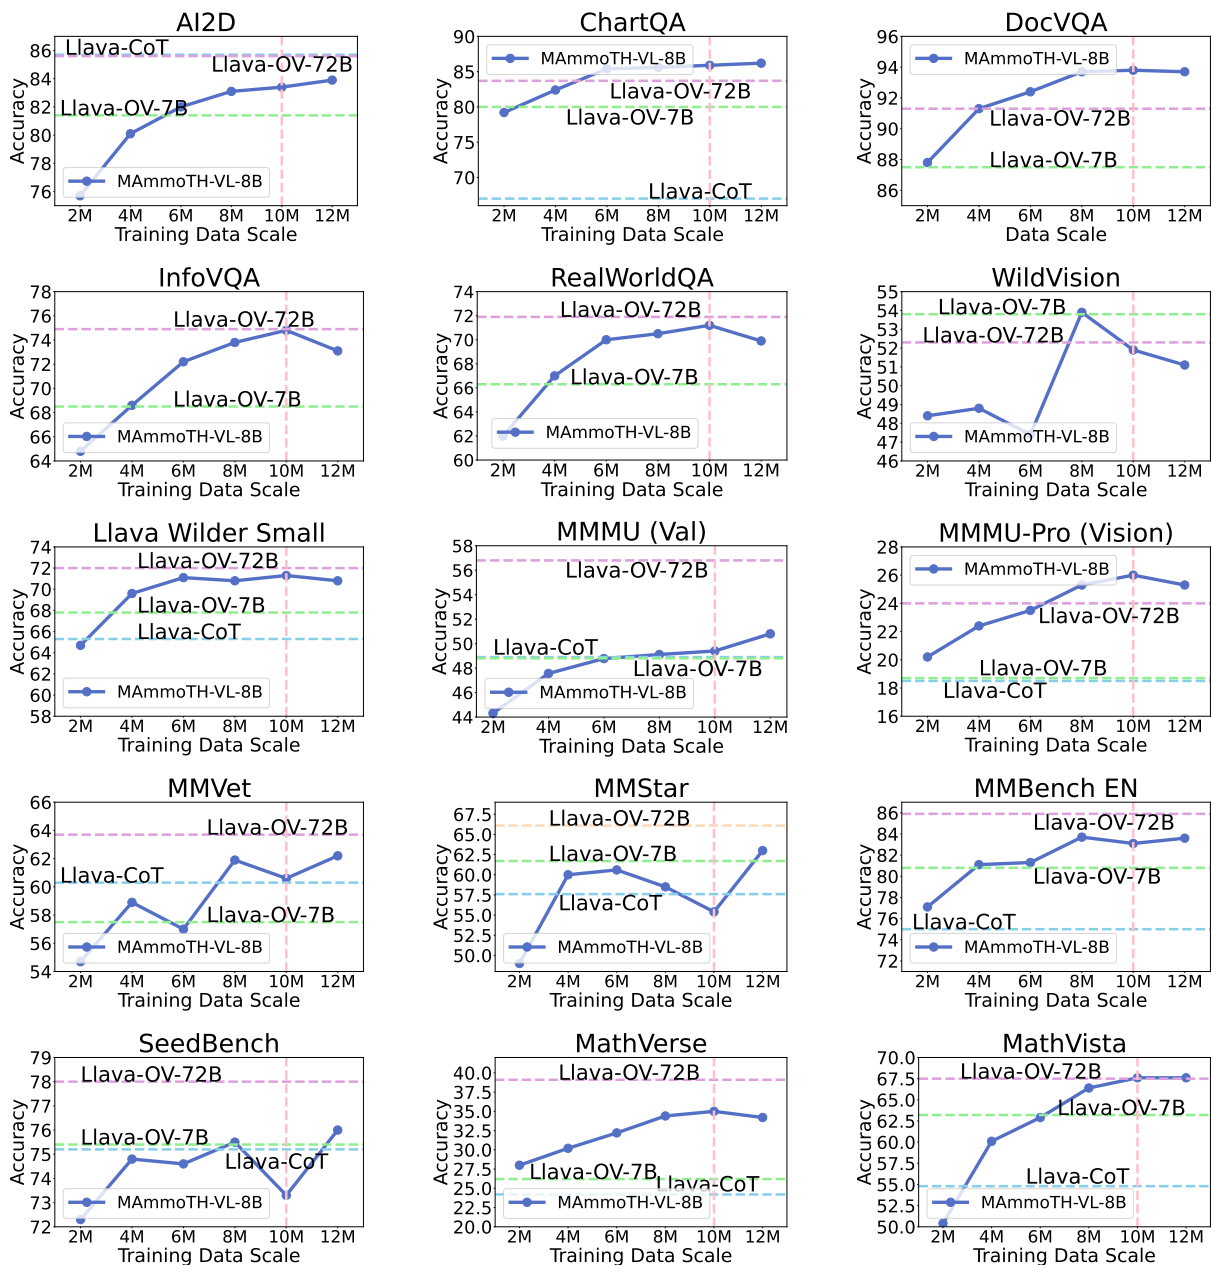


Figure A1: MAMmoTH-VL-8B’s overall performance across various benchmarks as the training dataset size increases.

## D.2 Specific performance on each benchmark for different mix ratio

To validate the effectiveness of merging original and rewritten datasets, we conduct experiments by training the model in five data schedules: first trained on the original dataset, second on the rewritten dataset, third on a combined dataset with a 3:7 ratio of original to rewritten data, fourth with a 7:3 ratio and fifth with a 5:5 ratio. The results in Table A3 and Figure A2 indicate that the model trained on the rewritten dataset achieved better average performance compared to the model trained on the original dataset. In addition, the model trained on the combined dataset with the merge ratio of 3:7 shows a slight performance gain which is better than rewritten. This demonstrates the effectiveness of merging datasets, as it enables the construction of a more diverse and comprehensive dataset.

Bench Name	Rewrite	Original	Mix 3:7	Mix 7:3	Mix 5:5
MMMU	40.9	<b>41.9</b>	41.5	41.3	41.7
MMSStar	<b>44.6</b>	43.3	43.4	42.3	43.7
SeedBench	67.9	<b>69.9</b>	68.7	69.3	68.9
MMMU-Pro Vision	13.7	13.0	<b>13.8</b>	13.5	13.5
MathVista	<b>42.0</b>	40.4	41.8	40.6	39.5
MMBench EN	65.1	67.8	66.1	<b>67.9</b>	66.4
MMVet	43.9	37.3	<b>45.5</b>	40.7	38.9
MathVerse	<b>22.6</b>	19.8	21.4	21.0	20.4
AI2D	61.8	<b>63.1</b>	62.9	62.5	62.8
ChartQA	<b>63.1</b>	56.5	61.1	56.8	56.6
InfoVQA	48.0	47.3	<b>49.0</b>	45.7	45.6
DocVQA	76.5	76.6	<b>77.4</b>	76.0	75.7
L-Wilder Small	59.8	56.4	<b>60.9</b>	56.8	57.4
WildVision	<b>42.2</b>	34.9	38.7	34.5	36.7
RealworldQA	56.0	56.1	<b>57.1</b>	55.7	54.8
Avg	49.9	48.3	<b>50.0</b>	48.3	48.2

Table A3: Benchmark Performance Of Models Trained On Data With Different Mix Ratios.

## D.3 Specific performance on each benchmark as trained data increased

Figure A1 show MAMmoTH-VL-8B’s overall performance across all single image benchmarks as the training dataset size increased, with each node representing 2 million data points. The results are compared against those of LLaVA-OneVision-7B, Llava-CoT-11B, and LLaVA-OneVision-72B. This trend clearly demonstrates that expanding the scale of instruction data has a significant positive effect on model performance. This observation suggests that as more diverse instruction data is introduced, the model’s ability to handle complex tasks is enhanced.




Figure A2: Effect of the mixing ratio of rewritten and original data. Performance generally improves with increased data rewriting. We select the 70% rewritten data level as a balanced choice.

#### D.4 Additional Details of different models for data rewriting

Bench Name	Original	Rewrite (Qwen2-VL-7B)	Rewrite (InternVL2-8B)	Rewrite (InternVL2-76B)
MMMU	40.4	40.6	<b>40.9</b>	40.78
MMStar	40.9	41.7	<b>41.7</b>	37.9
SeedBench	50.6	52.1	65.0	<b>67.0</b>
MMMU-Pro Vision	12.3	12.9	12.9	<b>15.3</b>
MathVista	36.4	38.8	37.4	<b>39.0</b>
MMBench EN	<b>65.8</b>	59.1	60.1	58.3
MMVet	38.6	38.1	38.6	<b>41.1</b>
MathVerse	17.6	<b>21.6</b>	19.8	20.6
AI2D	61.8	<b>62.3</b>	61.7	59.6
ChartQA	49.4	48.1	50.6	<b>58.7</b>
InfoVQA	43.8	43.1	43.7	<b>44.3</b>
DocVQA	<b>73.4</b>	70.8	71.3	72.2
L-Wilder Small	44.5	55.7	55.7	<b>60.5</b>
WildVision	32.7	32.0	30.8	<b>41.7</b>
RealWorldQA	56.5	55.1	<b>56.8</b>	53.5
Avg	46.8	47.3	48.4	<b>50.0</b>

Table A4: Performance On Different Benchmarks Of Models Trained On Data Rewritten By Different Models

Table A4 displays the performance on different benchmarks of models trained on data rewritten by different models. Figure A3 shows the filter rate of models on various data types. It can be seen that the retention rate of the data rewritten by larger model is much higher than that rewritten by models with smaller parameters, but it is worth noting that the retention rate of all three models in terms of ocr data is






Figure A3: Different Filter Rates For Rewritten Data Filtered By Different Models On Different Data Types.

very low, which highlights the current problems facing multimodal model.

## D.5 Additional Details of Verifying the consistency of model filter and manual filter

\	Model	Evaluator1	Evaluator2	Evaluator3
Model	-	0.73	0.70	0.63
Evaluator1	0.73	-	0.70	0.42
Evaluator2	0.70	0.70	-	0.53
Evaluator3	0.63	0.42	0.53	-

Table A5: Kappa Value Between Any Two.

Table A5 indicate that by averaging the Kappa values(McHugh, 2012), we obtain an average pairwise Kappa value of 0.55 among three evaluators. When the model's filtering decision is used to replace one of the evaluators, the average Kappa of the three new derived evaluator combinations is 0.64 (eg., calculate among Model, Evaluator1 and Evaluator2). Based on the Kappa value range (0.60 to 0.80 for good consistency), the agreement between model and manual judge is quite strong, suggesting that the process of model filtering is reliable and consistent with human judgment in this work.

## E Additional Details of Rewriting Step

### E.1 Prompts For Rewriting

#### Prompt for Score

You are given a Q&A conversation and a corresponding image. - Conversations: Your goal is to:

1. Evaluate Information Content and Complexity: Rate the combined information richness and complexity of the Q&A conversation and image on a scale from 1 to 5. Consider factors like the amount of detail, depth of content, and how well the conversation and image complement each other in conveying comprehensive information.

1: Minimal detail, shallow content, limited interaction between conversation and image.

2: Some detail, moderate depth, a basic connection between the conversation and image.

3: Good amount of detail, some complexity, moderate complementarity between conversation and image.

4: High level of detail, deep content, the conversation and image work well together to convey a thorough understanding.

5: Very rich in detail, highly complex, the conversation and image are seamlessly integrated to provide a comprehensive, insightful picture.

2. Evaluate Relevance: Rate the relevance of the conversation to the image on a scale from 1 to 5.

1: Very low relevance, the conversation and image are almost unrelated.

2: Low relevance, the conversation and image share some overlap, but one is mostly independent of the other.

3: Moderate relevance, there is some connection, but the conversation could stand alone without the image or vice versa.

4: High relevance, the conversation and image are closely tied, and both contribute significantly to each other's meaning.

5: Very high relevance, the conversation and image are inseparable, and one cannot fully be understood without the other.

Output Format:

1. Information Content Score (1-5):

2. Relevance (1-5):

Only provide the final result in the above structured format without any additional explanations.

## Rewrite Prompt for GeneralQA

Your objective is to craft a series of advanced <instruction, response >pairs derived from a specific image and its associated question-answer pair. First, choose an appropriate task type from the list below as the basis for your rewrite:

- 1. Coding & Debugging
- 2. Information Seeking
- 3. Creative Writing
- 4. Reasoning
- 5. Planning
- 6. Math
- 7. Editing
- 8. Data Analysis
- 9. Role Playing
- 10. Brainstorming
- 11. Advice Seeking

**\*\*Guidelines for New Instructions:\*\***

- Each instruction should be substantially more intricate than the original, introducing complex and thought-provoking elements.
- Focus on deeper aspects of the original topic, incorporating critical thinking, detailed analysis, or creative problem-solving.
- Where possible, connect the topic with other related fields or practical applications.

**\*\*Guidelines for Responses:\*\***

- Responses must be detailed and comprehensive, covering all aspects of the instruction.
- Include diverse perspectives or approaches, and integrate examples, case studies, or hypothetical scenarios.
- Discuss potential implications, challenges, or future developments in relation to the topic.

**\*\*Format:\*\***

Each <instruction, response>pair should be formatted as follows:

- <Instruction: [In-depth, sophisticated instruction linked to the initial Q&A]>
- <Response: [Extensive, expert-level response detailing the instruction's complexities]>

Aim to generate at least pair\_num pairs, ensuring each set is unique and progressively builds upon the knowledge from the initial Q&A. The goal is to elevate a simple concept into a sequence of intricate, expert-level discussions.

**\*\*Based on the following assets:\*\***

- **\*\*Image:\*\*** <image>
- **\*\*Initial Q&A:\*\*** qa\_text

## Rewrite Prompt for Chart

Your objective is to craft a series of advanced <instruction, response>pairs derived from a specific image that is associated question-answer pair. The image could be a screenshot of table, scenetext images or other charts. First, choose an appropriate task type from the list below as the basis for your rewrite:

- 1. Data Analysis
- 2. Information Seeking
- 3. Planning
- 4. Editing
- 5. Reasoning

Then you need to generate new <instruction, response>pairs should be progressively built upon the knowledge from the initial Q&A. The goal is to elevate a simple concept into a sequence of intricate, expert-level discussions.

Guidelines for New Instructions:

- Each instruction should be substantially more intricate than the original, introducing complex and thought-provoking elements.
- Focus on deeper aspects of the original topic, incorporating critical thinking, detailed analysis, or creative problem-solving.
- Where possible, connect the topic with other related fields or practical applications.

Guidelines for Responses:

- Responses must be detailed and comprehensive, covering all aspects of the instruction.
- Include diverse perspectives or approaches, and integrate examples, case studies, or hypothetical scenarios.
- Discuss potential implications, challenges, or future developments in relation to the topic.

Format:

Each <instruction, response>pair should be formatted as follows and use the starting "##Instruction##:" and "##Response##:"

- ##Instruction##: [In-depth, sophisticated instruction linked to the initial Q&A]
- ##Response##: [Extensive, expert-level response detailing the instruction's complexities]

Based on the following assets:

- \*\*Image:\*\* <image>
- \*\*Initial Q&A:\*\* qa\_text

## Rewrite Prompt for Math

Your task is to generate a detailed, step-by-step solution for a given math problem that includes an image. The answer provided is simple, but the goal is to fill in all the missing intermediate steps required to reach the final answer.

Please adhere to the following guidelines:

1. Response Crafting:

- Provide a comprehensive response that includes all intermediate steps, explanations of the reasoning process, and any relevant mathematical concepts or operations used to solve the problem.
- Ensure the response is exhaustive, covering each stage required to reach the final answer, while considering all details from the image.

2. Content Requirements:

- The response should be well-structured, with clear steps and explanations that thoroughly address the instruction.
- Do not include additional text or explanations outside of the required <response>.

### Example:

Question: What is the distance between points B and C?

Simple Answer: The answer is 9.

Revised Answer: <response: To find the distance between points B and C', we start by analyzing the given geometric configuration... [continue with detailed steps].>

Question:

Simple Answer:

## Rewrite Prompt for Caption

Your task is to generate a series of <instruction, response>pairs based on a given image caption while selecting one suitable task type from the following list for rewriting:

- 1. Coding & Debugging
- .....
- 11. Advice Seeking

The goal is to create challenging tasks that require a deep understanding of the visual content described in the caption, without explicitly mentioning too many details from the caption itself.

Please adhere to the following guidelines:

### 1. Instruction Generation:

- Select one task type from the above list and generate at least <instruction, response>pairs based on the chosen task.
- Ensure each instruction is unique, complex, and related to the given caption and task type.
- Avoid directly repeating specific details from the caption.
- Instructions should require critical thinking, analysis, or creative problem-solving based on a clear understanding of the visual scene.

### 2. Response Crafting:

- Provide detailed, comprehensive responses that demonstrate accurate interpretation of the visual information.
- Include in-depth explanations, multiple perspectives, or detailed steps where appropriate.
- Ensure responses clearly show an accurate grasp of the visual information implied by the instruction.

### 3. Content Requirements:

- Instructions should be varied, challenging, and explore different advanced aspects or potential sophisticated interpretations of the visual scene.
- Instructions should require the responder to infer and utilize visual information that may not be explicitly stated in the instruction.
- Responses should demonstrate expertise in both understanding the implied visual content and providing substantial, valuable information related to the task.
- Responses should be well-structured with clear sections or paragraphs.

### 4. Output Clarification:

- Do not include additional text or explanations outside of the required <instruction, response>pairs.
- Your output should only consist of the generated <instruction, response>pairs.

### ### Example:

Caption: a street scene with construction scaffolding, three individuals, a shopping cart filled with personal belongings, street signs, and a sidewalk. The construction scaffolding is blue and has text about the construction company and contact details. One individual is walking by, another person is standing and looking at something in their hands, and the third person is behind a shopping cart going through the items inside. There are no vehicles visible in the image. The shopping cart filled with personal belongings might suggest that the individual using it is homeless or transient. The time of day appears to be daytime, but no direct indicators of weather conditions are present, although the presence of personal umbrellas might suggest rain.

Task Category: Information Seeking

<Instruction: Analyze the social dynamics at play in this street scene, taking into account the individuals' body language, interactions, and apparent activities. How do their actions and postures reveal or conceal their relationships with each other and their environment?>

<Response: Upon close examination, it appears that the individual walking by is intentionally avoiding eye contact with the other two individuals, suggesting a sense of detachment or disengagement. The person standing and looking at something in their hands seems engrossed in their own world, oblivious to the surroundings. Meanwhile, the person behind the shopping cart is intensely focused on rummaging through the items, indicating a sense of urgency or desperation. The lack of eye contact or acknowledgement among the three individuals implies a lack of community or social bonding in this urban setting. Furthermore, the absence of vehicles and the presence of construction scaffolding create an atmosphere of disruption and transition, which may be contributing to the individuals' isolation. The fact that one individual is potentially homeless or transient adds a layer of complexity to the social dynamics, as they may be grappling with issues of poverty, displacement, or marginalization.>

Caption: { }

## Rewrite Prompt for OCR

I have an image and several corresponding Visual Question Answering (VQA) pairs. I'd like you to imagine a scenario and create a richer set of dialogues based on the information known from the VQA pairs. This dialogue set should include question-answering and instruction-based tasks.

The dialogue should revolve around the image, but the questions should not "spoil" too much information about the image content. This is because I intend to use this dialogue set to train a vision-language model, and excessive spoilers would make the learning process too easy.

In your output, please first provide a randomly imagined scenario (try to make it relevant to the image information), and then list the set of dialogues you've created.

When creating the dialogue:

1. Vary the types of questions and tasks (e.g., open-ended questions, specific queries, instructions)
2. Maintain a natural conversational flow
3. Include some challenging question or instruction that requires complex reasoning or analysis
4. Avoid directly stating key visual elements in the questions
5. Provide detailed and comprehensive responses for ALL answers
6. The AI assistant should act as if it's unaware of the scenario at the beginning of the conversation
7. Minimize the use of conversational fillers or emotional language in questions. The AI assistant's tone should be neutral and professional
8. The AI assistant should strive to be as helpful as possible, answering questions to the best of its ability based on the given information
9. The human should not ask questions that go beyond the scope of the information provided in the VQA pairs

The questioner is a curious human seeking help, and the answerer is an AI assistant capable of perceiving and analyzing images.

Format requirement:

Please format the dialogue as a series of alternating messages, with at most 5 rounds of conversation:

Scenario: [The scenario related to the image.]

Human: [First question, without conversational fillers]

Assistant: [Detailed and comprehensive answer to the first question, based only on VQA information]

Human: [Second question or instruction, without conversational fillers]

Assistant: [Detailed and comprehensive response to the second question or instruction, based only on VQA information]

...

Each message should be on a new line, clearly indicating the speaker. Do not mention the human or assistant's name in the dialogue. Ensure that all responses from the Assistant are detailed and comprehensive, providing as much relevant information as possible based on the VQA pairs.

Please create this rich dialogue set based on the brief VQA pairs I will provide. Do not include any additional content or explanations outside of the specified output format.

VQA: vqa

## Prompt for Data Filtering

**\*\*Q&A:\*\*** {}

Evaluate the provided question and answer based on the given image. Specifically:

1. **\*\*Relevance\*\***: Determine whether the question is directly related to the content of the image.
2. **\*\*Consistency\*\***: Verify if the answer is logically coherent and aligns with the question.
3. **\*\*Accuracy\*\***: Confirm whether both the question and answer accurately reflect the details or facts depicted in the image.

Provide your response as follows:

- Respond with **\*\*Yes\*\*** if the question and answer are relevant, consistent, and factually accurate.
- Respond with **\*\*No\*\*** if there are any issues with relevance, logical coherence, or factual accuracy.

Only respond with a single word Yes or No.

## E.2 Score Details Before and After Rewriting

Table A6 shows the comparison of Original and Rewritten data on Content and Relevance Scores. Rewritten data scores higher, indicating improved quality.

Source	Original Content & Relevance Score	Rewrite Content & Relevance Score
chrome-writting	3.0/3.4	4.0/4.7
OCRVQA	3.9/4.6	3.9/4.5
screen-qa	3.4/4.2	4.0/4.3
hme	3.3/4.3	4.0/4.8
textvqa	3.2/4.1	4.0/4.5
docvqa	3.9/4.8	4.0/4.8
st-vqa(cauldron,llava-format)	3.0/3.8	4/4.4
Ureader Chart	4.0/4.8	4.0/4.8
infographic-vqa	4.4/4.9	4.2/4.8
finqa	4.0/4.3	3.9/4.1
ureader-kg	4.0/4.7	4.1/4.7
chartqa(cauldron,llava-format)	4.0/4.9	4.0/4.9
vistext(cauldron)	4.0/4.8	4.0/4.8
chart2text(cauldron)	4.0/4.7	4.0/4.7
ureader-qa	4.0/4.7	4.1/4.7
Irv-chart	4.0/4.9	4.0/4.8
idefics375k	3.6/4.4	3.9/4.5
cambrian(filtered)	3.2/3.6	3.8/4.1
GQA	3.2/4.0	3.8/4.1
AlfWorld	3.0/3.6	3.2/3.8
IDK	3.0/3.6	3.8/4.2
CLlava Instruct	3.7/4.1	3.9/4.1
llava-zh	3.5/4.1	3.9/4.1
SVITCore	3.7/4.1	3.9/4.1
SVITCore mix	3.4/4.1	3.9/4.1
visual7w(cauldron,llava-format)	3.1/4.0	3.9/4.2
sharegpt4v	3.9/4.2	3.7/3.7
infographic(gpt4v)	4.3/5.0	4.1/4.5
sharegpt4o	4.0/4.4	3.8/3.9
sharegpt4v(coco)	3.9/4.0	3.7/3.7
sharegpt4v(llava)	3.7/3.9	3.7/3.8
sharegpt4v(sam)	3.9 / 3.9	3.8/3.7
Geo170K	3.9/4.7	4.0/4.8
MathVision	3.7/4.7	4.0/4.9
CLEVR-Math(MathV360K)	3.0/3.8	3.1/3.8
GEOS(MathV360K)	3.8/4.8	4.0/4.9
GeoQA+(MathV360K)	3.9/4.9	4.0/4.9
Geometry3K(MathV360K)	3.9/4.9	4/4.9
IconQA(MathV360K)	2.7/4.2	3.0/4.2
PMC-VQA(MathV360K)	3.9/4.3	4.0/4.6
Super-CLEVR(MathV360K)	3.4/4.0	3.3/4.0
TabMWP(MathV360K)	3.3/4.2	3.7/4.7
UniGeo(MathV360K)	3.9/4.9	4.0/4.9
VizWiz(MathV360K)	3.0/4.0	3.5/4.3
MapQA(MathV360K)	3.0/2.4	3.9/4.3
raven(cauldron)	3.3/4.2	3.9/4.6
M3IT+FLAN	3.2/4.0	3.4/4.1
WIT	3.9/4.5	4.1/4.7
ViQuAE	3.3/3.8	3.9/4.0
aokvqa(cauldron,llava-format)	3.0/3.7	3.6/4.1
vision-flan(filtered)	3.0/3.7	3.4/4.0
websight(cauldron)	3.7/4.4	3.8/4.4
vsr(cauldron,llava-format)	3.0/3.5	3.0/3.7
CLEVR	3.0/3.8	3.2/3.9
tallyqa(cauldron,llava-format)	2.9/3.7	3.1/4.0
scienceqa(nona-context)	3.4/4.0	3.5/4.1
PathVQA (164)	3.3/3.6	4.0/4.2
tqa(cauldron,llava-format)	2.6/3.6	2.8/3.9
vqarad(cauldron,llava-format)	4.0/4.7	4.0/4.7
<b>Average</b>	<b>3.5 / 4.2</b>	<b>3.8 / 4.4</b>

Table A6: Comparison of Original and Rewrite Average Content and Relevance Scores



Group A(37.9%)		
■ ReCTs (Liu et al., 2019)	■ IAM(cauldron) (Marti and Bunke, 2002)	■ Ureader OCR (Ye et al., 2023b)
■ RenderedText(cauldron) (Wendler, 2024)	■ IIIT 5K (Mishra et al., 2012)	■ llavar-gpt4-20k (Li et al., 2024b)
■ ArxivQA (Li et al., 2024d)	■ TextOCR(cleaned) (Singh et al., 2021)	■ Visualmrc(cauldron) (Tanaka et al., 2021)
■ DVQA (Kafle et al., 2018)	■ mPLUG-DocOwlchart (Hu et al., 2024)	■ Diagram-Image-to-Text(cauldron) (Li et al., 2024b)
■ HiTab(cauldron,llava-format) (Cheng et al., 2022)	■ RobuT-WTQ(cauldron,llava-format) (Zhao et al., 2023b)	■ RobuT-SQA(cauldron) (Zhao et al., 2023b)
■ PlotQA (Methani et al., 2020)	■ TobuT-WikiSQL(cauldron) (Zhao et al., 2023b)	■ TAT-QA (Zhu et al., 2021)
■ ALLaVA-Vflan (Chen et al., 2024a)	■ WildVision-Chat (Lu et al., 2024)	■ ALLaVA-Laion (Chen et al., 2024a)
■ Laion-GPT4V (Chen et al., 2024a)	■ LVIS-InstructV4 (Wang et al., 2023)	■ ShareGPT4V(new) (Self Collection)
■ ALLaVA-4V-Chinese(vflan) (Chen et al., 2024a)	■ Sherlock (Hessel et al., 2022)	■ ShareGPT4V (Self Collection)
■ Mavis-Math-Rule-Geo (Zhang et al., 2024d)	■ Irv-Normal(filtered) (Liu et al., 2023b)	■ ALLaVA-4V-Chinese(Laion) (Chen et al., 2024a)
■ TabMWP(cauldron) (Lu et al., 2023b)	■ Geo3k (Lu et al., 2021a)	■ Mavis-Math-Metagen (Zhang et al., 2024d)
■ ShareGPT4V(knowledge) (Chen et al., 2025)	■ AI2D(gpt4v) (Kembhavi et al., 2016)	■ Geomverse(cauldron) (Kazemi et al., 2024)
■ Arxiv-Chart-4o (Self Collection)	■ Design2Code (Si et al., 2024)	■ MultiUI (Liu et al., 2024c)
■ Coinstruct (Wu et al., 2025)	■ NLVR2 (Suhr et al., 2019)	■ Mimic CGD (Li et al., 2023a)
■ IconQA (Lu et al., 2021b)	■ HQ-Edit (Hui et al., 2024)	■ Raven (Zhang et al., 2019)
■ FlintstonesSV (Gupta et al., 2018)	■ VIST (Huang et al., 2016)	■ Contrast-Caption (Jiang et al., 2024)
■ M4 Instruct Video (Li et al., 2024c)	■ PororoSV (Li et al., 2019)	■ LLaVA-Video (Zhang et al., 2024h)
■ LLaVA-Video-NeXT-QA (Zhang et al., 2024h)	■ LLaVA-Video-ActivityNetQA (Zhang et al., 2024h)	■ LLaVA-Hound (Zhang et al., 2024e)
■ LLaVA-Video-PerceptionTest (Zhang et al., 2024h)	■ VideoChatGPT (Maaz et al., 2024)	■ Video-MME (Fu et al., 2024)
	■ EgoSchema (Mangalam et al., 2023)	

Figure A4: MAMmoTH-VL data source group. Group A: datasets that are directly kept without modification.

Group B(39.2%)		
■ ScreenQA (Hsiao et al., 2022)	■ Chrome-Writing (Li et al., 2024b)	■ OCRVQA (Mishra et al., 2019)
■ DocVQA(llava) (Mathew et al., 2021b)	■ HME (Yuan et al., 2022)	■ TextVQA (Singh et al., 2019)
■ Infographic-VQA (Mathew et al., 2021a)	■ ST-VQA(cauldron,llava-format) (Biten et al., 2019)	■ Ureader Chart (Ye et al., 2023a)
■ ChartQA(cauldron,llava-format) (Masry et al., 2022b)	■ FinQA (Chen et al., 2021b)	■ Ureader-KG (Ye et al., 2023a)
■ Ureader-QA (Ye et al., 2023a)	■ VisText(cauldron) (Tang et al., 2023)	■ Chart2Text(cauldron) (Obeid and Hoque, 2020)
■ Cambrian(filtered) (Tong et al., 2024)	■ Irv-chart (Liu et al., 2024a)	■ Idefics (Tong et al., 2024)
■ IDK (Cha et al., 2024)	■ GQA (Hudson and Manning, 2019)	■ AlfWorld (Tong et al., 2024)
■ SVITCore (Zhao et al., 2023a)	■ CLlava Instruct Self Collection	■ LLaVA-zh (OpenBMB, 2023)
■ ShareGPT4v (Chen et al., 2023a)	■ SVITCore-mix (Zhao et al., 2023a)	■ Visual7W(cauldron,llava-format) (Zhu et al., 2016)
■ ShareGPT4V(COCO) (Li et al., 2024b)	■ Infographic (Mathew et al., 2021a)	■ ShareGPT4o (Shanghai AI Laboratory, 2023)
■ Geo (Gao et al., 2023)	■ ShareGPT4V(LLAVA) (Li et al., 2024b)	■ ShareGPT4V(SAM) (Li et al., 2024b)
■ GEOS(MathV360K) (Seo et al., 2015)	■ MathVision (Wang et al., 2024b)	■ CLEVR-Math(MathV360K) (Lindström, 2022)
■ IconQA (MathV360K) (Lu et al., 2021b)	■ GeoQA+(MathV360K) (Chen et al., 2021a)	■ Geometry3K (MathV360K) (Lu et al., 2021a)
■ TabMWP (MathV360K) (Lu et al., 2023b)	■ PMC-VQA (MathV360K) (Zhang et al., 2024g)	■ Super-CLEVR (MathV360K) (Li et al., 2023f)
■ MapQA (MathV360K) (Chang et al., 2022)	■ UniGeo (MathV360K) (Chen et al., 2022)	■ VizWiz (MathV360K) (Gurari et al., 2018)
■ WIT (Srinivasan et al., 2021)	■ Raven (cauldron) (Zhang et al., 2019)	■ M3IT+FLAN (Li et al., 2023e)
■ Vision-Flan (filtered) (Xu et al., 2024c)	■ ViQuAE (Lerner et al., 2022)	■ A-OKVQA (cauldron, llava-format) (Schwenk et al., 2022)
■ CLEVR (Johnson et al., 2017)	■ WebSight (cauldron) (Laurençon et al., 2024)	■ VSR (cauldron, llava-format) (Liu et al., 2023a)
■ PathVQA (He et al., 2020)	■ TallyQA (cauldron, llava-format) (Acharya et al., 2019)	■ ScienceQA (nona-context) (Lu et al., 2022)
■ RCTW-17 (Shi et al., 2017)	■ TQA (cauldron, llava-format) (Kembhavi et al., 2017)	■ VQA-RAD (cauldron, llava-format) (Lau et al., 2018)

Figure A5: MAMmoTH-VL data source group. Group B: datasets that are used for rewriting.

Group C(22.9%)		
■ K12-Printing (Li et al., 2024b)	■ Iconqa (cauldron,llava-format) (Lu et al., 2021b)	■ CN-OCR-1 (Russakovsky et al., 2015)
■ Ureader-IE (Ye et al., 2023a)	■ Orand-Car-A (Russakovsky et al., 2015)	■ Sroie (Huang et al., 2019)
■ Q-Instruct (Wu et al., 2023a)	■ Inter-GPs (cauldron,llava-format) (Lu et al., 2021a)	■ Q-Align (Wu et al., 2023b)
■ SketchyVQA (Tu et al., 2023)	■ VizWiz (Gurari et al., 2018)	■ OODVQA (Tu et al., 2023)
■ Hatefulmemes (cauldron,llava-format) (Kielia et al., 2020)	■ CLEVR (cauldron,llava-format) (Johnson et al., 2017)	■ FigureQA (cauldron,llava-format) (Kahou et al., 2018)
■ VQAv2 (Goyal et al., 2017)	■ Screen2Words (cauldron) (Wang et al., 2021)	■ COCOQA (Ren et al., 2015)
■ image-textualization (filtered) (Pi et al., 2024)	■ LNQA (Changpinyo et al., 2022)	■ COCOCaption-Train (Lin et al., 2015)
■ CAP2QA (Cha et al., 2024)	■ COCOCaption-Val (Lin et al., 2015)	■ MSCOCO (Lin et al., 2015)
■ TextCaps (Sidorov et al., 2020)	■ Macaw-LLM (Lyu et al., 2023)	■ DetGPT (Pi et al., 2023)
■ LocalizedNarratives (Pont-Tuset et al., 2020)	■ Ureader-Cap (Ye et al., 2023a)	■ Ureader-Caption (Ye et al., 2023a)
■ MapQA (cauldron,llava-format) (Chang et al., 2022)	■ FigureQA (MathV360K) (Kahou et al., 2018)	■ Hand-Written-Arith (Self Collection)
	■ VisualGenome (Krishna et al., 2017)	■ RefCOCO (Yu et al., 2016)

Figure A6: MAMmoTH-VL data source group. Group C: data sources that are not included.

## F Case Study

### F.1 Incorrect Cases

#### OCR Incorrect Case



**Original Question:** what's the amount on the page?

**Original Answer:** 1999

**Incorrect Rewrite Question:** A tech-savvy customer is interested in purchasing a computer system that can support high-performance graphics. Based on the advertisement, what specific features of the Northgate 286 system would you highlight to convince them to make a purchase, and how do these features compare to the industry standards of 1987?

**Incorrect Rewrite Answer:** To convince the customer to purchase the Northgate 286 system, I would highlight the

following features that are relevant to high-performance graphics:

The system's "Super VGA Graphics" capability, which suggests that it can support high-resolution graphics. The inclusion of a "Northgate Graphics Card", which is specifically designed to enhance graphics performance. The "101-key Enhanced Keyboard" with its "OS/2 READY!" label, implying that the system is compatible with Operating System/2 (OS/2), which was known for its graphics capabilities. In 1987, these features would have been considered impressive, especially considering the system's competitive pricing. The Super VGA Graphics capability would have provided high-quality graphics performance, while the Northgate Graphics Card would have further enhanced the system's graphics capabilities. The OS/2 compatibility would have also been seen as a significant advantage, given the operating system's reputation for graphics support. It's worth noting that the system's 12MHz clock speed, 1MB of RAM, and 60MB hard disk drive would have also contributed to its overall performance, making it an attractive option for customers requiring high-performance graphics. By highlighting these features, I would be able to demonstrate the Northgate 286 system's capabilities and value proposition to the customer, making a compelling case for its purchase.


**Analysis:** Some details in the answer do not fully match the information in the image. The analysis is as follows:

1. Super VGA Graphics: The image does not mention the "Super VGA Graphics" functionality, so this point is inaccurate.
2. Northgate Graphics Card: The image does not mention the "Northgate Graphics Card" or any specific graphical capabilities, so this point is also incorrect.
3. OS/2 READY: The image indeed has an "OS/2 READY!" label, which is correct. However, this only indicates system compatibility with OS/2 and does not necessarily relate to high-performance graphical capabilities.
4. 101-key Enhanced Keyboard: The image indeed mentions that the system comes with a "101-key Click-Tactile Enhanced Keyboard," and this description is accurate.
5. 12MHz processor, 1MB memory, and 60MB hard drive: These configurations do align with the information in the image.

## Chart Incorrect Case

### Both U.S. and Germany are positive toward NATO

Favorable views of NATO



Source: Pew Research Center survey conducted in the U.S. Sept. 11-16, 2018, Q3. German results from Körber-Stiftung survey conducted Sept. 13-26, 2018. Trend data from the Global Attitudes Survey.

PEW RESEARCH CENTER

**Original Question:** What's the rightmost value of green graph? Answer the question using a single word or phrase

**Original Answer:** 64


**Incorrect Rewrite Question:** Calculate the average favorability rating for the U.S. and Germany from 2009 to 2018. Compare the results and discuss the significance of any differences.

**Incorrect Rewrite Answer:** To calculate the average favorability rating for the U.S., we add up the values for each year and divide by the number of years:  $(73 + 54 + 49 + 64) / 4 = 240 / 4 = 60\%$  For Germany, we follow the same process:  $(60 + 65 + 55 + 67) / 4 = 247 / 4 = 61.75\%$  The average favorability rating for the U.S. is 60%, while for Germany, it is 61.75%. The difference is relatively small, indicating that both countries have similar levels of support for NATO over the years. However, the slightly higher average for Germany could suggest a more consistent and stable support for NATO compared to the U.S. This difference could be significant in the context of NATO's decision-making process, where member countries' opinions and support play a crucial role. A more consistent support from Germany could influence the alliance's policies and strategies, potentially leading to a stronger and more unified NATO.

**Analysis:** Here is a detailed analysis of the errors:

1. Calculation Error for Averages: - For the U.S., the values in the graph are 73, 57, 49, and 64, but the calculation uses 73, 54, 49, and 64. This mismatch leads to an incorrect average. The actual average for the U.S. should be  $(73 + 57 + 49 + 64) / 4 = 243 / 4 = 60.75\%$ . - For Germany, the values in the graph are 60, 59, 55, and 67, but the calculation uses 60, 65, 55, and 67. The actual average for Germany should be  $(60 + 59 + 55 + 67) / 4 = 241 / 4 = 60.25\%$ .
2. Incorrect Comparison: - The answer claims that the U.S. average is 60% and Germany's average is 61.75%, but the recalculated averages (60.75% for the U.S. and 60.25% for Germany) reveal the opposite trend. The U.S. has a slightly higher average, not Germany.
3. Misinterpretation of Significance: - The conclusion states that Germany has "more consistent and stable support for NATO" compared to the U.S., based on the incorrect averages. However, the actual data shows that the averages are nearly the same, and there is no significant evidence of Germany having more consistent support than the U.S. The fluctuations in favorability ratings over the years are similar for both countries.
4. Failure to Acknowledge Trend Differences: - The analysis does not highlight the differing trends in the graph. For example, U.S. favorability starts high, dips, and then rises again, while Germany's favorability remains relatively stable but does not show the same sharp recovery seen in the U.S. This is a critical observation that the answer omits.
5. Use of Rounded Averages: - The response uses rounded averages (e.g., 60% for the U.S. and 61.75% for Germany) without clarifying that this is an approximation, which can lead to a misrepresentation of the actual differences.

## Math Incorrect Case



**Original Question:** Hint: Please answer the question and provide the correct option letter, e.g., A, B, C, D, at the end. Question: Find the measure of angle 3. Choices: (A) 28 (B) 32 (C) 38 (D) 42

**Original Answer:** The answer is D.

**Incorrect Rewrite Question:** Hint: Please answer the question and provide the correct option letter, e.g., A, B, C, D, at the end. Question: Find the measure of angle 3. Choices: (A) 28 (B) 32 (C) 38 (D) 42

**Incorrect Rewrite Answer:** To find the measure of  $\angle 3$ , we need to use the properties of triangles and the fact that the sum of the angles in a triangle is always  $180^\circ$ . 1. First, consider the triangle on the left with angles  $28^\circ$ ,  $82^\circ$ , and  $\angle 1$ . - The sum of the angles in this triangle is  $180^\circ$ . - Therefore, we can write the equation:

$$28^\circ + 82^\circ + \angle 1 = 180^\circ$$

- Solving for  $\angle 1$ :

$$\angle 1 = 180^\circ - 28^\circ - 82^\circ = 70^\circ$$

2. Next, consider the triangle on the right with angles  $68^\circ$ ,  $\angle 2$ , and  $\angle 3$ . - The sum of the angles in this triangle is also  $180^\circ$ . - Therefore, we can write the equation:

$$68^\circ + \angle 2 + \angle 3 = 180^\circ$$

3. Notice that  $\angle 1$  and  $\angle 2$  are supplementary because they form a straight line. - Therefore:

$$\angle 1 + \angle 2 = 180^\circ$$

- Substituting the value of  $\angle 1$  we found earlier:

$$70^\circ + \angle 2 = 180^\circ$$

- Solving for  $\angle 2$ :

$$\angle 2 = 180^\circ - 70^\circ = 110^\circ$$

4. Now, substitute  $\angle 2$  back into the equation for the sum of angles in the triangle on the right:

$$68^\circ + 110^\circ + \angle 3 = 180^\circ$$

- Solving for  $\angle 3$ :

$$\angle 3 = 180^\circ - 68^\circ - 110^\circ = 42^\circ$$

Therefore, the measure of  $\angle 3$  is  $42^\circ$ , which corresponds to option (D). Revised Answer: The answer is D.

**Analysis:** The key mistake in the reasoning lies in the incorrect assumption about the relationship between  $\angle 1$  and  $\angle 2$ . Here is the analysis:

**Incorrect Assumption:**

It was assumed that  $\angle 1 + \angle 2 = 180^\circ$ , implying that  $\angle 1$  and  $\angle 2$  are supplementary angles.

This is incorrect because  $\angle 1$  and  $\angle 2$  are actually vertical (opposite) angles, as they are formed by the intersection of two straight lines.

## Caption Incorrect Case



**Original Caption:** The image displays a white, protective armband designed to hold a smartphone. Specifically, the smartphone inserted into the armband appears to be a Samsung device. The design of the armband is structured to be worn around the upper arm, probably intended for use during physical activities such as running or working out. The front section of the armband is clear, allowing for visibility and usage of the smartphone's touchscreen. Here are the detailed components and features observed in the image: 1. Smartphone: 1) Brand and Model: The phone is from the Samsung brand, identifiable by the name imprinted on its front. The exact model is not explicitly stated, but its design characteristics suggest it could be from the Galaxy series. 2) Time Display: The phone's display shows the time as 12:45. This time might be relevant to showcasing the phone's screen capabilities. 3) Date Display: The date shown is Wednesday, September 3. This could either be a static image used for promotional purposes or a real-time capture. 4) UI Elements: The display shows typical smartphone UI elements, including the status bar with signal strength, WiFi connectivity, battery percentage, and possibly notifications. 2. Armband: 1) Material and Color: The armband material appears to be a blend of fabric and synthetic materials, predominantly in white color. 2) Clear Front Cover: The front cover of the armband is clear, providing visibility and functionality to the smartphone's screen. 3) Strap and Buckle: To the side, it has an adjustable strap with a buckle mechanism. The strap appears to be made of black elastic material, ensuring that the armband can fit various arm sizes securely. The buckle is metallic, providing a durable mechanism for securing the armband. 3. Design and Ergonomics: 1) Stitching and Build Quality: The stitching around the armband looks precise and robust, indicating good build quality aimed at withstanding regular usage. 2) Port Accessibility: The design does not obstruct the smartphone's various ports and buttons, ensuring that users can access volume controls, home buttons, or charging ports without removing the phone from the armband. 4. Functional Aspects: 1) Purpose: The primary function of the armband is to hold the smartphone securely during physical activities. The clear cover allows users to interact with their device without removing it. 2) Practical Features: The adjustable strap ensures a snug fit for different arm sizes, while the metallic buckle provides added security to prevent the armband from coming loose.

**Incorrect Rewrite Question:** Based on market trends for smartphone accessories and the functionalities observed in the armband, analyze the potential target market for this product and predict its sales outlook in the next quarter, assuming an MSRP of \$29.99 and availability through major online retailers.

**Incorrect Rewrite Answer:** Analyzing the market trends for smartphone accessories reveals a growing demand for hands-free solutions that cater to active lifestyles. Given the armband's functionalities, such as its adjustable strap, clear front cover, and secure buckle mechanism, it is poised to capture a significant share of the market. Target demographics likely include fitness enthusiasts, athletes, and individuals who enjoy outdoor activities. Assuming the MSRP of \$29.99, a competitive price point compared to similar products, and considering the armband's availability through major online retailers, we can predict the sales outlook for the next quarter. Historical sales data from similar products suggest a strong demand during the first quarter of the year, particularly in January and February, as consumers adopt New Year's resolutions related to fitness and wellness. Projecting a conservative growth rate of 10% month-over-month and accounting for potential seasonality, we estimate the following sales figures: Month 1: 10,000 units, Month 2: 11,000 units, Month 3: 12,100 units. These numbers translate to quarterly revenue of approximately \$348,000, indicating a promising sales outlook for the armband in the next quarter.


**Analysis:** Generated questions have very little relevance to the image.

The examples show the comparison of cases before and after incorrect rewriting. The exceptions that occur during the rewriting process include generating information that is not in the image, incorrectly recognizing of information in the image, making errors in the mathematical logic reasoning process, and generating QA that is not related to the image.

## F.2 Good Cases

Analysis: Cases show the comparison of cases before and after good rewriting. Good rewrites bring more complex and diverse instructions, contain more knowledge, and expand the current monotonous datasets.

### OCR Good Case



**Original Question:** what year is this?

**Original Answer:** 1962-63

**Rewrite Question:** According to the poster, what type of cruise experience can travelers expect during the Winter-Spring season of 1962-63, and how is this experience reflected in the visual elements of the advertisement?

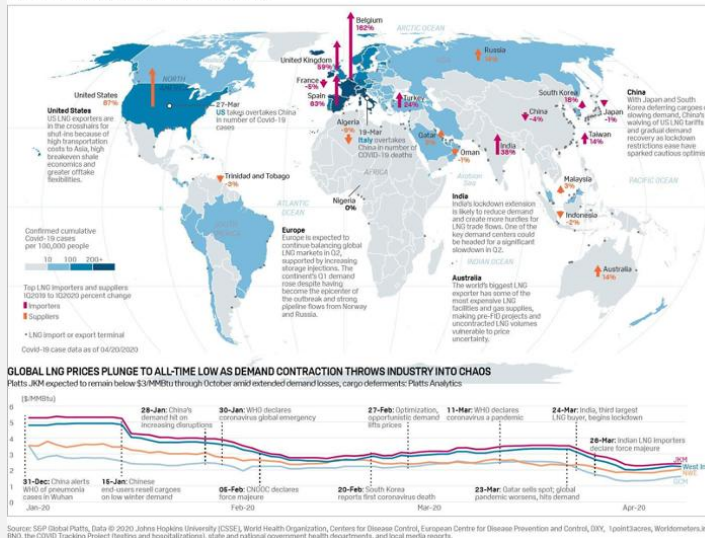
**Rewrite Answer:** Travelers can expect a vibrant and exotic getaway to the West Indies and South America during the Winter-Spring season of 1962-63, as promised by Seven Sun-Way Cruises. The visual elements of the poster effectively convey this experience by featuring three individuals in colorful traditional attire, playing musical instruments on a sunny beach setting. The use of tropical colors, palm trees, and a cruise ship on the horizon suggests a fun, lively, and adventurous atmosphere, which aligns with the idea of a warm and exotic getaway. The poster's focus on musicians and a picturesque seascape implies that the cruise experience will be filled with entertainment, relaxation, and opportunities to soak up the sun.



## Chart Good Case

### GLOBAL LNG MARKETS BRACE FOR UNPRECEDENTED SHOCK IN 2020 ON COVID-19

The spread of coronavirus across the world has generated an unprecedented global health and economic crisis, and presented the LNG industry with a demand shock like no other in its history. The outbreak has hit global LNG demand, shifted contract disputes, disrupted trade flows and derailed project investment plans amid uncertainty over the length and depth of the crisis. With cargo deferrals and cancellations leading to a flurry of spot supply tenders and widespread national lockdowns offering little demand support on the horizon, curtailments at production facilities are seen as the only way to help rebalance demand and supply.




**Original Question:** On which day WHO declared Covid 19 global emergency?

**Original Answer:** 30-Jan

**Rewrite Question:** Analyze the global LNG market trends depicted in the image and discuss the potential long-term impacts of the COVID-19 pandemic on LNG demand and supply. Consider the role of major LNG importers and exporters, and how their strategies might evolve in response to the pandemic-induced market changes.

**Rewrite Answer:** The image illustrates the significant disruption caused by the COVID-19 pandemic in the global LNG market. The pandemic has led to a demand shock, with many countries experiencing reduced economic activity and lower energy consumption. This has resulted in a flurry of spot supply tenders and widespread national lockdowns, which have further exacerbated the demand contraction. Major LNG importers such as Japan, South Korea, and China have seen significant reductions in LNG demand due to industrial slowdowns and decreased energy requirements. For instance, Japan and South Korea have deferred cargoes, while China has waived US LNG tariffs to support its economic recovery. These actions reflect the immediate response to the pandemic's impact on demand. On the supply side, major LNG exporters like Qatar, Australia, and the United States have had to navigate the challenges posed by reduced demand and lower prices. Qatar, the world's largest LNG exporter, has had to sell spot cargoes as the pandemic worsens, highlighting the need for flexibility in supply contracts. In the long term, the pandemic is likely to accelerate the shift towards more flexible LNG supply contracts, as both buyers and sellers seek to mitigate risks associated with demand uncertainty. Additionally, the pandemic may spur investments in LNG infrastructure in regions with growing demand, such as Southeast Asia and Latin America, as countries seek to diversify their energy sources and enhance energy security. Overall, the COVID-19 pandemic has underscored the importance of agility and adaptability in the LNG market. As the global economy recovers, LNG demand is expected to rebound, but the industry will need to navigate a new landscape characterized by evolving market dynamics and changing energy consumption patterns.



**Original Question:** What is the measure of angle QPS?

**Original Answer:** The answer is 61

**Rewrite Question:** What is the measure of angle QPS?

**Rewrite Answer:** To find the measure of angle  $\angle QPS$ , we need to use the properties of the quadrilateral  $PQRS$ . The sum of the interior angles of any quadrilateral is  $360^\circ$ .

Given the angles:

$$\angle P = x^\circ,$$

$$\angle Q = (2x - 16)^\circ,$$

$$\angle R = 2x^\circ,$$

$$\angle S = (x + 10)^\circ,$$

we can set up the equation for the sum of the interior angles:

$$x + (2x - 16) + 2x + (x + 10) = 360.$$

Combine like terms:

$$x + 2x + 2x + x - 16 + 10 = 360,$$

$$6x - 6 = 360.$$

Add 6 to both sides:

$$6x = 366.$$

Divide by 6:

$$x = 61.$$

Now, we need to find the measure of angle  $\angle QPS$ , which is given as  $x$ :

$$x = 61^\circ.$$

Therefore, the measure of angle  $\angle QPS$  is  $61^\circ$ .

## Caption Good Case



**Caption:** The image captures a group of four musicians, who appear to be performing indoors. The setting suggests a relaxed and intimate atmosphere, possibly a small venue or a casual gathering. The background features a wooden plank wall, giving the environment a rustic and cozy feel. To the upper left, there is a modern hanging light fixture that consists of several exposed bulbs suspended from cords, adding a touch of warmth to the scene. Starting from the left, the first musician is playing a violin. He is dressed in a green shirt and appears to be focused on his instrument. Next to him, slightly to the right and in the foreground, is a female vocalist. She has long hair, which is tied back, and she is wearing a black and orange striped sweater. She is singing into a microphone held by a stand. To her right is another musician who is playing a string instrument, possibly a banjo or a mandolin. He has a beard and is dressed casually in an earthy-toned shirt, and he is standing a bit further back compared to the others. Finally, on the far right is a man playing an acoustic guitar. He is wearing a beanie and a casual jacket. He is also singing into a microphone, contributing to the vocal harmony. Near his microphone, a small portion of a red and green circular sign can be seen, which might be part of the venue's signage or decor. On the far right, there is a drum set, though no drummer is visible in the image, suggesting that the band may have more members not captured in this particular frame or they are performing without a drummer at this moment. The musicians appear to be deeply engaged in their performance, with expressions suggesting concentration and genuine involvement in their music. The use of acoustic instruments and their casual attire suggest that their musical style might be folk or indie. In summary, the image captures a quartet of musicians performing with a violin, a banjo or mandolin, an acoustic guitar, and vocals, in a setting that conveys a warm, intimate, and possibly rustic atmosphere.

**Rewrite Question:** Determine the artist's possible musical genre, based on the instruments, attire, and atmospheric elements present in the scene. What supporting visual evidence justifies this inference?


**Rewrite Answer:** Based on the visual elements, it appears that the musical genre is likely folk or indie. The use of acoustic instruments, such as the violin, banjo or mandolin, and acoustic guitar, suggests a focus on traditional or roots music. The casual attire and earthy-toned shirts worn by the musicians also support this inference, implying a relaxed, non-mainstream approach to music. The wooden plank wall and rustic decor in the background reinforce this interpretation, evoking a sense of warmth and authenticity often associated with folk or indie music. Furthermore, the use of microphones and the absence of electric instruments suggest a focus on acoustic performance, which is consistent with these genres.

## G Model Capabilities and Qualitative Examples

In [Appendix G](#), we present a diverse collection of question-and-answer pairs drawn from real-world scenarios to evaluate our model's practical capabilities. These examples showcase the model's ability to handle complex, unconstrained queries while providing meaningful, contextually appropriate responses. The analysis of these Q&A pairs reveals not only the model's comprehension and response generation capabilities but also offers concrete evidence of its real-world applicability. Through these examples, we gain valuable insights into both the model's current strengths and areas for potential improvement in future iterations.

## Comparison with Llava-OneVision-7B Case 1

No. 3: Two runners  $A$  and  $B$  start at the origin and run along the positive  $x$ -axis, with  $B$  running 3 times as fast as  $A$ . An observer, standing one unit above the origin, keeps  $A$  and  $B$  in view. What is the maximum angle of sight  $\theta$  between the observer's view of  $A$  and  $B$ ? (See Fig. 16-24.) [image 1]



Options:

- |                |                |
|----------------|----------------|
| (A) $60^\circ$ | (B) $40^\circ$ |
| (C) $20^\circ$ | (D) $55^\circ$ |
| (E) $35^\circ$ | (F) $25^\circ$ |
| (G) $30^\circ$ | (H) $50^\circ$ |
| (J) $15^\circ$ | (I) $45^\circ$ |

**MAMmoTH-VL-8B(✓)**: To find the maximum angle of sight  $\theta$  between the observer's view of  $A$  and  $B$ , we need to analyze the geometry of the situation.

1. Understanding the Problem: - Runner  $A$  starts at the origin and runs along the positive  $x$ -axis. - Runner  $B$  runs 3 times as fast as  $A$  and also starts from the origin. - The observer is standing one unit above the origin.

2. Setting Up the Geometry: - Let's denote the position of runner  $A$  at time  $t$  as  $A(t)$  and the position of runner  $B$  as  $B(t)$ . - Since  $B$  runs 3 times as fast as  $A$ , if  $A$  runs a distance  $d$  in time  $t$ ,  $B$  will run  $3d$  in the same time  $t$ .

3. Finding the Coordinates: - At any time  $t$ , the coordinates of  $A$  are  $(t, 0)$ . - The coordinates of  $B$  are  $(3t, 0)$ .

4. Calculating the Angles: - The angle of sight  $\theta$  is the difference between the angles formed by the lines connecting the observer to  $A$  and the observer to  $B$ . - The angle formed by the line connecting the observer to  $A$  is  $\alpha = \tan^{-1}\left(\frac{1}{t}\right)$ . - The angle formed by the line connecting the observer to  $B$  is  $\beta = \tan^{-1}\left(\frac{1}{3t}\right)$ .

5. Maximizing the Angle of Sight: - To maximize  $\theta$ , we need to find the maximum value of  $\alpha - \beta$ . - Using calculus, we can find the critical points by taking the derivative of  $\alpha - \beta$  with respect to  $t$  and setting it to zero.

6. Solving for  $t$ : - After solving the equation, we find that the maximum angle of sight occurs when  $t = \frac{1}{\sqrt{3}}$ .

7. Calculating the Maximum Angle:

- Substituting  $t = \frac{1}{\sqrt{3}}$  into the expressions for  $\alpha$  and  $\beta$ , we get:

$$\alpha = \tan^{-1}(\sqrt{3}) = 60^\circ$$

$$\beta = \tan^{-1}\left(\frac{\sqrt{3}}{3}\right) = 30^\circ$$

- Therefore, the maximum angle of sight  $\theta$  is:


$$\theta = \alpha - \beta = 60^\circ - 30^\circ = 30^\circ$$

8. Matching with Options: The closest option to  $30^\circ$  is (G)  $30^\circ$ . Thus, the maximum angle of sight between the observer's view of  $A$  and  $B$  is  $30^\circ$ . Answer: (G)

**Llava-OneVision-7B(X)**: The correct answer is D

## Comparison with Llava-OneVision-7B Case2

No. 25: Which of the following best explains the overall trend shown in the [image 1] ?



Options:

- |  |   |
|--|---|
| (A) Political instability leading to population decline  | (B) The spread of pathogens across the Silk Road    |
| (C) Development of new trade routes                      | (D) Climate change affecting the Silk Road          |
| (E) Migrations to areas of Central Asia for resettlement | (F) Technological advancements in transportation    |
| (G) Invasions by Mongol tribes                           | (H) Large-scale famine due to crop failures         |
| (I) Economic prosperity and population growth            | (J) Rise of religious conflicts along the Silk Road |

Previous Next

1 / 170 Go

**MAMmoTH-VL-8B(✓)**: To determine the best explanation for the overall trend shown in the graph, we need to analyze the population trends of the Roman and Han empires over time.

1. Identify the Trends: The graph shows two lines representing the population of the Roman and Han empires. The Roman line starts at a higher population and declines over time. The Han line starts at a lower population and also declines over time.

2. Evaluate the Options:

(A) Political instability leading to population decline: This option suggests that political instability could lead to a decline in population. However, it does not explain why the Roman population started higher and declined more significantly than the Han population.

(B) The spread of pathogens across the Silk Road: This option implies that the spread of diseases could cause population decline. While this could be a contributing factor, it does not fully explain the observed trends.

(C) Development of new trade routes: This option suggests that the development of new trade routes could lead to population growth or stability. However, the graph shows a decline in population for both empires, which contradicts this option.

(D) Climate change affecting the Silk Road: This option indicates that environmental factors could impact population trends. While climate change could contribute to population decline, it does not fully explain the specific trends observed in the graph.

(E) Migrations to areas of Central Asia for resettlement: This option suggests that migrations could lead to population changes. However, it does not explain why the Roman population started higher and declined more significantly than the Han population. ....

3. Conclusion:

- Based on the analysis, none of the options fully explain the observed trends in the graph. However, option (B) The spread of pathogens across the Silk Road is the most plausible explanation, as it suggests that disease could have contributed to the decline in population for both empires. Therefore, the answer is: Answer: B

**Llava-OneVision-7B(✗)**: The overall trend shown in the image is a decline in population. The graph shows that both the Roman and Han populations decreased over time, with the Roman population declining more sharply than the Han population. This decline could be due to political instability leading to population decline, as option (A) suggests. However, without more context or information about the specific events that occurred during this time period, it is difficult to determine the exact cause of the decline. Nonetheless, based on the information provided in the image, the best explanation for the overall trend shown in the image is option (A).

## Comparison with Llava-OneVision-7B Case3

Flashcard sets, textbooks, questions

**No. 21:** [image 1] shows the results from looking at the diagnostic accuracy of a new rapid test for HIV in 100,000 subjects, compared to the Reference standard ELISA test. The rows of the table represent the test result and the columns the true disease status (as confirmed by ELISA). What is the Sensitivity of the new rapid test for HIV? Report the answer to 3 decimal places.

	HIV+	HIV-	Total
Test +	378	397	775
Test -	2	98,823	98,825
	380	99,220	100,000

**Options:**

- (A)0.950
- (B)0.985
- (C)0.975
- (D)0.970
- (E)0.990
- (F)0.980
- (G)0.965
- (H)0.960
- (I)0.995

**MAMmoTH-VL-8B(✓):** To find the sensitivity of the new rapid test for HIV, we need to use the information provided in the table. Sensitivity is defined as the proportion of true positive results out of all the actual positive cases. From the table, we can see the following values:

- True positive (TP): 378
- False negative (FN): 2
- True negative (TN): 98,823
- False positive (FP): 99,220

The formula for sensitivity is:

$$\text{Sensitivity} = \frac{\text{True Positive (TP)}}{\text{True Positive (TP)} + \text{False Negative (FN)}}$$

Plugging in the values from the table:

$$\text{Sensitivity} = \frac{378}{378 + 2} = \frac{378}{380}$$

Now, we calculate the decimal value:

$$\frac{378}{380} \approx 0.995$$

Therefore, the sensitivity of the new rapid test for HIV is approximately 0.995.

Answer: (I)0.995

**Llava-OneVision-7B(✗):** The sensitivity of the new rapid test for HIV is 0.985. Therefore, the answer is option B.  
Answer:B

## Multi-Image and Video-1



**User:** What are the accusations of the five species?

**MAMmoTH-VL-8B:** The grasshopper is an herbivore, feeding on the grass. The mouse is a carnivore, feeding on the grasshopper. The hawk is a carnivore, feeding on the mouse. The grass is a producer, providing energy for the grasshopper. The mushroom is a decomposer, breaking down dead organic matter and recycling nutrients back into the ecosystem.

**User:** Based on the five species in the images, draw a food chain. Explain the role of each species in the food chain.

**MAMmoTH-VL-8B:**

The food chain can be represented as follows:

Grass → Grasshopper → Mouse → Hawk

In this food chain, the grass is the primary producer, converting sunlight into energy through photosynthesis. The grasshopper is a primary consumer that feeds on the grass, obtaining energy from it. The mouse is a secondary consumer that feeds on the grasshopper, obtaining energy from it. The hawk is a tertiary consumer that feeds on the mouse, obtaining energy from it.

Each species plays a crucial role in the food chain. The grass provides energy for the grasshopper, which in turn provides energy for the mouse. The mouse provides energy for the hawk, which is at the top of the food chain. Without each species, the food chain would collapse, and the ecosystem would be disrupted.

## Multi-Image and Video-2



**User:** Please describe the video in detail.

**MAMmoTH-VL-8B:** The video begins with a close-up of two individuals, one wearing a white t-shirt and the other in a sleeveless top with a floral pattern. They are standing in a kitchen with a brick wall background, which includes various kitchen items such as bottles and utensils. The scene transitions to a close-up of hands meticulously arranging strips of bacon on a metal baking tray placed on a wooden cutting board. The hands carefully place each strip, ensuring they are evenly spaced. The focus remains on the preparation process, highlighting the precision involved in arranging the bacon. The video then shows a person placing the tray into an oven, with text appearing on the screen that reads '400°F 10-15 minutes,' indicating the cooking instructions. The final scene returns to the two individuals in the kitchen, now engaged in conversation while standing near a counter filled with various ingredients and kitchen tools, including bottles, vegetables, and a bowl.

## Chart & Doc Understanding-1



**User:** List the weather for Monday to Friday

**MAmmoTH-VL-8B:**

Monday: Sunny

Tuesday: Thunderstorms in the afternoon

Wednesday: Mostly cloudy

Thursday: Partly cloudy

Friday: Rain

## Chart & Doc Understanding-2

This is a lot of 12 point text to test the ocr code and see if it works on all types of file format.

The quick brown dog jumped over the lazy fox. The quick brown dog jumped over the lazy fox. The quick brown dog jumped over the lazy fox. The quick brown dog jumped over the lazy fox.

**User:** Read text in the image

**MAmmoTH-VL-8B:** This is a lot of 12 point text to test the ocr code and see if it works on all types of file format. The quick brown dog jumped over the lazy fox. The quick brown dog jumped over the lazy fox. The quick brown dog jumped over the lazy fox. The quick brown dog jumped over the lazy fox.




## Multi-Discipline Knowledge and Mathematical Reasoning-1



**User:** Answer the following question based on the provided image: What country do these planes belong to?

**MAmmoTH-VL-8B:** The planes in the image belong to the United States. This can be inferred from several clues within the image: the design of the aircraft, which is consistent with American military planes from the World War II era; the presence of the U.S. Air Force roundel on the fuselage of the aircraft on the right; and the style of the hangar, which resembles those used by the U.S. military. Additionally, the informational plaques and the overall presentation suggest that this is a museum dedicated to American aviation history.



**User:** Find the measure of angle 3.

**MammoTH-VL-8B:** To find the measure of angle 3, we need to use the properties of triangles and the fact that the sum of angles in a triangle is always 180 degrees.

1. First, observe the triangle on the left with angles  $28^\circ$  and  $82^\circ$ . - The sum of these two angles is:

$$28^\circ + 82^\circ = 110^\circ$$

2. Since the sum of all angles in a triangle must be  $180^\circ$ , the third angle in this triangle (which is angle 1) can be found by subtracting the sum of the known angles from  $180^\circ$ :

$$180^\circ - 110^\circ = 70^\circ$$

3. Now, consider the triangle on the right with angles  $68^\circ$  and the angle we just found ( $70^\circ$ ). - The sum of these two angles is:

$$68^\circ + 70^\circ = 138^\circ$$

4. Again, since the sum of all angles in a triangle must be  $180^\circ$ , the third angle in this triangle (which is angle 3) can be found by subtracting the sum of the known angles from  $180^\circ$ :

$$180^\circ - 138^\circ = 42^\circ$$

Therefore, the measure of angle 3 is  $42^\circ$ .



User: Please perfectly describe this cartoon illustration in as much detail as possible

**MAMmoTH-VL-8B:** The image is a cartoon illustration featuring two anthropomorphic animal characters. On the left, there's a character resembling a gray wolf with large ears, wearing a white shirt with a black vest and red-striped pants. This character has a surprised expression on its face.

On the right side of the image, there's another character that appears to be a blue fox with a similar anthropomorphic design. This character is standing behind a transparent wall or barrier, which seems to be made of ice or glass. The fox character is wearing blue shoes and has a confident, almost smug expression on its face.

Above the fox character, there's a speech bubble containing text that reads: "Meh... Don't worry about it. I'm a New Englander, so I'm used to it." The background of the image shows a store entrance with a sign that says "WELCOME" in green letters. The floor of the store is tiled, and there are reflections visible on the glass doors.

The style of the illustration is cartoonish with bold outlines and bright colors. The overall tone of the image seems to be humorous, as indicated by the speech bubble and the expressions of the characters.

