

---

# MMT-Bench: A Comprehensive Multimodal Benchmark for Evaluating Large Vision-Language Models Towards Multitask AGI

---

Kaining Ying<sup>\*1</sup> Fanqing Meng<sup>\*21</sup> Jin Wang<sup>\*3</sup> Zhiqian Li<sup>13</sup> Han Lin<sup>21</sup> Yue Yang<sup>21</sup> Hao Zhang<sup>1</sup>  
Wenbo Zhang<sup>4</sup> Yuqi Lin<sup>15</sup> Shuo Liu<sup>1</sup> Jiayi Lei<sup>12</sup> Quanfeng Lu<sup>1</sup> Runjian Chen<sup>13</sup> Peng Xu<sup>13</sup>  
Renrui Zhang<sup>1</sup> Haozhe Zhang<sup>5</sup> Peng Gao<sup>1</sup> Yali Wang<sup>6</sup> Yu Qiao<sup>1</sup> Ping Luo<sup>31</sup> Kaipeng Zhang<sup>†1</sup>  
Wenqi Shao<sup>†1</sup>

## Abstract

Large Vision-Language Models (LVLMs) show significant strides in general-purpose multimodal applications such as visual dialogue and embodied navigation. However, existing multimodal evaluation benchmarks cover a limited number of multimodal tasks testing rudimentary capabilities, falling short in tracking LVLM development. In this study, we present MMT-Bench, a comprehensive benchmark designed to assess LVLMs across massive multimodal tasks requiring expert knowledge and deliberate visual recognition, localization, and reasoning. MMT-Bench comprises 31,325 meticulously curated multi-choice visual questions from various multimodal scenarios such as vehicle driving and embodied navigation, covering 32 core meta-tasks and 162 sub-tasks in multimodal understanding. Due to its extensive task coverage, MMT-Bench enables the evaluation of LVLMs using a task map, facilitating the discovery of in- and out-of-domain tasks. Evaluation results involving 32 LVLMs such as the proprietary GPT-4o, GeminiProVision, and open-sourced InternVL-Chat, underscore the significant challenges posed by MMT-Bench. We anticipate that MMT-Bench will inspire the community to develop next-generation multimodal foundation models aimed at achieving general-purpose multimodal intelligence.

---

<sup>\*</sup>Equal contribution <sup>1</sup>Shanghai Artificial Intelligence Laboratory <sup>2</sup>Shanghai Jiao Tong University <sup>3</sup>The University of Hong Kong <sup>4</sup>The University of Adelaide <sup>5</sup>Zhejiang University <sup>6</sup>Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences. Correspondence to: Wenqi Shao <shaowenqi@pjlab.org.cn>, Kaipeng Zhang <zhangkaipeng@pjlab.org.cn>.

## 1. Introduction

In recent years, Large Vision-Language Models (LVLMs) (Zhang et al., 2023a; Yang et al., 2023b; Liu et al., 2023c) have emerged as powerful tools for advancing artificial intelligence, demonstrating remarkable progress in various domains such as visual dialogue, video analysis and document understanding. Driven by diverse and high-quality instruction fine-tuning data mined from various fields, LVLMs will continue to advance towards multitask AGI (Team, 2023a; Bai et al., 2023). As pointed out in Levels of AGI (Morris et al., 2023), the breadth (generality) of tasks is a fundamental criterion for different levels of AGI. A multitask AGI model can perform a wide range of tasks across different domains with human-like proficiency, which could revolutionize many fields such as personalized education (Latif et al., 2023) and medical diagnosis (Singhal et al., 2023). Therefore, it is crucial to build a comprehensive evaluation benchmark to track multitask AGI development.

However, evaluating LVLMs significantly lags behind their development (Morris et al., 2023; Yue et al., 2023b; Liu et al., 2024b). A line of work attempts to bridge this gap by proposing various multimodal evaluation benchmarks. Examples include LVLM-eHub (Xu et al., 2023), MMBench (Liu et al., 2023d), MME (Fu et al., 2023), and SEED-Bench (Li et al., 2023a), which propose dimensions of multimodal capabilities and corresponding test samples. However, these benchmarks have limited coverage of multimodal tasks while testing rudimentary capabilities like visual recognition and text-scarce OCR. Therefore, they cannot fulfil the requirement of the breadth of tasks (Morris et al., 2023). Moreover, recent LVLMs continue to excel in these benchmarks. For instance, InternLM-XComposer2 (Dong et al., 2024) achieved 2242.7/2800 and 79.6/100 overall performance on MME and MMBench, respectively. Other works, such as MathVista (Lu et al., 2023) and MMMU (Yue et al., 2023a), focus on discipline knowledge understanding and reasoning but are constrained to visual questions with scientific diagram images, limiting their breadth for benchmarking multitask AGI.

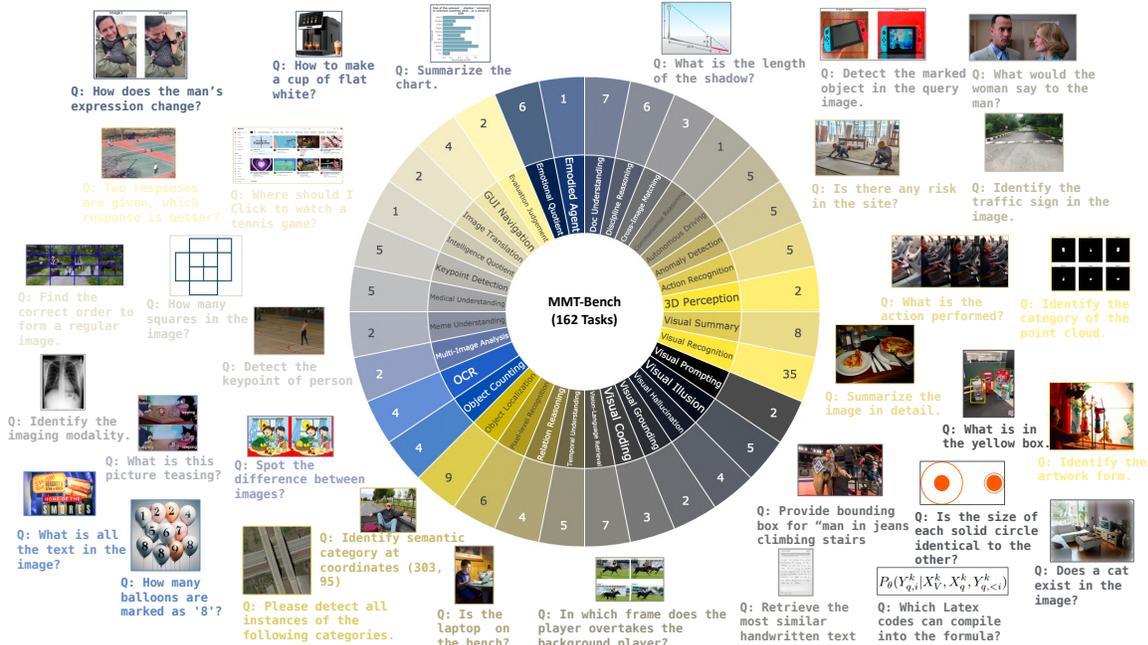


Figure 1. Visualization of MMT-Bench. Our MMT-Bench consists of 32 meta-tasks (middle ring) which are decomposed into 162 subtasks (outer ring). For each meta-task, we denote the number of subtasks in it and illustrate one example of the pair of the image and the question (see task hierarchy in Table A3 to Table A5 of Appendix). MMT-Bench can be comprehensive enough to evaluate the multitask performance of LVLMs.

To address this challenge, we introduce MMT-Bench, a new benchmark designed to comprehensively assess LVLMs in multimodal multitask understanding. The **breadth** of MMT-Bench features in three aspects. First, MMT-Bench is meticulously curated and comprises 32K multi-choice visual questions covering 32 core meta-tasks and a total of 162 subtasks (Fig. 1), which is 8.1 times larger than MMBench (Liu et al., 2023d). Second, it encompasses 13 image types such as natural scenes, synthetic images, depth maps, text-rich images, paintings, screenshots, point clouds, medical images, et al. (Fig. 2). Such diversity demands the model to be capable enough to interpret various visual inputs. Third, MMT-Bench spans multimodal scenarios such as vehicle driving, GUI navigation, and embodied AI, testing 14 kinds of multimodal capabilities including visual recognition, localization, reasoning, OCR, counting, 3D perception, temporal understanding, et al. (Fig. 2).

We assess 32 publicly available LVLMs under various input modes for best evaluation performance. Our findings highlight the significant challenges posed by MMT-Bench. For instance, the best performing GPT-4o (OpenAI, 2024) only achieves 65.5/100 and 59.5/100 overall scores across all subtasks and subtasks except for visual recognition tasks, respectively, indicating significant room for improvement towards multitask AGI. Thanks to the extensive coverage of multimodal tasks, MMT-Bench enables the evaluation of LVLMs using a task map. This facilitates the discovery of both in- and out-of-domain tasks, providing valuable in-

sights for multimodal commercial applications and ongoing efforts to enhance LVLMs. We summarize the findings as follows:

- GPT-4o leads other SOTA LVLMs with an absolute advantage. Meanwhile, the open-source LVLm InternVL-Chat-v1.2 (Chen et al., 2023b; 2024) ranks as the runner-up, outperforming other open-source and closed-source methods.
- The comprehensive error analyses conducted on 162 multimodal tasks reveal that top-performing LVLMs such as InternVL-Chat, GPT4V, and GeminiProVision are predominantly prone to perception, reasoning, and knowledge errors.
- The taskonomy analysis shows that current LVLMs perform well in tasks related to visual recognition and description which are in-domain tasks, yet fall short in tasks related to localization and pixel perception which are out-of-domain tasks.
- BLIP2 (Li et al., 2023b) that does not undergo instruction tuning even outperforms most LVLMs that are tuned by millions of instruction-following data, implying that instruction-tuning with data in some tasks even hurts the generalization on other tasks.
- Certain tasks show improved performance with specific prompting methods, such as multi-image and

Table 1. The comparison between MMT-Bench and existing evaluation benchmarks. MMT-Bench consists of massive samples and multimodal tasks compared with other benchmarks. I, T, V, and P respectively represent image, text, video, and point cloud.

Benchmark	Data Collection					
	# Sample	# Meta-task	# Task	# Modality	Source	Answer Type
SEED-Bench (Li et al., 2023a)	19K	12	12	I + T + V	Annotated	Multi-Choice
MMBench (Liu et al., 2023d)	3K	2	20	I + T	Repurposed	Multi-Choice
MM-VET (Yu et al., 2023)	0.2K	6	N/A	I + T	Repurposed	Multi-Choice
MMMU (Yue et al., 2023b)	11.5K	6	30	I + T	Annotated	Multi-Choice/Open
Tiny LVLm-eHub (Shao et al., 2023)	2.1K	5	42	I + T	Repurposed	Multi-Choice/Open
MMT-Bench	31K	32	162	I + T + V + P	Repurposed	Multi-Choice

coordinate-related tasks, as well as those involving visual referring prompts. However, most models do not exhibit improved performance with visual prompting, suggesting potential areas for future enhancement.

- Model performance significantly improves with an increase in size (7B to 13B) for both LLaVA-v1.5 and LLaVA-v1.5-Xtuner. Upgrading LLMs, from InternLM to InternLM2, also enhances the performance of LLaVA.

Overall, the contributions of this work are three-fold. i) We build a new evaluation benchmark called MMT-Bench for multimodal multitask comprehension, allowing us to measure the progress on the path to multitask AGI. ii) We evaluate various publicly available LVLms on MMT-Bench, revealing that current LVLms including GPT-4o, InternVL-Chat, GPT-4V, and GeminiProVision achieve plain performance in multitask intelligence. iii) We present a taskonomy analysis by evaluating LVLms on a task map built upon MMT-Bench, facilitating the discovery of both in- and out-of-domain tasks relative to current LVLms. We anticipate that MMT-Bench will inspire the community to push the boundaries of LVLm research and development, driving us closer to the realization of truly intelligent multimodal systems. The MMT-Bench is open-sourced at <https://github.com/OpenGVLab/MMT-Bench>.

## 2. Related Work

**LVLm.** As the Large Language Models (LLMs) continue to garner impressive achievements (Bai et al., 2023; Team, 2023b; Touvron et al., 2023a;b; Zheng et al., 2023; Chung et al., 2022), academic emphasis is increasingly shifting towards the exploration and development of Large Visual Language Models (LVLms), to bolster the multimodal understanding and generative capabilities of models. Some notable open-source LVLms, such as mPLUG-Owl2 (Ye et al., 2023b), LLaVA (Liu et al., 2023c), and LLaMA-Adapter (Gao et al., 2023; Zhang et al., 2023b), have adopted LLMs as their backbone, processing visual features through these LLMs, ultimately achieving an innovative integration of text

and visuals. In addition, closed-source models like Gemini (Team, 2023a) and GPT-4V (Yang et al., 2023c) have demonstrated remarkable results across numerous tasks, making groundbreaking contributions. We aim to undertake an in-depth and comprehensive exploration of LVLms and their capabilities by testing them on massive multimodal tasks.

**LVLm Evaluation.** Recently, LVLms have demonstrated remarkable capabilities to handle many visual-language tasks, which makes previous single-task benchmarks (Antol et al., 2015; Hudson & Manning, 2019; Krishna et al., 2017b; Lin et al., 2014b; Marino et al., 2019) insufficient to provide comprehensive evaluations of current LVLms. To this end, current LVLm evaluation benchmarks aimed to provide relatively holistic evaluations for the overall reasoning capabilities of LVLms, such as OwlEval (Ye et al., 2023a), LVLm-eHub (Xu et al., 2023), SEED-Bench (Li et al., 2023a), LAMM (Yin et al., 2023), MM-Vet (Yu et al., 2023) and MMBench (Liu et al., 2023d). However, these benchmarks only covered a small range of multimodal tasks and vision-language skills, making them not comprehensive enough to assess multitask AGI capabilities. Besides, recent studies also presented benchmarks of LVLms which required expert-level domain knowledge, such as Mathvista (Lu et al., 2023) and MMMU (Yue et al., 2023a). In comparison, our proposed MMT-Bench covers an extensive range of multimodal reasoning capabilities with sufficient test samples from various modalities as shown in Table 1, which requires expert knowledge and deliberate visual recognition, localization, reasoning, and planning. Our MMT-Bench poses significant challenges for the current state-of-the-art LVLms.

**Multitask Analysis.** Characterizing various tasks and establishing inter-task relationships is an effective means for multitask analysis (Ilharco et al., 2023; Achille et al., 2019; Zamir et al., 2018; Wallace et al., 2021), with wide applications in areas such as meta-learning and transfer learning. A substantial amount of research has been conducted in Taskonomy (Zamir et al., 2018). It utilizes transfer learning to model the structure of the space of visual tasks, thereby harnessing the interconnections among visual tasks to avoid

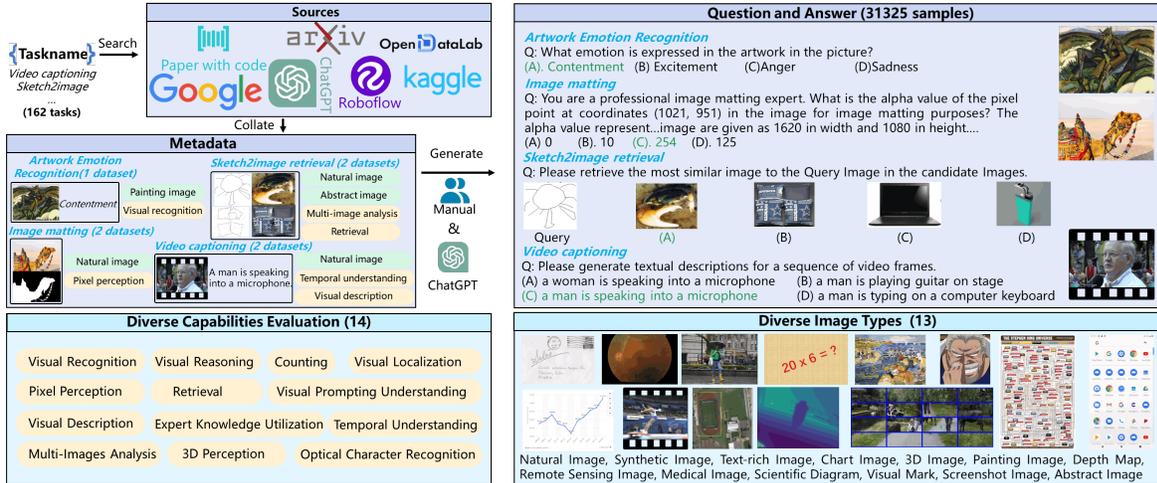


Figure 2. An illustration of our pipeline for data collection. First, given a task name, we retrieve its related datasets from the internet. Then we collate them in a uniform data format - metadata. Finally, we generate questions with choices and answers from metadata using manually designed rules or ChatGPT. Our benchmarks cover capabilities evaluation with diverse image types.

redundancy in learning. Task2Vec (Achille et al., 2019) extracts fisher information as task vectors, which is used in meta-learning. In our paper, thanks to the vast amount of task data collected, we evaluate LVLMs on a task map and conclude challenging tasks for the current LVLMs.

### 3. MMT-Bench

#### 3.1. Hierarchical Task Structure

We utilize a hierarchical structure to include as more as multimodal tasks to build the MMT-Bench. First, all co-authors come up with meta-tasks for multimodal understanding by brainstorming. We then collect 32 meta-tasks by deduplication and filtering for important tasks as depicted in Fig. 1. Second, we decompose each meta-task into several subtasks. The subtask is kept in the MMT-Bench by three criteria. i) Whether the subtask examines the basic multimodal capability. ii) Whether the subtask challenges the current LVLMs. iii) Whether the test sample for the subtask can be publicly accessible. After selection, MMT-Bench comprises 162 subtasks, which is 3.8 times larger than TinyLVLM-eHub which previously contained the most tasks (Shao et al., 2023). The detailed comparison between MMT-Bench and previous benchmarks is provided in Table 1. We also present the whole hierarchical structure in Table A3 of the Appendix.

#### 3.2. Data Collection

We design an efficient pipeline (see Fig. 2) to construct multi-choice visual questions evaluation data for each subtask and the data collection is completed by dozens of co-authors specializing in artificial intelligence.

**Datasets Search.** We conduct comprehensive searches for related datasets using various sources such as Google, Paper With Code, Kaggle, and ChatGPT, based on the name of the subtask. After downloading the datasets, we meticulously assess their suitability for evaluating the subtask, ensuring usability and relevance. While most tasks have multiple datasets available, a few may only have one dataset publicly accessible.

**Metadata Construction.** We define a uniform format, the metadata, to collate downloaded datasets. It enables the further generation of visual questions and answers. Each sample of metadata consists of images and meta-information. The meta-information (see Fig. 2) includes the necessary information to generate questions and answers for the evaluation and also includes manual annotations of required capabilities and the type of visual prompt (i.e., input image). For evaluation efficiency, in each task, we keep the maximum number of samples at 200 by random sampling, and each dataset comprises the same number of samples.

**Question and Answer Generation.** For each subtask, we generate multi-choice (maximum eight choices depending on the task) visual questions with choices and answers from their metadata. Specifically, depending on a specific task, we manually design rules or use ChatGPT with well-designed prompts for efficient and high-quality generation. For example, in sketch2image retrieval, we use the corresponding image as a ground-truth answer and generate other choices by randomly sampling other images from metadata. In video captioning, we use ChatGPT to write confused wrong choices.

**Dataset Statistics.** MMT-Bench comprises 31, 325 meticulously curated multi-choice questions with 13 input image

types such as natural scenes, synthetic images, text-rich images, medical images, et al. (see Fig. 2), covering 32 core meta-tasks and 162 subtasks for multitask multimodal understanding. Compared to previous LVLMs benchmarks (Yue et al., 2023a; Xu et al., 2023) addressing limited image types and skills, questions in MMT-Bench span diverse multimodal scenarios such as GUI navigation and document understanding, testing 14 kinds of capabilities including visual recognition, localization, reasoning, OCR, counting, 3D perception, temporal understanding, et al., as shown in Fig. 2. These features ensure that MMT-Bench meets the requirement of task breadth for evaluating multitask AGI.

## 4. Experiments

In this section, we conduct a comprehensive evaluation of 32 LVLMs on the MMT-Bench. Sec. 4.1 presents the selected LVLMs zoo and the evaluation methods. The quantitative evaluation of each meta-task is provided in Sec. 4.2. We present the analysis of specific tasks with different prompt methods in Sec. 4.3. Furthermore, we give an error analysis of three representative LVLMs in Sec. 4.4.

### 4.1. Evaluation Details

**Selected LVLMs.** For completeness, we assess 32 representative LVLMs, including closed source LVLMs (GPT-4o, GPT-4o, QWen-VL-Max, QWen-VL-Plus, etc.) and open source LVLMs varying in parameters, vision encoders (InternVL (Chen et al., 2023b), EVA-CLIP-ViT (Sun et al., 2023), CLIP-ViT (Radford et al., 2021)), and LLMs (QWen (Bai et al., 2023), InternLM (Team, 2023b), LLaMA (Touvron et al., 2023a;b), Vicuna (Zheng et al., 2023), Flan-T5 (Chung et al., 2022)). For details, see Appendix D.1.

**Evaluation Methods.** In MMT-Bench, samples are in a multi-choice format, e.g., ‘What is this? Options: (A) Dog (B) Cat’. To extract the choice from LVLMs’ responses, we follow OpenCompass’ protocol (Contributors, 2023a): 1) Check if the response includes option letters (A/B); 2) Check for option content (‘dog’/‘cat’); 3) Use ChatGPT for extraction. If these steps fail, we set the model selection as option letter Z to avoid random assignment (Yue et al., 2023a). Accuracy is the primary metric.

### 4.2. Overall Evaluation

This section evaluates LVLMs on MMT-Bench alongside *Random Choice* and *Frequent Choice* baselines. We report the overall score for all meta-tasks as well as the best performance on each meta-task in Table 2. The detailed results of each sub-task are provided in the Sec. L of the Appendix. Various prompt settings for all tasks are investigated. We summarize the key findings as follows.

**i) The Comprehensive Challenge of MMT-Bench:** The

benchmark poses significant challenges, with even advanced models like GPT-4o, InternVL-Chat and GeminiProVision achieving just 65.5%, 63.4% and 61.6% accuracy, respectively, indicating substantial room for improvement. Notably, removing its strongest area, Visual Recognition (VR), where it scores 88.0%, GPT-4o’s overall performance drops to 59.5%, below satisfactory. The varied task dimensions of the MMT-Bench demand wide-ranging capabilities for optimal performance, emphasizing the benchmark’s extensive and rigorous criteria. **ii) The comparison between Open-source LVLMs and close-source LVLMs.** The performance of most open-source models lags behind that of closed-source models. However, leading open-sourced LLM InternVL-Chat-V1.2-34B have demonstrated remarkable performance, outperforming sophisticated proprietary models such as GPT-4V and GeminiProVision in overall accuracy. This achievement suggests that by scaling model size, optimizing training regimes, and leveraging diverse high-quality data, open-sourced LVLMs can rival and even exceed the capabilities of advanced proprietary models. It brings a sense of pride to the open-source community and paves the way for more high-performance yet cost-effective solutions in academia and industry. **iii) The Influence of LLMs and Model Scaling.** As shown in Table 2, model performance significantly improves with an increase in size (7B to 13B) for both llava-v1.5 and llava-v1.5-tuner. Upgrading LLMs, from internlm to internLM2, also enhances the performance of LLaVA, suggesting that larger or improved LLMs boost multi-task performance, with unchanged training data and visual encoders. **iv) Model Performance across Different Meta-Tasks.** Most LVLMs excel in Visual Recognition (VR) tasks and Visual Captioning (VC), highlighting the ability of LVLMs to recognize ‘what’ an object is and describe the content shown in the image. However, for fine-grained perception tasks (localization, pixel-level perception, etc) or complex reasoning tasks (image evaluation judgment), most LVLMs struggle. **v) BLIP2 impresses in open-source models without instruction-following training, outdoing LLaVA models trained with extensive instruction-following data.** Although instruction-tuned models can give responses aligning better with human preference than BLIP2 in open-set QA on some tasks (Liu et al., 2023c), they perform worse than BLIP2 in close-set settings in MMT-Bench. This reflects MMT-Bench’s multi-task challenges and hints at using the taxonomy of MMT-Bench to expand the dataset in supervised fine-tuning for future advancement.

### 4.3. Specific Task and Prompt Methods Analysis

In this section, we evaluate specific tasks using different prompts for LVLMs.

**Prompting LVLMs with multi-images vs single-image.**

Here we explore the effects of exploiting multi-image

MMT-Bench

Table 2. Quantitative results for 32 LVLMs across 32 meta-tasks are summarized. Accuracy is the metric, and the Overall score is computed across all subtasks, excluding visual recognition (VR) as denoted by \*. The maximum value of each meta-task is bolded. Meta-tasks are abbreviated for brevity, with full terms in Sec. C of the appendix.

Model	Overall Overall*	VR VI	Loc MemU	OCR VPU	Count AND	HLN KD	IR VCR	3D IEJ	VC MIA	VG CIM	DU TU	AR VP	PLP MedU	I2IT AUD	RR DKR	IQT EA	Emo GN
Frequency Guess	31.7 32.2	30.0 52.1	28.2 32.8	30.4 29.3	28.2 44.4	43.4 33.7	29.9 27.0	26.5 30.0	28.2 46.5	29.1 28.5	37.6 29.1	30.0 29.5	29.4 30.9	30.8 29.7	33.5 29.4	18.0 28.0	30.1 29.0
Random Guess	28.5 28.9	27.1 50.8	28.1 25.5	27.2 31.4	25.0 36.5	41.6 32.2	24.3 28.0	25.5 25.0	24.8 48.5	24.8 26.8	30.3 27.0	25.4 28.8	26.6 27.8	21.2 26.8	33.4 25.4	10.5 27.5	25.4 24.4
GPT-4o	<b>65.5</b> <b>59.5</b>	<b>88.0</b> 39.1	<b>67.4</b> <b>84.2</b>	66.5 <b>77.0</b>	54.6 51.6	69.4 51.6	47.8 59.9	49.5 43.0	86.4 33.2	50.6 49.2	74.4 38.7	57.0 <b>57.8</b>	<b>46.3</b> 78.3	<b>40.2</b> 54.7	<b>74.6</b> 63.2	15.5 <b>88.5</b>	<b>57.8</b> <b>46.1</b>
InternVL-Chat-v1.2-34B	63.4 58.2	81.3 56.7	59.4 79.5	60.5 60.4	<b>61.6</b> 53.8	<b>76.8</b> 58.2	59.7 86.0	45.5 <b>47.5</b>	82.3 <b>82.8</b>	49.4 56.8	68.3 35.3	52.6 47.3	37.9 71.8	34.0 57.8	54.3 51.2	11.0 84.0	53.2 41.2
QwenVLMax	62.4 56.6	83.0 40.9	54.1 79.2	<b>66.9</b> 54.1	56.9 59.2	74.1 52.2	41.8 81.0	46.0 40.8	86.2 81.8	45.9 60.7	<b>77.9</b> 39.3	53.7 46.8	42.1 74.1	36.8 53.1	61.4 51.9	12.5 85.5	51.5 31.5
Qwen-VL-Plus	62.3 56.6	82.6 41.1	55.3 78.0	65.6 54.6	57.1 <b>59.8</b>	74.4 52.3	39.6 81.5	46.5 39.0	<b>86.5</b> 81.5	43.6 61.3	77.3 38.7	53.4 45.2	42.8 73.3	36.0 59.5	60.2 51.2	11.0 84.5	51.1 32.6
GeminiProVision	61.6 55.1	84.7 47.5	43.6 75.8	59.5 50.9	56.4 47.4	65.9 49.5	<b>68.4</b> <b>86.5</b>	45.2 35.0	80.1 70.2	33.0 33.3	71.6 40.5	<b>57.4</b> <b>46.0</b>	40.3 <b>82.6</b>	31.5 <b>59.5</b>	58.5 49.2	11.0 74.5	55.2 33.4
GPT4V	61.1 55.1	84.0 28.6	56.9 81.2	55.6 66.2	49.9 56.2	68.6 49.6	41.2 82.5	36.8 44.5	83.4 44.5	38.5 67.5	67.8 38.8	55.0 49.3	44.2 69.6	38.2 50.3	66.0 59.8	14.5 85.0	55.1 40.4
LLaVA-NEXT-34B	60.8 56.3	76.7 56.3	61.0 79.0	64.1 63.1	58.3 57.8	75.5 58.9	35.2 83.0	48.5 43.8	85.9 77.8	<b>56.2</b> 43.0	69.1 35.2	50.6 47.7	39.1 61.9	27.5 55.1	61.8 51.2	13.0 76.5	53.2 41.4
XComposer2	55.7 50.0	75.3 52.6	47.9 71.2	43.9 56.1	51.0 56.2	69.5 41.5	32.4 83.0	40.5 43.8	73.7 80.8	42.6 61.2	62.0 36.6	46.3 36.3	43.9 53.5	31.5 48.8	50.5 43.8	8.0 50.5	53.6 29.4
BLIP2	54.8 49.1	75.1 55.6	54.1 76.2	48.1 39.8	29.8 43.7	66.1 <b>60.2</b>	27.4 77.0	47.8 29.8	78.7 62.8	33.5 <b>73.0</b>	43.0 <b>42.7</b>	51.1 43.2	46.1 60.1	28.2 44.6	53.0 77.0	14.0 80.5	43.1 33.4
Yi-VL-34B	54.2 48.6	74.6 43.6	47.0 75.2	58.0 49.4	52.0 56.0	73.6 43.1	27.1 77.5	38.8 38.2	74.0 66.0	41.5 48.0	56.4 30.7	40.4 40.0	38.7 59.6	19.5 48.2	57.2 48.7	14.0 68.5	44.3 32.4
Monkey-Chat	53.4 46.0	79.0 55.3	40.1 69.5	51.0 43.6	43.6 44.6	63.1 36.3	26.8 85.5	46.5 26.0	68.9 58.8	27.5 61.7	51.1 36.8	49.3 33.3	32.2 68.0	29.5 43.6	61.8 38.1	11.0 46.0	45.1 29.8
DeepSeek-VL-7B	53.2 46.5	75.6 44.6	42.0 66.8	61.1 48.9	44.1 46.6	64.2 34.4	27.3 81.0	47.2 35.5	69.1 35.0	38.4 67.2	51.9 30.6	44.8 31.2	41.7 69.7	26.0 48.8	50.1 38.9	12.5 37.0	49.5 36.8
Yi-VL-6B	53.2 47.5	73.5 53.1	49.4 67.8	53.1 43.1	50.5 50.0	70.2 47.9	24.9 82.5	43.5 39.2	63.4 52.0	42.1 43.3	55.2 31.7	43.8 42.5	36.8 60.6	29.5 46.9	54.6 43.2	13.0 47.0	46.8 34.8
LLaVA-NEXT-13B	53.0 46.8	74.0 57.4	35.6 71.2	51.8 48.4	54.5 46.5	70.0 80.5	28.5 30.1	<b>50.0</b> 32.2	75.0 36.5	44.6 59.8	53.6 33.9	46.5 27.0	30.9 55.8	27.0 55.7	54.9 40.7	14.5 50.0	45.8 41.0
TransCore-M	52.7 46.9	73.6 55.6	40.5 76.8	50.4 51.9	54.5 43.7	71.9 38.6	27.5 85.5	45.0 34.2	75.6 52.8	35.1 65.8	45.3 29.7	46.9 28.8	38.3 61.1	25.0 46.5	53.2 38.4	15.0 39.5	46.3 35.6
QWen-VL-Chat	52.5 45.4	77.5 50.9	33.7 74.2	46.9 42.4	46.7 40.2	63.9 35.9	27.5 86.0	45.0 30.0	73.0 49.2	26.5 58.3	51.5 37.3	50.9 30.8	32.7 67.1	30.5 45.4	57.4 35.6	13.5 55.0	45.4 30.2
Claude3V_Haiku	52.2 46.4	74.3 44.6	44.8 40.0	54.4 43.6	48.4 39.3	66.5 46.9	33.4 81.5	38.2 33.5	67.6 34.8	26.9 52.5	69.8 33.7	46.2 45.2	37.1 62.7	23.8 34.1	48.6 44.2	15.5 49.5	42.0 35.1
XComposer	52.1 45.6	75.4 53.4	40.4 63.8	44.1 40.6	39.9 40.6	66.5 43.4	49.7 78.0	47.0 29.0	72.1 66.2	27.2 52.3	36.6 33.1	47.9 28.3	39.6 55.6	24.5 40.8	50.2 39.3	14.0 38.5	45.9 45.2
mPLUG-Owl2	52.0 45.0	76.5 58.5	45.8 59.0	44.5 40.1	47.6 49.4	63.4 32.9	27.6 85.5	45.2 30.0	66.6 55.0	33.0 57.7	42.4 31.9	45.2 27.3	41.6 63.4	25.5 45.5	52.0 38.1	<b>18.0</b> 35.0	42.0 27.8
RBDash-v1-13B	51.8 46.1	72.2 57.1	42.2 67.5	53.6 51.4	51.6 45.7	66.6 33.2	26.3 78.0	40.8 39.0	75.5 32.0	36.9 64.2	48.1 31.6	47.1 25.5	38.3 59.3	22.5 46.3	55.9 38.1	14.0 53.5	43.4 32.4
LLaVA-v1.5-13B	51.7 45.7	73.8 58.1	38.8 66.0	51.8 43.9	55.1 48.3	65.8 31.4	27.2 79.0	39.8 35.8	70.4 28.5	37.4 62.5	45.7 33.3	46.6 27.5	37.6 58.6	28.0 46.6	58.2 39.4	13.5 40.5	45.3 37.5
CogVLM-Chat	51.6 44.2	77.7 52.4	24.7 75.5	48.5 39.8	49.8 43.4	66.0 28.2	26.1 82.0	42.2 28.0	69.8 70.8	28.8 45.8	49.1 35.5	46.3 28.3	33.2 65.9	23.8 44.9	61.6 36.9	14.0 48.0	50.3 29.9
ShareGPT4V-7B	51.5 45.1	74.2 54.5	36.0 70.5	47.8 47.1	50.9 48.2	62.4 26.3	27.8 83.0	45.2 27.8	71.6 38.0	35.4 64.3	47.9 32.1	46.2 30.0	39.2 60.8	21.8 46.1	59.8 38.9	14.0 42.0	44.3 28.9
LLaVA-NEXT-7B	51.1 44.6	<b>73.3</b> 68.8	29.5 42.9	52.0 44.4	52.5 24.4	66.1 79.5	26.1 30.8	43.2 32.5	69.8 67.8	37.0 32.4	49.7 27.8	47.9 30.8	30.8 55.5	19.0 53.5	56.1 39.5	13.5 47.5	47.4 31.6
LLaVA-v1.5-13B-XTuner	51.1 45.1	72.5 54.4	40.7 66.5	46.8 47.9	54.1 52.0	66.5 28.8	26.4 82.0	47.5 39.2	68.8 37.0	35.6 56.8	47.0 28.3	44.2 28.3	38.3 49.1	26.0 44.4	52.4 37.3	14.0 33.5	51.0 40.9
LLaVA-InternLM2-7B	50.8 44.4	73.3 52.3	38.9 62.5	49.5 45.1	51.8 57.2	67.8 35.2	27.7 83.0	49.5 34.2	66.4 55.8	36.9 58.2	37.7 26.8	43.7 18.5	35.1 57.8	14.2 45.1	58.0 33.7	0.0 35.5	51.1 35.2
LLaVA-v1.5-7B-XTuner	50.2 43.9	72.5 60.1	41.1 56.5	46.0 42.6	49.9 47.2	62.1 28.4	26.0 80.5	45.5 32.2	66.4 41.2	35.3 63.2	42.8 29.9	45.8 24.2	42.5 52.5	25.5 43.4	53.9 37.2	11.5 32.0	44.2 30.5
SharedCaptioner	49.9 43.2	72.8 55.1	41.8 53.8	47.8 45.4	46.2 38.3	63.1 33.6	27.0 82.5	44.2 20.2	61.9 57.8	27.0 56.8	39.5 32.6	46.7 28.7	33.5 59.4	25.0 44.7	59.5 38.4	14.5 35.0	39.9 29.6
LLaVA-InternLM-7B	49.7 43.9	70.1 57.5	38.7 58.2	47.6 45.6	46.0 46.5	62.0 33.2	25.5 75.5	42.0 33.0	65.0 57.0	26.5 59.7	43.9 28.0	45.6 27.3	38.3 52.0	25.0 42.2	52.4 38.1	14.0 46.5	47.0 37.6
LLaVA-v1.5-7B	49.5 43.1	72.8 57.6	34.3 70.5	45.0 33.3	47.5 49.1	61.6 31.6	26.1 81.0	44.8 27.8	68.1 37.5	34.0 62.3	40.8 31.7	46.6 27.5	36.0 56.8	22.2 45.1	58.0 35.6	12.5 42.5	42.5 20.4
LLaMA-Adapter-v2-7B	40.4 34.1	62.3 36.4	32.5 40.5	35.0 33.8	30.1 30.4	46.5 34.9	24.1 71.0	33.8 33.2	34.8 42.2	25.2 35.8	30.2 31.1	43.9 25.8	33.1 52.0	18.2 29.1	44.9 32.0	11.0 25.0	36.0 29.9
VisualGLM-6B	38.6 33.9	55.0 28.9	33.1 44.8	33.8 27.1	31.1 34.5	39.2 35.2	26.0 65.0	36.8 28.0	40.5 35.8	31.1 48.2	39.1 30.8	39.2 23.5	32.4 44.0	26.8 26.2	43.8 29.6	14.0 37.5	33.1 21.1

prompts and single-image prompts on the performance of LVLMs. To this end, we summarized 28 tasks in our MMT-Bench, which usually require multiple images as input, such

as image retrieval and video captioning. For multi-images prompting, we first evaluated LVLMs which are inherently designed to support multiple images as input (dubbed *Multi-*

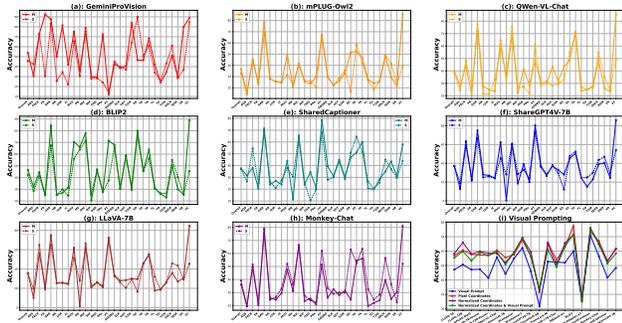


Figure 3. (a)-(h): Comparing the performance of LVLMs between settings of multiple-images prompt (denoted as **M**) and single-image prompt (denoted as **S**). Please check Appendix D.2 for the full task names of task name abbreviations. (i): Comparison of different prompting methods for visual referring prompting-related tasks. Here we select 14 subtasks from the MMT-Bench. We only report the average accuracy here. Zoom in for better view.

*Images LVLMs*), including mPLUG-Owl2, QWen-VL-chat and Gemini-Pro-Vision. Besides, we also assessed LVLMs which mainly learned on single-image prompts (dubbed *Single-Image LVLMs*) for more comprehensive comparisons, including BLIP2, SharedCaptioner, ShareGPT4V-7B, Monkey and LLaVA-v1.5-7B. Following previous studies (Dai et al., 2023; Li et al., 2023c), we input each image individually to Single-Image LVLMs and concatenated all output visual embeddings before feeding into LLMs. The designed multi-image prompts for Multi-Images LVLMs and Single-Image LVLMs are summarized in Appendix Sec.D.2. As for single-image prompting, we manually combine multiple images into one image and feed it into LVLMs (see examples in Fig. 1).

The detailed performance comparisons are presented in Fig. 3(a)-(h). We have several observations: i) Multi-images tasks posed significant challenges to current LVLMs, where the best accuracy achieved by GeminiProVision is only 53.8. ii) For Multi-Images LVLMs, providing multiple images as prompts instead of a single image boosted the overall performance on these tasks, demonstrating their capabilities to extract beneficial information from multiple images. For instance, for the task of face retrieval (**FR**), the performance of GeminiProVision increased from 30.5 to 92.5 when providing multiple images as visual prompts. iii) For Single-Image LVLMs, multi-image prompts also help improve the overall performance of most models, except for Monkey. To our surprise, BLIP2 achieved significant performance gain when switching to a multi-image prompt setting, especially on tasks like general action recognition (**GAR**) and video captioning (**VC**). These results highlight the potential of LVLMs to learn more robust unified representations of multiple modalities.

### Most LVLMs Show Poor Generalization in Visual Re-

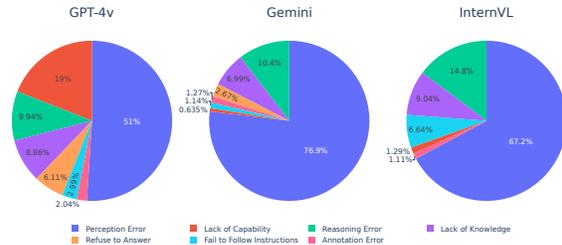


Figure 4. Distribution of error types for GPT-4V, GeminiProVision and InternVL-Chat-V1.2.

**Referring Prompting.** Visual referring prompting is an impressive prompting technique that entails direct image edits like drawing bounding boxes or masks to guide LVLMs to focus on specific regions (Yang et al., 2023b). We select 14 tasks (see Sec. D.3) involving visual referring prompting to explore the influence of different prompting methods on the final results. We compared three additional settings: using text prompts for bounding boxes in normalized  $([0,1])$  and pixel  $([0, h \text{ or } w])$  formats, and combining visual and text prompts. As depicted in Fig. 3(i), visual prompting (blue curve) significantly lags behind other settings, a disparity mainly attributed to the lack of visual prompting data in most LVLMs during the Supervised Fine-Tuning (SFT) stage.

### 4.4. Error Analysis

To analyze the error distribution of LVLMs on the MMT-Bench, we examined three LVLMs: GPT-4V, GeminiProVision, and InternVL-Chat-V1.2 (InternVL). Specifically, we randomly selected up to 5 incorrectly answered questions per subtask for each model. Task-specific experts among the co-authors then analyzed these error samples to identify the underlying reasons for the mistakes, yielding the error distribution presented in Fig. 4. For definitions and case studies of these six error types, please refer to Sec. G in the appendix.

As shown in Fig 4, perception error stands out as the most common type of error across all models, with GPT-4V exhibiting a significantly lower perception error rate (51%) compared to GeminiProVision (76.9%) and InternVL (67.2%), indicating its superior performance in perception tasks. Reasoning error emerges as the second most prevalent error type, with InternVL having the highest reasoning error rate (14.8%), followed by GeminiProVision (10.4%) and GPT-4V (9.94%), highlighting the challenges all models face in complex reasoning tasks.

Additionally, the proportion of lack of knowledge errors is similar across the three models, ranging from 6.99% to 9.0%. It suggests that insufficient knowledge is a common issue. However, GPT-4V has notably higher error rates in

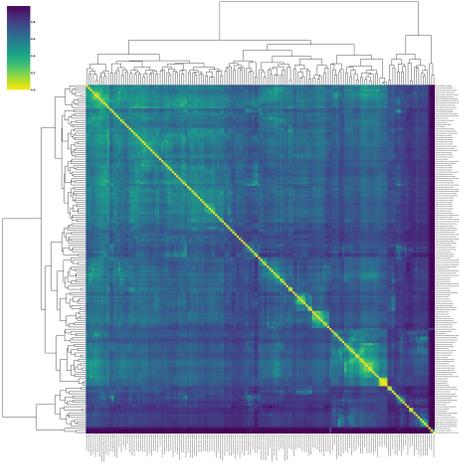


Figure 5. Visualization of task maps and hierarchical clustering with task map. Please zoom in for better visualizations.

lack of capability (19%) and Refusing to Answer (6.11%) compared with the other models, which may be attributed to its more honest approach in acknowledging its limitations and refusing to answer certain questions.

InternVL stands out for its high error rate in failing to follow instructions (6.64%), significantly surpassing GPT-4V (2.99%) and GeminiProVision (1.14%), indicating its struggle in comprehending and executing instructions effectively. On the other hand, annotation error contributes the least to the overall error distribution, implying that the quality of data annotation is high and has a minimal impact on model performance.

To enhance the performance of these large language models, future improvements should focus on addressing the specific error types identified. By targeting perception and reasoning capabilities, tackling the lack of knowledge, and refining the ability to follow instructions, developers can work towards creating more accurate and reliable language models. GPT-4V’s honest approach to its limitations also highlights the importance of transparency in AI systems, which can be further explored and incorporated into future model designs.

## 5. Taskonomy Analysis

Thanks to the extensive coverage of tasks in the MMT-Bench, we can evaluate the multimodal performance of LVLMs on a task map. In this way, the roles of different tasks in multimodal capability can be systematically interpreted by analyzing relationships between tasks in the map.

### 5.1. Analytical Tools

**Task map.** To investigate the relationships between subtasks, we quantify each subtask as a task vector by fol-

Table 3. The relationship between task distance threshold  $\delta$  (normalized by the maximum task distance on the task map) and the consistency of LVLMs performance ranking  $\tau_\delta$ . We see that LVLMs have a more consistent performance ranking when two tasks get closer to each other.

$\delta$	1	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{6}$	$\frac{1}{8}$
$\tau_\delta$	0.29	0.31	0.32	0.41	0.60

lowing (Ilharco et al., 2023). Formally, a task vector is defined by the weight variation between the weight finetuned on task data  $D^t$  and the initial weight  $W_0$  of a probing model, as given by  $V^t = \arg \min_W \mathcal{L}(W|D^t) - W_0$  where the subscript  $t$  denotes the task and  $\mathcal{L}$  is the task loss. Three steps are adopted to obtain  $V^t$ . First, we use pre-trained QwenVL-Chat as the probing model because QwenVL-Chat achieves good results on most subtasks, which helps acquire promising task vectors. Second, we construct task data  $D^t$  by adapting all multi-choice VQA samples into the instruction-following data for each subtask. Third, unlike TaskVec (Ilharco et al., 2023) that finetunes the whole model, we finetune QwenVL-Chat for 3 epochs using LoRA fine-tuning (Hu et al., 2021) for all 162 subtasks, which reduces the length of task vector from 9.6B to 3.5M and consumes less storage resources. With task vector, a task map can be constructed as  $\mathcal{G} = \{G^{st}\}_{s,t=1}^T$  where  $G^{st} = 1 - \cos(V^s, V^t)$  denotes the cosine distance between task  $s$  and  $t$  and  $T = 162$  denoted the total number of subtasks. By definition, we know that  $0 \leq G^{st} \leq 2$ .

**Ranking correlation: Kendall’s tau  $\tau$ .** To quantitatively evaluate LVLMs on a task map, we use the metric of Kendall’s tau  $\tau$  to measure the ranking correlation between performance sequences of LVLMs on different subtasks. The intuition is that model  $A$  would be superior to model  $B$  on task  $t$  if model  $A$  performs better than model  $B$  on task  $s$  when task distance  $G^{st}$  is small. The Kendall’s tau  $\tau$  is defined as  $\tau^{st} = \frac{2}{M(M-1)} \sum_{1 \leq m < n \leq M} \text{sign}((P_m^s - P_n^s)(P_m^t - P_n^t))$  where  $P_m^s$  denotes the performance of model  $m$  on task  $s$  and  $M$  is the number of LVLMs. The function  $\text{sign}(\cdot)$  returns  $-1$  if the argument is negative and  $1$  otherwise. When  $\tau^{st} = 1$ , LVLMs have completely consistent performance ranking on task  $s$  and  $t$ .

### 5.2. Findings on Task Map

**LVLMs obtain a more consistent performance ranking on tasks closer to each other.** We assess whether LVLMs achieve consistent performance on two tasks close to each other. To measure this consistency, we employ the Kendall tau metric as introduced in Sec. 5.1. Specifically, we consider all subtask pairs in which two tasks are closer to each other and calculate their average Kendall’s tau  $\tau$ , which can be given by  $\tau_\delta = \frac{1}{T} \sum_{s=1}^T \frac{1}{|\Delta_s|} \sum_{t \in \Delta_s} \tau^{st}$

Table 4. The number of tasks within each cluster after hierarchical clustering, and the Kendall’s tau  $\tau$  between the average performance of the model on these tasks and the overall performance of the model.

Cluster	1	2	3	4	5	6	7	8	9	10	11	12
# Tasks	11	53	16	16	9	8	7	16	4	9	10	3
$\tau$	0.54	0.73	0.57	0.48	-0.05	0.62	0.63	0.34	0.12	0.57	0.38	0.59
Acc	40.4	64.7	61.9	39.9	55.9	30.0	33.1	40.2	31.4	61.2	33.2	50.7

where  $\Delta_s = \{t : G^{st} \leq \delta\}$  and  $\delta$  is a threshold used to control the proximity between two tasks. As shown in Table 3, as the threshold  $\delta$  decreases, the task distance becomes smaller, and  $\tau_\delta$  increases. This suggests that LVLMs obtain a more consistent performance ranking on tasks closer to each other. Hence, the performance of LVLMs on a new task can be predicted if it is close to one of the MMT-Bench subtasks.

**Out-of-Domain (OoD) tasks discovery.** The OoD tasks mean tasks that the current model struggles to handle. Discovering OoD tasks can provide insights for future evaluation efforts and the development of stronger LVLMs. Since model performance on different tasks is related to task distances, we hypothesize that OoD tasks would be grouped in local regions on the task map. Therefore, we conduct hierarchical clustering on the task map to find OoD tasks. Specifically, 162 subtasks are grouped into 12 clusters as shown in Fig. 5. We use two criteria to identify clusters containing OoD tasks. First, LVLMs would achieve poor performance on OoD tasks. In this regard, we calculate the average multimodal performance within each task cluster over all LVLM models. Second, the performance of LVLMs on OoD tasks would be inconsistent with the overall multimodal score in Table 2 because LVLMs with competitive overall scores would even fail to solve OoD tasks. Hence, we calculate the average ranking correlation  $\tau$  within each cluster. We present these statistics in Table 4 and provide a detailed analysis with the clustering results in Appendix A.

We can see that clusters 8, 9, and 11 achieve low multimodal accuracy and ranking correlation  $\tau$ . In sec 4.2, we find that the model struggles with handling fine-grained visual tasks, such as detection. Through the analysis of these clusters, we similarly find that current multimodal large models cannot perform fine-grained visual cognition and understanding of positional and spatial relationships, such as localization and detection tasks. Moreover, they exhibit poor performance in tasks related to new data structures or types of images, showing a lack of proficiency in handling tasks related to GUI and special data structures like tables.

**In-domain tasks discovery.** In-domain tasks are tasks that most current multimodal large models can handle correctly. Discovering in-domain tasks guides the commercial application of LVLMs in specific scenarios. Different from OoD

tasks, we identify in-domain tasks by looking for clusters with large ranking correlation  $\tau$  and high multimodal accuracy. From Table 4, we can see that clusters 2, 3, and 10 achieve relatively high accuracy and large ranking correlation  $\tau$ . We observe that current multimodal large models possess strong high-level visual comprehension capabilities, enabling them to effectively handle visual recognition tasks, even when dealing with specialized images such as medical images, which is also found in sec 4.2. Moreover, they benefit from the powerful LLMs to accurately describe images. We provide a detailed analysis along with the clustering results in Appendix A.

## 6. Conclusion and Discussion

In this work, we introduce MMT-Bench, a comprehensive benchmark designed to evaluate LVLMs in multimodal multitask understanding. The breadth of MMT-Bench is highlighted by its meticulously curated dataset of 31, 325 multi-choice questions covering 162 multimodal tasks. Our evaluation reveals significant challenges for current LVLMs posed by our MMT-Bench. We present a taskonomy analysis of LVLMs on the task map, allowing us to predict the performance of a new task. Our goal with MMT-Bench is to measure the progress on the path to multitask AGI. We shall acknowledge that MMT-Bench may not be sufficient as a standard for determining whether multitask AGI has been achieved, as it is impossible to include all multimodal tasks. However, we believe that it should be necessary for a multitask AGI to achieve strong performance on MMT-Bench. We will continue to expand the task set of MMT-Bench. We believe that MMT-Bench will inspire further research and development in LVLMs, bringing us closer to the realization of truly intelligent multimodal systems.

## Acknowledgements

This paper is partially supported by the National Key R & D Program of China No.2022ZD0160101 & No.2022ZD0161000.

## Impact Statement

The development and widespread adoption of MMT-Bench as a benchmark for evaluating large vision-language models (LVLMs) have the potential to significantly impact the field of artificial intelligence. While MMT-Bench offers valuable insights and guidance for advancing LVLM research, it is important to consider its broader impact, including ethical considerations and potential societal consequences. One potential positive impact of MMT-Bench is its role in driving advancements in LVLM technology, leading to improved performance and capabilities in various multimodal tasks. This could benefit numerous applications, such as visual

dialogue, video analysis, and document understanding, ultimately enhancing user experiences and productivity.

However, it is crucial to recognize and address potential negative impacts as well. One of the primary limitations of MMT-Bench is its reliance on curated data, which may inadvertently introduce biases based on the sources and methodologies used for data collection. For example, the performance of each meta-task is obtained by taking the average over all subtasks, which may lead to biased assessment because meta-tasks comprise different numbers of subtasks. Moreover, the selection of tasks and subtasks in MMT-Bench may only partially capture the diversity of real-world scenarios, leading to a limited understanding of LVLMS' capabilities across different domains and populations. Furthermore, the data collection process might disproportionately represent certain demographics or contexts, which can lead to biased evaluations of LVLMS' performance.

The other concern is that the benchmark's emphasis on performance metrics such as overall scores and task-specific accuracies may oversimplify the evaluation process and obscure nuanced differences in LVLMS' performance. This could mask disparities in model performance across demographic groups or domains, contributing to the perpetuation of biases and inequities in AI systems. We are dedicated to collecting as many multimodal tasks as possible into our MMT-Bench for unbiased evaluation.

## References

- Moeimouto face dataset. <http://www.nurs.or.jp/~nagadomi/animeface-character-dataset/>.
- Wikiart. [WikiArt.org](http://WikiArt.org), 2022.
- Acharya, M., Kafle, K., and Kanan, C. Tallyqa: Answering complex counting questions. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 8076–8084, 2019.
- Achille, A., Lam, M., Tewari, R., Ravichandran, A., Maji, S., Fowlkes, C. C., Soatto, S., and Perona, P. Task2vec: Task embedding for meta-learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6430–6439, 2019.
- AI, ., ., Young, A., Chen, B., Li, C., Huang, C., Zhang, G., Zhang, G., Li, H., Zhu, J., Chen, J., Chang, J., Yu, K., Liu, P., Liu, Q., Yue, S., Yang, S., Yang, S., Yu, T., Xie, W., Huang, W., Hu, X., Ren, X., Niu, X., Nie, P., Xu, Y., Liu, Y., Wang, Y., Cai, Y., Gu, Z., Liu, Z., and Dai, Z. Yi: Open foundation models by 01.ai, 2024.
- Alam, F., Alam, T., Hasan, M. A., Hasnat, A., Imran, M., and Ofli, F. Medic: A multi-task learning dataset for disaster image classification. *Neural Computing and Applications*, 35:2609–2632, 2023.
- ALESSIO, C. Animals-10. <https://www.kaggle.com/datasets/alessiocorrado99/animals10?select=translate.py>, 2020.
- anchen li. Sketch2code. <https://github.com/mzbac/sketch2code>, 2018.
- Andriluka, M., Pishchulin, L., Gehler, P., and Schiele, B.
- Aneja, D., Colburn, A., Faigin, G., Shapiro, L., and Mones, B. Modeling stylized character expressions via deep learning. In *Asian Conference on Computer Vision*, pp. 136–153. Springer, 2016.
- Anguita, D., Ghio, A., Oneto, L., Parra, X., Reyes-Ortiz, J. L., et al. A public domain dataset for human activity recognition using smartphones. In *Esann*, volume 3, pp. 3, 2013.
- Anthropic. Claude, 2023. URL <https://www.anthropic.com>. Accessed: 2023-04-18.
- Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C. L., and Parikh, D. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- Artem. Image season recognition. <https://www.kaggle.com/c/image-season-recognition/data>, 2013.
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., and Zhou, J. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023.
- BANERJEE, S. Animal image dataset (90 different animals). <https://www.kaggle.com/datasets/iamsouravbanerjee/animal-image-dataset-90-different-animals>, 2022.
- Bell, S., Upchurch, P., Snavely, N., and Bala, K. Opensurfaces: A richly annotated catalog of surface appearance. *ACM Transactions on graphics (TOG)*, 32(4):1–17, 2013.
- Beltramelli, T. pix2code: Generating code from a graphical user interface screenshot. *arXiv preprint arXiv:1705.07962*, 2017.
- Bergmann, P., Fauser, M., Sattlegger, D., and Steger, C. Mvtec ad—a comprehensive real-world dataset for unsupervised anomaly detection. In *Proceedings of the*

- IEEE/CVF conference on computer vision and pattern recognition*, pp. 9592–9600, 2019.
- Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Guo, Y., et al. Improving image generation with better captions. *Computer Science*. <https://cdn.openai.com/papers/dall-e-3.pdf>, 2(3): 8, 2023.
- BHATHENA, J. Weather image recognition. <https://www.kaggle.com/datasets/jehanbathena/weather-dataset>, 2021.
- Bitton-Guetta, N., Bitton, Y., Hessel, J., Schmidt, L., Elovici, Y., Stanovsky, G., and Schwartz, R. Breaking common sense: Whoops! a vision-and-language benchmark of synthetic and compositional images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2616–2627, 2023.
- Bossard, L., Guillaumin, M., and Van Gool, L. Food-101—mining discriminative components with random forests. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VI 13*, pp. 446–461. Springer, 2014.
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., and Beijbom, O. nusenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11621–11631, 2020.
- Cai, M., Liu, H., Mustikovela, S. K., Meyer, G. P., Chai, Y., Park, D., and Lee, Y. J. Making large multimodal models understand arbitrary visual prompts. *arXiv preprint arXiv:2312.00784*, 2023.
- Chao, Y.-W., Liu, Y., Liu, X., Zeng, H., and Deng, J. Learning to detect human-object interactions. In *2018 IEEE winter conference on applications of computer vision (wacv)*, pp. 381–389. IEEE, 2018.
- Chen, D. L. and Dolan, W. B. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (ACL-2011)*, Portland, OR, June 2011.
- Chen, L., Li, J., Dong, X., Zhang, P., He, C., Wang, J., Zhao, F., and Lin, D. Sharegpt4v: Improving large multi-modal models with better captions. *arXiv preprint arXiv:2311.12793*, 2023a.
- Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., Li, B., Luo, P., Lu, T., Qiao, Y., and Dai, J. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*, 2023b.
- Chen, Z., Wang, W., Tian, H., Ye, S., Gao, Z., Cui, E., Tong, W., Hu, K., Luo, J., Ma, Z., et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024.
- Cheng, M.-M., Mitra, N. J., Huang, X., Torr, P. H., and Hu, S.-M. Global contrast based salient region detection. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):569–582, 2014a.
- Cheng, Y., Fu, H., Wei, X., Xiao, J., and Cao, X. Depth enhanced saliency detection method. In *Proceedings of international conference on internet multimedia computing and service*, pp. 23–27, 2014b.
- Chi, Z., Huang, H., Xu, H.-D., Yu, H., Yin, W., and Mao, X.-L. Complicated table structure recognition. *arXiv preprint arXiv:1908.04729*, 2019.
- CHUKS, P. Fake/real logo detection dataset. <https://www.kaggle.com/datasets/prosperchuks/fakereal-logo-detection-dataset>, 2023.
- Chung, H. W., Hou, L., Longpre, S., Zoph, B., Tay, Y., Fedus, W., Li, E., Wang, X., Dehghani, M., Brahma, S., Webson, A., Gu, S. S., Dai, Z., Suzgun, M., Chen, X., Chowdhery, A., Narang, S., Mishra, G., Yu, A., Zhao, V., Huang, Y., Dai, A., Yu, H., Petrov, S., Chi, E. H., Dean, J., Devlin, J., Roberts, A., Zhou, D., Le, Q. V., and Wei, J. Scaling instruction-finetuned language models, 2022. URL <https://arxiv.org/abs/2210.11416>.
- CIJOV, A. Images alike. <https://www.kaggle.com/datasets/alincijov/images-alike>, 2021.
- Contributors, O. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>, 2023a.
- Contributors, T.-M. Transcore-m. <https://github.com/PCIResearch/TransCore-M>, 2023b.
- Contributors, X. Xtuner: A toolkit for efficiently fine-tuning llm. <https://github.com/InternLM/xtuner>, 2023c.
- Dai, W., Li, J., Li, D., Tiong, A. M. H., Zhao, J., Wang, W., Li, B., Fung, P., and Hoi, S. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023.
- Danci, M.-D. Architecturalstyle recognition. <https://github.com/dumitru/architectural-style-recognition>, 2019.

- Davison, A. K., Lansley, C., Costen, N., Tan, K., and Yap, M. H. Samm: A spontaneous micro-facial movement dataset. *IEEE transactions on affective computing*, 9(1): 116–129, 2016.
- Ding, H., Liu, C., He, S., Jiang, X., and Loy, C. C. Mevis: A large-scale benchmark for video segmentation with motion expressions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2694–2703, 2023.
- Ding, J., Xue, N., Xia, G.-S., Bai, X., Yang, W., Yang, M., Belongie, S., Luo, J., Datcu, M., Pelillo, M., and Zhang, L. Object detection in aerial images: A large-scale benchmark and challenges. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021a. doi: 10.1109/TPAMI.2021.3117983.
- Ding, M., Yang, Z., Hong, W., Zheng, W., Zhou, C., Yin, D., Lin, J., Zou, X., Shao, Z., Yang, H., et al. Cogview: Mastering text-to-image generation via transformers. *Advances in Neural Information Processing Systems*, 34: 19822–19835, 2021b.
- Doersch, C., Gupta, A., Markeeva, L., Recasens, A., Smaira, L., Aytar, Y., Carreira, J., Zisserman, A., and Yang, Y. Tap-vid: A benchmark for tracking any point in a video. *Advances in Neural Information Processing Systems*, 35: 13610–13626, 2022.
- Dong, X., Zhang, P., Zang, Y., Cao, Y., Wang, B., Ouyang, L., Wei, X., Zhang, S., Duan, H., Cao, M., Zhang, W., Li, Y., Yan, H., Gao, Y., Zhang, X., Li, W., Li, J., Chen, K., He, C., Zhang, X., Qiao, Y., Lin, D., and Wang, J. Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. *arXiv preprint arXiv:2401.16420*, 2024.
- Dr. Hamid Rahim Sheikh, Dr. Alan C. Bovik, D. L. C. D. Z. W. Live image quality assessment database. <https://live.ece.utexas.edu/research/quality/subjective.htm>, 2006.
- El Korchi, A. and Ghanou, Y. 2d geometric shapes dataset–for machine learning and pattern recognition. *Data in Brief*, 32:106090, 2020.
- Ertler, C., Mislej, J., Ollmann, T., Porzi, L., Neuhold, G., and Kuang, Y. The mapillary traffic sign dataset for detection and classification on a global scale. In *European Conference on Computer Vision*, pp. 68–84. Springer, 2020.
- Everingham, M., Eslami, S. M. A., Van Gool, L., Williams, C. K. I., Winn, J., and Zisserman, A. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, January 2015.
- Fan, D.-P., Ji, G.-P., Sun, G., Cheng, M.-M., Shen, J., and Shao, L. Camouflaged object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2777–2787, 2020.
- Fu, C., Chen, P., Shen, Y., Qin, Y., Zhang, M., Lin, X., Yang, J., Zheng, X., Li, K., Sun, X., et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023.
- GAJARE, N. Rock images. <https://www.kaggle.com/datasets/neelgajare/rocks-dataset>, 2022.
- Gao, P., Han, J., Zhang, R., Lin, Z., Geng, S., Zhou, A., Zhang, W., Lu, P., He, C., Yue, X., Li, H., and Qiao, Y. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023.
- Gbeminiyi, A. Multi-class weather dataset for image classification. *Mendeley Data*, 6:15–23, 2018.
- Geiger, A., Lenz, P., and Urtasun, R. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 3354–3361. IEEE, 2012.
- GERRY. 30 musical instruments -image classification. <https://www.kaggle.com/datasets/gpiosenka/musical-instruments-image-classification/data>, 2021.
- GERRY. 100 sports image classification. <https://www.kaggle.com/datasets/gpiosenka/sports-classification>, 2023.
- GHIMIRE, B. Landscape color and grayscale images. <https://www.kaggle.com/datasets/theblackmamba31/landscape-image-colorization/code>, 2021.
- Hanhe Lin, F. W. S. Helmet dataset. <https://osf.io/4pwj8/>, 2020.
- Haq, N. U., Fraz, M. M., Hashmi, T., and Shahzad, M. Orientation aware weapons detection in visual data: a benchmark dataset. *Computing*, 104(12):2581–2604, 2022.
- HARI, P. Movie posters. <https://www.kaggle.com/datasets/phiitm/movie-posters>, 2020.
- Hsieh, M.-R., Lin, Y.-L., and Hsu, W. H. Drone-based object counting by spatially regularized regional proposal network. In *Proceedings of the IEEE international conference on computer vision*, pp. 4145–4153, 2017.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.

- Hu, Y., Li, T., Lu, Q., Shao, W., He, J., Qiao, Y., and Luo, P. Omnimedvqa: A new large-scale comprehensive evaluation benchmark for medical lvlm. *arXiv preprint arXiv:2402.09181*, 2024.
- Huang, G. B., Ramesh, M., Berg, T., and Learned-Miller, E. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- Huang, M.-L., Liao, Y.-C., Shiau, K.-L., and Tseng, Y.-L. Traditional chinese god image dataset: A glimpse of chinese culture. *Data in Brief*, 46:108861, 2023.
- Huang, T.-H., Ferraro, F., Mostafazadeh, N., Misra, I., Agrawal, A., Devlin, J., Girshick, R., He, X., Kohli, P., Batra, D., et al. Visual storytelling. In *Proceedings of the 2016 conference of the North American chapter of the association for computational linguistics: Human language technologies*, pp. 1233–1239, 2016.
- Huang, Z., Chen, K., He, J., Bai, X., Karatzas, D., Lu, S., and Jawahar, C. Icdar2019 competition on scanned receipt ocr and information extraction. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pp. 1516–1520. IEEE, 2019.
- Hudson, D. A. and Manning, C. D. 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*, pp. 6700–6709, 2019.
- Huttunen, H. Tau vehicle type recognition competition. <https://www.kaggle.com/c/vehicle>, 2019.
- Hwang, E. and Shwartz, V. Memecap: A dataset for captioning and interpreting memes. *arXiv preprint arXiv:2305.13703*, 2023.
- ICARO. Best artworks of all time. <https://www.kaggle.com/datasets/ikarus777/best-artworks-of-all-time>, 2019.
- Idrees, H., Zamir, A. R., Jiang, Y.-G., Gorban, A., Laptev, I., Sukthankar, R., and Shah, M. The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding*, 155:1–23, 2017.
- Ilharco, G., Ribeiro, M. T., Wortsman, M., Gururangan, S., Schmidt, L., Hajishirzi, H., and Farhadi, A. Editing models with task arithmetic. *the 11th International Conference on Learning Representation (ICLR 2023)*, 2023.
- INNAT. Van gogh paintings. <https://www.kaggle.com/datasets/ipythonx/van-gogh-paintings>, 2022.
- Jain, V. and Learned-Miller, E. Fddb: A benchmark for face detection in unconstrained settings. Technical report, UMass Amherst technical report, 2010.
- JENSEN, M. B. Lisa traffic light dataset. <https://www.kaggle.com/datasets/mbornoe/lisa-traffic-light-dataset>, 2018.
- Jhamtani, H. and Berg-Kirkpatrick, T. Learning to describe differences between pairs of similar images. *arXiv preprint arXiv:1808.10584*, 2018.
- JOHANN. Vehicle type recognition. <https://www.kaggle.com/code/theoneandonlyp/vehicle-type-recognition>, 2023.
- Ju, R., Liu, Y., Ren, T., Ge, L., and Wu, G. Depth-aware salient object detection using anisotropic center-surround difference. *Signal Processing: Image Communication*, 38:115–126, 2015.
- Karatzas, D., Shafait, F., Uchida, S., Iwamura, M., i Bigorda, L. G., Mestre, S. R., Mas, J., Mota, D. F., Almazan, J. A., and De Las Heras, L. P. Icdar 2013 robust reading competition. In *2013 12th international conference on document analysis and recognition*, pp. 1484–1493. IEEE, 2013.
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- Kiela, D., Firooz, H., Mohan, A., Goswami, V., Singh, A., Ringshia, P., and Testuggine, D. The hateful memes challenge: Detecting hate speech in multimodal memes. *Advances in neural information processing systems*, 33: 2611–2624, 2020.
- Kondo, Y., Ukita, N., Yamaguchi, T., Hou, H.-Y., Shen, M.-Y., Hsu, C.-C., Huang, E.-M., Huang, Y.-C., Xia, Y.-C., Wang, C.-Y., Lee, C.-Y., Huo, D., Kastner, M. A., Liu, T., Kawanishi, Y., Hirayama, T., Komamizu, T., Ide, I., Shinya, Y., Liu, X., Liang, G., and Yasui, S. MVA2023 Small Object Detection Challenge for Spotting Birds: Dataset, Methods, and Results. In *2023 18th International Conference on Machine Vision and Applications (MVA)*, 2023. <https://www.mva-org.jp/mva2023/challenge>.
- Kong, Y., Jia, Y., and Fu, Y. Learning human interaction by interactive phrases. In *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part I 12*, pp. 300–313. Springer, 2012.

- Kosti, R., Alvarez, J. M., Recasens, A., and Lapedriza, A. Context based emotion recognition using emotic dataset. *IEEE transactions on pattern analysis and machine intelligence*, 42(11):2755–2766, 2019.
- Krause, J., Johnson, J., Krishna, R., and Fei-Fei, L. A hierarchical approach for generating descriptive image paragraphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 317–325, 2017.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017a.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.-J., Shamma, D. A., et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017b.
- KUMAR, V. Face mask detection. <https://www.kaggle.com/datasets/vijaykumar1799/face-mask-detection>, 2021.
- KUMARUJJAWAL1234. Religious symbols-image classification. <https://www.kaggle.com/datasets/kumarujawal123456/famous-religious-symbols>, 2023.
- Kylberg, G. The kylberg texture dataset v. 1.0. External report (Blue series) 35, Centre for Image Analysis, Swedish University of Agricultural Sciences and Uppsala University, Uppsala, Sweden, September 2011. URL <http://www.cb.uu.se/~gustaf/texture/>.
- LABS, D. Electronics object image dataset — computer parts. <https://www.kaggle.com/datasets/dataclusterlabs/electronics-mouse-keyboard-image-dataset>, 2023a.
- LABS, D. Transparent object images — indoor object dataset. <https://www.kaggle.com/datasets/dataclusterlabs/transparent-object-detection>, 2023b.
- Lai, X., Tian, Z., Chen, Y., Li, Y., Yuan, Y., Liu, S., and Jia, J. Lisa: Reasoning segmentation via large language model. *arXiv preprint arXiv:2308.00692*, 2023.
- Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.
- LARXEL. Helmet detection. <https://www.kaggle.com/datasets/andrewmvd/helmet-detection>, 2020a.
- LARXEL. Road sign detection. <https://www.kaggle.com/datasets/andrewmvd/road-sign-detection>, 2020b.
- Latif, E., Mai, G., Nyaaba, M., Wu, X., Liu, N., Lu, G., Li, S., Liu, T., and Zhai, X. Agi: Artificial general intelligence for education, 2023.
- Lazebnik, S., Schmid, C., and Ponce, J. A sparse texture representation using local affine regions. *IEEE transactions on pattern analysis and machine intelligence*, 27(8):1265–1278, 2005.
- Le, Y. and Yang, X. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7):3, 2015.
- Lee, J., Kim, S., Kim, S., Park, J., and Sohn, K. Context-aware emotion recognition networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2019.
- Li, B., Wang, R., Wang, G., Ge, Y., Ge, Y., and Shan, Y. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023a.
- Li, D., Rodriguez, C., Yu, X., and Li, H. Word-level deep sign language recognition from video: A new large-scale dataset and methods comparison. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 1459–1469, 2020a.
- Li, J., Zhang, J., and Tao, D. Deep automatic natural image matting. *arXiv preprint arXiv:2107.07235*, 2021.
- Li, J., Zhang, J., Maybank, S. J., and Tao, D. Bridging composite and real: towards end-to-end deep image matting. *International Journal of Computer Vision*, 130(2):246–266, 2022.
- Li, J., Li, D., Savarese, S., and Hoi, S. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models, 2023b.
- Li, K., Wang, Y., He, Y., Li, Y., Wang, Y., Liu, Y., Wang, Z., Xu, J., Chen, G., Luo, P., et al. Mvbench: A comprehensive multi-modal video understanding benchmark. *arXiv preprint arXiv:2311.17005*, 2023c.
- Li, X., Wei, T., Chen, Y. P., Tai, Y.-W., and Tang, C.-K. Fss-1000: A 1000-class dataset for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2869–2878, 2020b.

- Li, Y., Yang, X., Sun, P., Qi, H., and Lyu, S. Celeb-df (v2): a new dataset for deepfake forensics [j]. *arXiv preprint arXiv*, 2019.
- Li, Z., Yang, B., Liu, Q., Ma, Z., Zhang, S., Yang, J., Sun, Y., Liu, Y., and Bai, X. Monkey: Image resolution and text label are important things for large multi-modal models. *arXiv preprint arXiv:2311.06607*, 2023d.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pp. 740–755. Springer, 2014a.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pp. 740–755. Springer, 2014b.
- Liu, B., Fu, J., Kato, M. P., and Yoshikawa, M. Beyond narrative description: Generating poetry from images by multi-adversarial training. In *Proceedings of the 26th ACM international conference on Multimedia*, pp. 783–791, 2018a.
- Liu, F., Emerson, G. E. T., and Collier, N. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 2023a.
- Liu, H., Li, C., Li, Y., and Lee, Y. J. Improved baselines with visual instruction tuning, 2023b.
- Liu, H., Li, C., Wu, Q., and Lee, Y. J. Visual instruction tuning, 2023c.
- Liu, H., Li, C., Li, Y., Li, B., Zhang, Y., Shen, S., and Lee, Y. J. Llava-next: Improved reasoning, ocr, and world knowledge, 2024a.
- Liu, S., Ying, K., Zhang, H., Yang, Y., Lin, Y., Zhang, T., Li, C., Qiao, Y., Luo, P., Shao, W., et al. Convbench: A multi-turn conversation evaluation benchmark with hierarchical capability for large vision-language models. *arXiv preprint arXiv:2403.20194*, 2024b.
- Liu, W., W. Luo, D. L., and Gao, S. Future frame prediction for anomaly detection – a new baseline. In *2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018b.
- Liu, X., Liu, W., Mei, T., and Ma, H. A deep learning-based approach to progressive vehicle re-identification for urban surveillance. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14*, pp. 869–884. Springer, 2016a.
- Liu, Y., Duan, H., Zhang, Y., Li, B., Zhang, S., Zhao, W., Yuan, Y., Wang, J., He, C., Liu, Z., et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023d.
- Liu, Y., Zhu, M., Li, H., Chen, H., Wang, X., and Shen, C. Matcher: Segment anything with one shot using all-purpose feature matching. *arXiv preprint arXiv:2305.13310*, 2023e.
- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- Liu, Z., Luo, P., Qiu, S., Wang, X., and Tang, X. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1096–1104, 2016b.
- Lu, C., Shi, J., and Jia, J. Abnormal event detection at 150 fps in matlab. In *Proceedings of the IEEE international conference on computer vision*, pp. 2720–2727, 2013.
- Lu, H., Liu, W., Zhang, B., Wang, B., Dong, K., Liu, B., Sun, J., Ren, T., Li, Z., Yang, H., Sun, Y., Deng, C., Xu, H., Xie, Z., and Ruan, C. Deepseek-vl: Towards real-world vision-language understanding, 2024.
- Lu, P., Bansal, H., Xia, T., Liu, J., Li, C., Hajishirzi, H., Cheng, H., Chang, K.-W., Galley, M., and Gao, J. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023.
- Lv, Y., Zhang, J., Dai, Y., Li, A., Liu, B., Barnes, N., and Fan, D.-P. Simultaneously localize, segment and rank the camouflaged objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11591–11601, 2021.
- MA, A. T. Anime characters personality and facial images. <https://www.kaggle.com/datasets/tianyimasf/anime-characters>, 2023.
- Ma, W. and Liang, S. Polar: Posture-level action recognition dataset. In *2019 6th International Conference on Systems and Informatics (ICSAI)*, pp. 427–433. IEEE, 2019.
- Machajdik, J. and Hanbury, A. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the 18th ACM international conference on Multimedia*, pp. 83–92, 2010.

- Mallikarjuna, P., Targhi, A. T., Fritz, M., Hayman, E., Caputo, B., and Eklundh, J.-O. The kth-tips2 database. *Computational Vision and Active Perception Laboratory, Stockholm, Sweden*, 11:12, 2006.
- Marino, K., Rastegari, M., Farhadi, A., and Mottaghi, R. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pp. 3195–3204, 2019.
- Marti, U.-V. and Bunke, H. The iam-database: an english sentence database for offline handwriting recognition. *International Journal on Document Analysis and Recognition*, 5:39–46, 2002.
- Martin, D., Fowlkes, C., Tal, D., and Malik, J. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int’l Conf. Computer Vision*, volume 2, pp. 416–423, July 2001.
- Masry, A., Long, D. X., Tan, J. Q., Joty, S., and Hoque, E. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022.
- Mathew, M., Bagal, V., Tito, R., Karatzas, D., Valveny, E., and Jawahar, C. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 1697–1706, 2022.
- MENON, S. S. Animals 151. <https://www.kaggle.com/datasets/sharansmenon/animals141>, 2022.
- Mishra, A., Alahari, K., and Jawahar, C. V. Scene text recognition using higher order language priors. In *BMVC*, 2012.
- MOHAMED, M. Garbage classification (12 classes). <https://www.kaggle.com/datasets/mostafaabla/garbage-classification>, 2021.
- Mohamed, Y., Khan, F. F., Haydarov, K., and Elhoseiny, M. It is okay to not be okay: Overcoming emotional bias in affective image captioning by contrastive data collection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, volume abs/2204.07660, 2022.
- Morris, M. R., Sohl-dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C., and Legg, S. Levels of agi: Operationalizing progress on the path to agi. *arXiv preprint arXiv:2311.02462*, 2023.
- Mouchere, H., Viard-Gaudin, C., Zanibbi, R., and Garain, U. Icfhr 2014 competition on recognition of on-line handwritten mathematical expressions (crohme 2014). In *2014 14th International Conference on Frontiers in Handwriting Recognition*, pp. 791–796. IEEE, 2014.
- mrayinteractive. The quick, draw! dataset. <https://github.com/googlecreativelab/quickdraw-dataset>, 2014.
- Ng, X. L., Ong, K. E., Zheng, Q., Ni, Y., Yeo, S. Y., and Liu, J. Animal kingdom: A large and diverse dataset for animal behavior understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 19023–19034, June 2022.
- Nilsback, M.-E. and Zisserman, A. Automated flower classification over a large number of classes. In *2008 Sixth Indian conference on computer vision, graphics & image processing*, pp. 722–729. IEEE, 2008.
- NOGRA, J. A. Face mask usage. <https://www.kaggle.com/datasets/jamesnogra/face-mask-usage>, 2022.
- Obeid, J. and Hoque, E. Chart-to-text: Generating natural language descriptions for charts by adapting the transformer model. *arXiv preprint arXiv:2010.09142*, 2020.
- OLAFENWA, M. Idenprof. <https://github.com/OlafenwaMoses/IdenProf>, 2018.
- OpenAI. Gpt-4o. <https://openai.com/index/hello-gpt-4o/>, 2024.
- Parmar, P. and Morris, B. Action quality assessment across multiple actions. In *2019 IEEE winter conference on applications of computer vision (WACV)*, pp. 1468–1476. IEEE, 2019.
- Parmar, P. and Tran Morris, B. Learning to score olympic events. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pp. 20–28, 2017.
- Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., and Wang, B. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.
- Philbin, J., Chum, O., Isard, M., Sivic, J., and Zisserman, A. Object retrieval with large vocabularies and fast spatial matching. In *2007 IEEE conference on computer vision and pattern recognition*, pp. 1–8. IEEE, 2007.
- Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., and Van Gool, L. The 2017 davis challenge on video object segmentation. *arXiv:1704.00675*, 2017.

- Qi, J., Gao, Y., Hu, Y., Wang, X., Liu, X., Bai, X., Belongie, S., Yuille, A., Torr, P. H., and Bai, S. Occluded video instance segmentation: A benchmark. *International Journal of Computer Vision*, 130(8):2022–2039, 2022.
- Quattoni, A. and Torralba, A. Recognizing indoor scenes. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 413–420. IEEE, 2009.
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning transferable visual models from natural language supervision, 2021.
- Rahman, R., Hasan, R., Farhad, A. A., Laskar, M. T. R., Ashmafee, M. H., and Kamal, A. R. M. Chartsumm: A comprehensive benchmark for automatic chart summarization of long and short summaries. *arXiv preprint arXiv:2304.13620*, 2023.
- Ranjan, V., Sharma, U., Nguyen, T., and Hoai, M. Learning to count everything. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3394–3403, 2021.
- Rawles, C., Li, A., Rodriguez, D., Riva, O., and Lillicrap, T. Android in the wild: A large-scale dataset for android device control. *arXiv preprint arXiv:2307.10088*, 2023.
- RBDash-Team. Rbdash. <https://github.com/RBDash-Team/RBDash>, 2023.
- ROMAN, K. Facial emotion recognition dataset. <https://www.kaggle.com/datasets/tapakah68/facial-emotion-recognition>, 2023.
- Rosenfeld, A., Solbach, M. D., and Tsotsos, J. K. Totally looks like-how humans compare, compared to machines. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1961–1964, 2018.
- Rosler, A., Cozzolino, D., Verdoliva, L., Riess, C., Thies, J., and Nießner, M. Faceforensics++: Learning to detect manipulated facial images. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1–11, 2019.
- SAHA, B. cricket-football-baseball classification. <https://www.kaggle.com/datasets/imbikramsaha/cricket-football-baseball>, 2022.
- Salman Ibne Eunos, JU Komol, R. A. N. S. H. A. M. A. Rock classification dataset. <https://www.kaggle.com/datasets/salmaneunos/rock-classification>, 2021.
- SANYAL, S. Weapon detection dataset. <https://www.kaggle.com/datasets/snehilsanyal/weapon-detection-test>, 2023.
- Sasaki, R., Fujinami, M., and Nakai, H. Comprehensive image dataset for enhancing object detection in chemical experiments. *Data in Brief*, 52:110054, 2024.
- SEKAR, S. Waste classification data. <https://www.kaggle.com/datasets/techsash/waste-classification-data>, 2019.
- Sermanet, P., Xu, K., and Levine, S. Unsupervised perceptual rewards for imitation learning. *arXiv preprint arXiv:1612.06699*, 2016.
- SETH, K. Fruits and vegetables image recognition dataset. <https://www.kaggle.com/datasets/kritikseth/fruit-and-vegetable-image-recognition>, 2022.
- Shao, W., Hu, Y., Gao, P., Lei, M., Zhang, K., Meng, F., Xu, P., Huang, S., Li, H., Qiao, Y., et al. Tiny lvlm-ehub: Early multimodal experiments with bard. *arXiv preprint arXiv:2308.03729*, 2023.
- SHETTY, S. Image colorization. <https://www.kaggle.com/datasets/shravankumar9892/image-colorization>, 2018.
- Silberman, N., Hoiem, D., Kohli, P., and Fergus, R. Indoor segmentation and support inference from rgb-d images. In *Computer Vision—ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part V 12*, pp. 746–760. Springer, 2012.
- Singh, A. asingh33/CNNGestureRecognizer: CNNGestureRecognizer, November 2017. URL <https://doi.org/10.5281/zenodo.1064825>.
- Singh, S. S. Teaching machines to code: neural markup generation with visual attention. *arXiv preprint arXiv:1802.05415*, 2018.
- Singhal, K., Azizi, S., Tu, T., Mahdavi, S. S., Wei, J., Chung, H. W., Scales, N., Tanwani, A., Cole-Lewis, H., Pfohl, S., et al. Large language models encode clinical knowledge. *Nature*, 620(7972):172–180, 2023.
- Smith, B., Yin, Q., Feiner, S., and Nayar, S. Gaze Locking: Passive Eye Contact Detection for Human-Object Interaction. In *ACM Symposium on User Interface Software and Technology (UIST)*, pp. 271–280, Oct 2013.
- Srivastava, N., Mansimov, E., and Salakhudinov, R. Unsupervised learning of video representations using lstms. In

- International conference on machine learning*, pp. 843–852. PMLR, 2015.
- stepanje. Metal parts defect detection dataset. <https://github.com/stepanje/MPDD>, 2021.
- Sun, P., Kretschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2446–2454, 2020.
- Sun, Q., Fang, Y., Wu, L., Wang, X., and Cao, Y. Evalclip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.
- sunshine\_joe. FlickrSportLogos-10. <https://github.com/sunshine-joe/FlickrSportLogos-10>, 2018.
- SUPERPOTATO9. Ai recognition dataset. <https://www.kaggle.com/datasets/superpotato9/dalle-recognition-dataset>, 2024.
- TAMRAKAR, A. E waste image dataset. <https://www.kaggle.com/datasets/akshat103/e-waste-image-dataset>, 2023.
- Team, G. Gemini: A family of highly capable multimodal models, 2023a.
- Team, I. Internlm: A multilingual language model with progressively enhanced capabilities. <https://github.com/InternLM/InternLM>, 2023b.
- Team, Q. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023c.
- tensorflow. Flower photos. [https://www.tensorflow.org/tutorials/load\\_data/images](https://www.tensorflow.org/tutorials/load_data/images).
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., and Lample, G. Llama: Open and efficient foundation language models, 2023a.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., Fuller, B., Gao, C., Goswami, V., Goyal, N., Hartshorn, A., Hosseini, S., Hou, R., Inan, H., Kardas, M., Kerkez, V., Khabsa, M., Kloumann, I., Korenev, A., Koura, P. S., Lachaux, M.-A., Lavril, T., Lee, J., Liskovich, D., Lu, Y., Mao, Y., Martinet, X., Mihaylov, T., Mishra, P., Molybog, I., Nie, Y., Poulton, A., Reizenstein, J., Rungta, R., Saladi, K., Schelten, A., Silva, R., Smith, E. M., Subramanian, R., Tan, X. E., Tang, B., Taylor, R., Williams, A., Kuan, J. X., Xu, P., Yan, Z., Zarov, I., Zhang, Y., Fan, A., Kambadur, M., Narang, S., Rodriguez, A., Stojnic, R., Edunov, S., and Scialom, T. Llama 2: Open foundation and fine-tuned chat models, 2023b.
- TYAGI, A. music instruments classification. <https://www.kaggle.com/datasets/aayushme/music-instruments-classification/data>, 2020.
- tz28. Chinese-number-gestures-recognition. <https://github.com/tz28/Chinese-number-gestures-recognition>, 2018.
- Uy, M. A., Pham, Q.-H., Hua, B.-S., Nguyen, D. T., and Yeung, S.-K. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *International Conference on Computer Vision (ICCV)*, 2019.
- VERMA, A. Disaster images dataset. <https://www.kaggle.com/datasets/varpit94/disaster-images-dataset>, 2021.
- Wah, C., Branson, S., Welinder, P., Perona, P., and Belongie, S. The caltech-ucsd birds-200-2011 dataset. 2011.
- Wallace, B., Wu, Z., and Hariharan, B. Can we characterize tasks without labels or features? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1245–1254, 2021.
- Wang, H., Ge, S., Lipton, Z., and Xing, E. P. Learning robust global representations by penalizing local predictive power. *Advances in Neural Information Processing Systems*, 32, 2019.
- Wang, L., Lu, H., Wang, Y., Feng, M., Wang, D., Yin, B., and Ruan, X. Learning to detect salient objects with image-level supervision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 136–145, 2017.
- Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., Xu, J., Xu, B., Li, J., Dong, Y., Ding, M., and Tang, J. Cogvlm: Visual expert for pretrained language models. 2023.
- Wang, Z., Yang, J., Jin, H., Shechtman, E., Agarwala, A., Brandt, J., and Huang, T. S. Deepfont: Identify your font from an image. In *Proceedings of the 23rd ACM international conference on Multimedia*, pp. 451–459, 2015.

- Wu, B., Yu, S., Chen, Zhenfang, T. J. B., and Gan, C. Star: A benchmark for situated reasoning in real-world videos. In *Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., and Xiao, J. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1912–1920, 2015.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.
- Xie, B., Zhang, S., Zhou, Z., Li, B., Zhang, Y., Hessel, J., Yang, J., and Liu, Z. Funqa: Towards surprising video comprehension. *arXiv preprint arXiv:2306.14899*, 2023.
- Xie, E., Wang, W., Wang, W., Ding, M., Shen, C., and Luo, P. Segmenting transparent objects in the wild. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pp. 696–711. Springer, 2020.
- Xu, J., Mei, T., Yao, T., and Rui, Y. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5288–5296, 2016.
- Xu, L., Jin, S., Zeng, W., Liu, W., Qian, C., Ouyang, W., Luo, P., and Wang, X. Pose for everything: Towards category-agnostic pose estimation. October 2022.
- Xu, P., Shao, W., Zhang, K., Gao, P., Liu, S., Lei, M., Meng, F., Huang, S., Qiao, Y., and Luo, P. Lvlm-ehub: A comprehensive evaluation benchmark for large vision-language models. *arXiv preprint arXiv:2306.09265*, 2023.
- Yan, W.-J., Wu, Q., Liu, Y.-J., Wang, S.-J., and Fu, X. Casme database: A dataset of spontaneous micro-expressions collected from neutralized faces. In *2013 10th IEEE international conference and workshops on automatic face and gesture recognition (FG)*, pp. 1–7. IEEE, 2013.
- Yang, J., Zhang, H., Li, F., Zou, X., Li, C., and Gao, J. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*, 2023a.
- Yang, L., Fan, Y., and Xu, N. Video instance segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5188–5197, 2019.
- Yang, S., Luo, P., Loy, C. C., and Tang, X. Wider face: A face detection benchmark. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- Yang, X., Yan, J., Liao, W., Yang, X., Tang, J., and He, T. Scrdet++: Detecting small, cluttered and rotated objects via instance-level feature denoising and rotation loss smoothing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- Yang, Z., Li, L., Lin, K., Wang, J., Lin, C.-C., Liu, Z., and Wang, L. The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9 (1):1, 2023b.
- Yang, Z., Li, L., Lin, K., Wang, J., Lin, C.-C., Liu, Z., and Wang, L. The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 9 (1):1, 2023c.
- Yang, Z., Liu, J., Han, Y., Chen, X., Huang, Z., Fu, B., and Yu, G. Appagent: Multimodal agents as smartphone users. *arXiv preprint arXiv:2312.13771*, 2023d.
- Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023a.
- Ye, Q., Xu, H., Ye, J., Yan, M., Hu, A., Liu, H., Qian, Q., Zhang, J., Huang, F., and Zhou, J. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration, 2023b.
- Yin, Z., Wang, J., Cao, J., Shi, Z., Liu, D., Li, M., Sheng, L., Bai, L., Huang, X., Wang, Z., et al. Lamm: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark. *arXiv preprint arXiv:2306.06687*, 2023.
- Young, P., Lai, A., Hodosh, M., and Hockenmaier, J. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78, 2014.
- Yu, H., Xu, Y., Zhang, J., Zhao, W., Guan, Z., and Tao, D. Ap-10k: A benchmark for animal pose estimation in the wild. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- Yu, L., Poirson, P., Yang, S., Berg, A. C., and Berg, T. L. Modeling context in referring expressions. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II 14*, pp. 69–85. Springer, 2016.
- Yu, W., Yang, Z., Li, L., Wang, J., Lin, K., Liu, Z., Wang, X., and Wang, L. Mm-vet: Evaluating large multi-modal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023.

- Yu, X., Gong, Y., Jiang, N., Ye, Q., and Han, Z. Scale match for tiny person detection. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 1257–1265, 2020.
- Yuan, Y., Liu, X., Dikubab, W., Liu, H., Ji, Z., Wu, Z., and Bai, X. Syntax-aware network for handwritten mathematical expression recognition. *arXiv preprint arXiv:2203.01601*, 2022.
- Yue, X., Ni, Y., Zhang, K., Zheng, T., Liu, R., Zhang, G., Stevens, S., Jiang, D., Ren, W., Sun, Y., Wei, C., Yu, B., Yuan, R., Sun, R., Yin, M., Zheng, B., Yang, Z., Liu, Y., Huang, W., Sun, H., Su, Y., and Chen, W. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. *arXiv preprint arXiv:2311.16502*, 2023a.
- Yue, X., Ni, Y., Zhang, K., Zheng, T., Liu, R., Zhang, G., Stevens, S., Jiang, D., Ren, W., Sun, Y., et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. *arXiv preprint arXiv:2311.16502*, 2023b.
- Zamir, A. R., Sax, A., Shen, W., Guibas, L. J., Malik, J., and Savarese, S. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3712–3722, 2018.
- Zhan, Y., Xiong, Z., and Yuan, Y. Rsvg: Exploring data and models for visual grounding on remote sensing data. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–13, 2023. doi: 10.1109/TGRS.2023.3250471.
- Zhang, P., Dong, X., Wang, B., Cao, Y., Xu, C., Ouyang, L., Zhao, Z., Ding, S., Zhang, S., Duan, H., Zhang, W., Yan, H., Zhang, X., Li, W., Li, J., Chen, K., He, C., Zhang, X., Qiao, Y., Lin, D., and Wang, J. Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*, 2023a.
- Zhang, R., Han, J., Liu, C., Gao, P., Zhou, A., Hu, X., Yan, S., Lu, P., Li, H., and Qiao, Y. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023b.
- Zhang, T., Zhang, X., Li, J., Xu, X., Wang, B., Zhan, X., Xu, Y., Ke, X., Zeng, T., Su, H., et al. Sar ship detection dataset (ssdd): Official release and comprehensive data analysis. *Remote Sensing*, 13(18):3690, 2021.
- Zhang, W., Zhu, M., and Derpanis, K. G. From actemes to action: A strongly-supervised representation for detailed action understanding. In *Proceedings of the IEEE international conference on computer vision*, pp. 2248–2255, 2013.
- Zhang, Y., Zhou, D., Chen, S., Gao, S., and Ma, Y. Single-image crowd counting via multi-column convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 589–597, 2016.
- Zhang, Y., Pan, J., Zhou, Y., Pan, R., and Chai, J. Grounding visual illusions in language: Do vision-language models perceive illusions like humans? In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 5718–5728, 2023c.
- Zhang, Z., Luo, P., Loy, C.-C., and Tang, X. Learning social relation traits from face images. In *Proceedings of the IEEE international conference on computer vision*, pp. 3631–3639, 2015.
- Zhao, T., Zhang, T., Zhu, M., Shen, H., Lee, K., Lu, X., and Yin, J. VI-checklist: Evaluating pre-trained vision-language models with objects, attributes and relations. *arXiv preprint arXiv:2207.00221*, 2022.
- Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., and Tian, Q. Scalable person re-identification: A benchmark. In *Proceedings of the IEEE international conference on computer vision*, pp. 1116–1124, 2015.
- Zheng, L., Chiang, W.-L., Sheng, Y., Zhuang, S., Wu, Z., Zhuang, Y., Lin, Z., Li, Z., Li, D., Xing, E. P., Zhang, H., Gonzalez, J. E., and Stoica, I. Judging llm-as-a-judge with mt-bench and chatbot arena, 2023.
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017.
- Zhu, P., Wen, L., Du, D., Bian, X., Fan, H., Hu, Q., and Ling, H. Detection and tracking meet drones challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):7380–7399, 2021.
- Zhujian Yang, Tomas Lozano-Perez, J. M. W. L. Kitchen worlds. <https://github.com/Learning-and-Intelligent-Systems/kitchen-worlds>, 2022.

In this appendix, we provide further details as follows:

- Sec. A: Presents hierarchical clustering and more analyses on the task map constructed from our MMT-Bench.
- Sec. B: Includes details on sample size, visual input types, and capabilities of LVLMs evaluated for each subtask.
- Sec. C: Enumerates task abbreviations used throughout the paper.
- Sec. D: Presents detailed model configurations and experimental details in multi-images and visual prompting.
- Sec. E: Compares the performance on tasks involving pixel coordinates and normalized coordinates.
- Sec. F: Compares the performance of LVLMs on different image types and multimodal capabilities.
- Sec. G: Illustrates error cases of GPT-4V, GeminiProVision, and InternVL-Chat on 32 meta-tasks in MMT-Bench.
- Sec. H: Gives the comparison of MMT-Bench with other benchmarks on OCR-Related Tasks.
- Sec. I: Presents some details about the benchmark construction.
- Sec. J: Discusses the OpenCompass protocol used in MMT-Bench and other alternatives.
- Sec. K: Gives the computational resources used in evaluation.
- Sec. L: Provides the detailed performance of 30 models across all 162 subtasks on MMT-Bench.

## A. Task Map

We perform hierarchical clustering on the taskmap, as shown in Fig. 5. When selecting the number of clustering clusters as 12, we analyze the clustering results of the task map and the model performance on the corresponding tasks. Here, we list the names of the tasks within each cluster in Table A2.

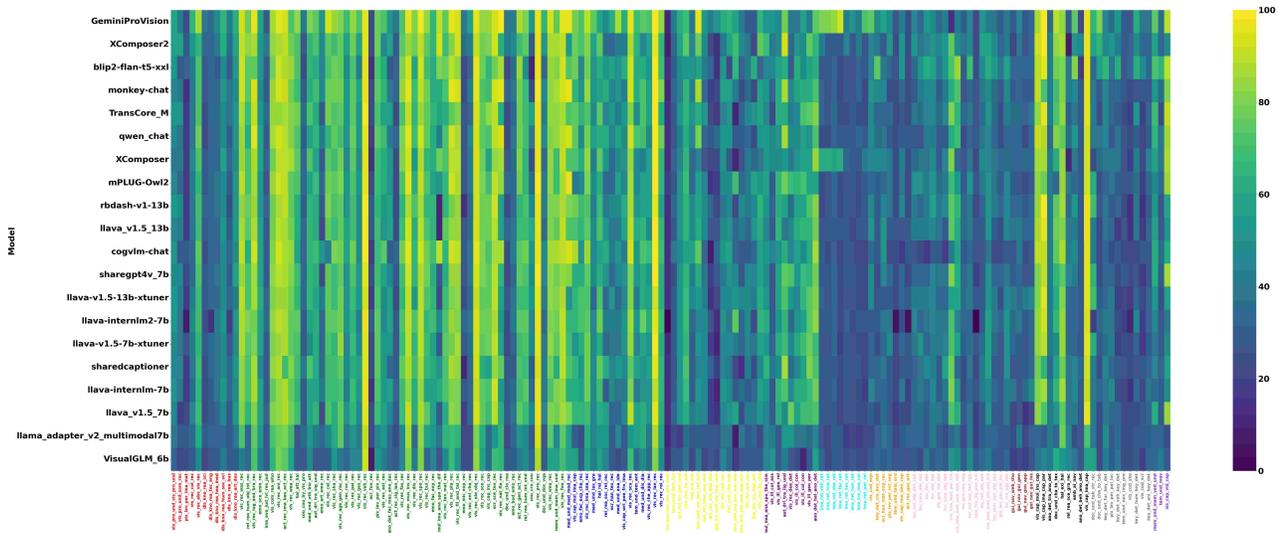


Figure A1. Visualization of model performance on different tasks. Different colours signify the respective categories formed after clustering, arranged from left to right, starting from the first category through to the twelfth. Please zoom in for better visualizations.

**Out-of-Domain (OoD) tasks discovery.** We can see that clusters 8, 9, and 11 achieve low multimodal accuracy and ranking correlation  $\tau$ . From these clusters, we find that current multimodal large models lack the ability to perform fine-grained visual cognition and understanding of positional and spatial relationships, such as localization and detection tasks. Moreover, they exhibit poor performance in tasks related to new data structures or types of images, showing a lack of proficiency in handling tasks related to GUI and special data structures like tables.

- Cluster 8 mainly involves detection, tracking, and localization tasks, all of which are related to the localization of objects within images. This indicates that current large multimodal models lack fine-grained visual cognition and understanding of positional and spatial relationships.
- Tasks in cluster 9 are centered around GUI navigation, a novel task type requiring strong visual understanding, object localization, and expert knowledge in operating mobile devices (Yang et al., 2023d). This suggests that current large multimodal models need further optimization for GUI-related tasks.
- Apart from detection and localization tasks, cluster 11 also includes tasks involving the recognition of special images or their conversion into structured text. The former requires models to possess spatial cognition and fine-grained visual capabilities, while the latter demands robust OCR abilities and extensive knowledge (such as understanding and outputting the basic structure of code or tables). Our testing LVLMs currently fall short in this aspect.

**In-Domain tasks discovery.** From Table 4, we can see that clusters 2, 3, and 10 achieve relatively high accuracy and large ranking correlation  $\tau$ . We observe that current multimodal large models possess strong high-level visual comprehension capabilities, enabling them to effectively handle visual recognition tasks, even when dealing with specialized images such as medical images. Moreover, they benefit from the powerful LLMs to accurately describe images.

- Cluster 2 mainly comprises visual recognition tasks, which require the model to possess certain high-level visual capabilities, yet these tasks are relatively simple. Examining Table 2 and Fig. A1, we observe that the model’s performance within this cluster is generally good. This validates that the current multimodal large models possess fundamental abilities for visual-semantic understanding, allowing them to fulfil recognition tasks.
- Cluster 3 mainly includes visual recognition tasks as well, yet extends to cover sophisticated visual understanding tasks that require primary specialist knowledge, such as medicine and emotion. Within this cluster, the model demonstrates large  $\tau$  and high accuracy, suggesting that current multimodal models pay attention to tasks necessitating the infusion of domain-specific knowledge, beyond just natural images. This implies a certain ability to handle problems in specialized fields.
- In Cluster 10, LVLMs achieve good performance on tasks related to the visual description of the image. It indicates that current large multimodal models can describe the image well. It would stem from the fact that these models are typically tuned by massive image-text pairs.

## B. Hierarchical Structure of MMT-Bench

In Table A3 to Table A5, we present all 32 meta-tasks from MMT-Bench, encompassing a total of 162 subtasks. These tables include details on sample size, visual input types, capabilities of LVLMs evaluated and data source for each subtask.

## C. Task Abbreviations

Given the extensive number of tasks and models tested within the benchmark, we employ abbreviations to condense the manuscript. The abbreviations used throughout the paper are shown in Table A1.

## D. More Experimental Details

### D.1. LVLMs Model Details

Table A6 summarizes the LVLMs information used in this paper, including the corresponding parameter sizes, visual encoders, and LLMs. Note that we use follow OpenCompass’ protocol (Contributors, 2023a) to conduct the evaluation process. The inference time varies with different models. For instance, the smaller LLaVA-v1.5-7B (Liu et al., 2023b) model takes only 12 minutes to complete the evaluation using 8 GPUs, while the larger InternVL-Chat-V1.2-34B model (Chen et al., 2023b) requires 79 minutes and around 80GB of memory. Our open-source codebase supports multi-GPU distributed inference, effectively accelerating the inference process.

Table A1. The Abbreviations of terms mentioned in this paper and their corresponding full terms.

Abbreviation	Full Term	Abbreviation	Full Term
Meta-Task			
VR	Visual Recognition	VI	Visual Illusion
Loc	Localization	MemU	Meme Understanding
OCR	OCR	VPU	Visual Prompt Understanding
Count	Counting	AND	Anomaly Detection
HLN	Hallucination	KD	Keypoint Detection
IR	Image Retrieval	VCR	Visual Commonsense Reasoning
3D	3D	IEJ	Image Evaluation Judgement
VC	Visual Captioning	MIA	Multiple Image Analysis
VG	Visual Grounding	CIM	Cross Image Matching
DU	Doc Understanding	TU	Temporal Understanding
AR	Action Recognition	VCo	Visual Code
PLP	Pixel Level Perception	MedU	Medical Understanding
I2IT	Image-to-image Translation	AUD	Autonomous Driving
RR	Relation Reasoning	DKR	Discipline Knowledge Reasoning
IQT	Intelligence Quotient Test	EA	Embodied AI
Emo	Emotion	GN	GUI Navigation
Subtask			
AQS	Action Quality Assessment	SODRD	Salient Object Detection RGBD
FECR	Facial Expression Change Recognition	SLR	Sign Language Recognition
FR	Face Retrieval	SOT	Single Object Tracking
GAR	General Action Recognition	S2IR	Sketch2image Retrieval
HR	Handwritten Retrieval	SD	Spot the Diff
I2IR	Image2image Retrieval	SS	Spot the Similarity
IC	Image Colorization	TA	Temporal Anticipation
MVU	Meme Video Understanding	TL	Temporal Localization
ME	MEVIS	TO	Temporal Ordering
MIC	Multiple Image Captioning	T2IR	Text2image Retrieval
NIP	Next Image Prediction	3DCR	3D CAD Recognition
OSD	One-shot Detection	3DIR	3D Indoor Recognition
Pre	Person Reid	VR	Vehicle Retrieval
PT	Point Tracking	VC	Video Captioning

## D.2. Multi-Images Prompt Experimental Details

In terms of the 28 tasks requiring multiple images as input, please see Table A7-A10 for the specific task names given task abbreviations. Besides, we also present the designed prompt examples for Single-Image LVLMs and Multi-Images LVLMs in Table A7-A10 for reference.

## D.3. Visual Referring Prompting Experimental Details

In Section 4.3, we explore the differential efficacy of visual prompting compared to alternative prompting strategies across a spectrum of 14 distinct tasks. These encompass human interaction understanding, social relation recognition, human-object interaction recognition, animal keypoint detection, vehicle keypoint detection, human keypoint detection, clothes keypoint detection, scene text recognition, interactive segmentation, instance captioning, multiple instance captioning, one-shot detection, single object tracking, and counting by visual prompting.

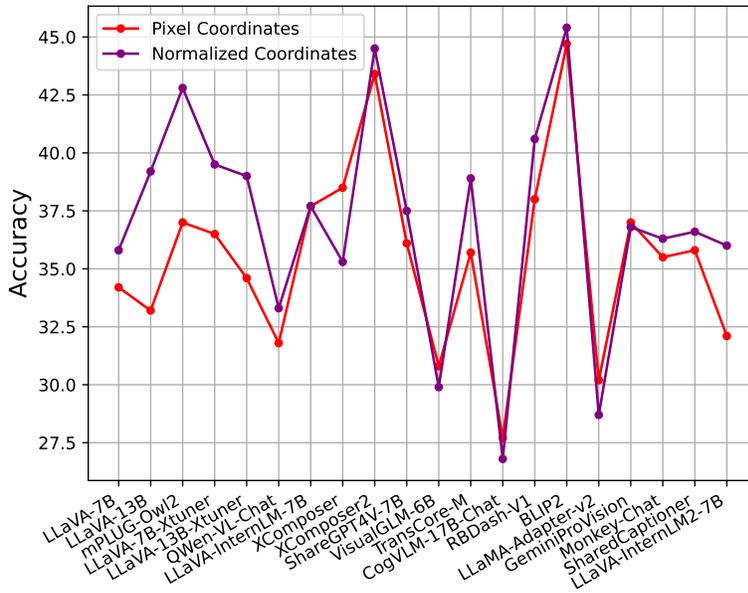


Figure A2. Comparison of coordinate formats for detection tasks across 19 MMT-Bench subtasks, reporting average accuracy.

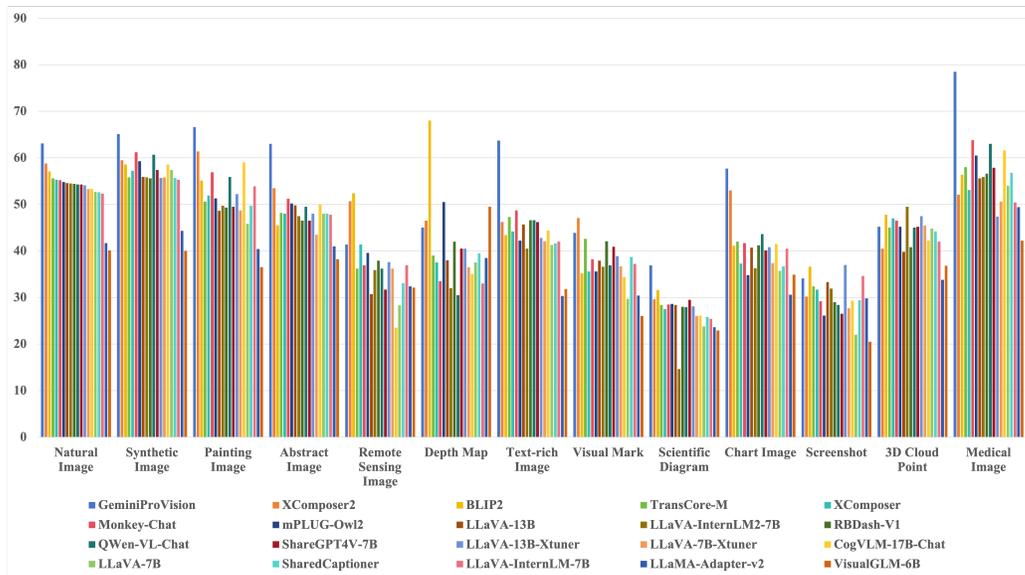


Figure A3. The performance of 20 LVLMs across 13 types of visual input.

## E. Pixel Coordinates vs. Normalized Coordinates

In Fig. A2, we analyze the performance across 19 detection-related tasks, specifically point tracking, image matting, pixel recognition, polygon localization, pixel localization, depth estimation, MEVIS, remote sensing object detection, rotated object detection, small object detection, camouflage object detection, salient object detection in RGB-D, transparent object detection, face detection, object detection, salient object detection in RGB, referring detection, reason segmentation, and image dense captioning. These tasks span Localization, Pixel-level Perception, and Visual Captioning, comparing outcomes under two different coordinate formats. Notably, GeminiProVision lags behind top open-source LVLMs like BLIP2 and XComposer2, which have been extensively trained with detection data. The preference for normalized coordinates among most models is attributed to their use in the training instruction templates.

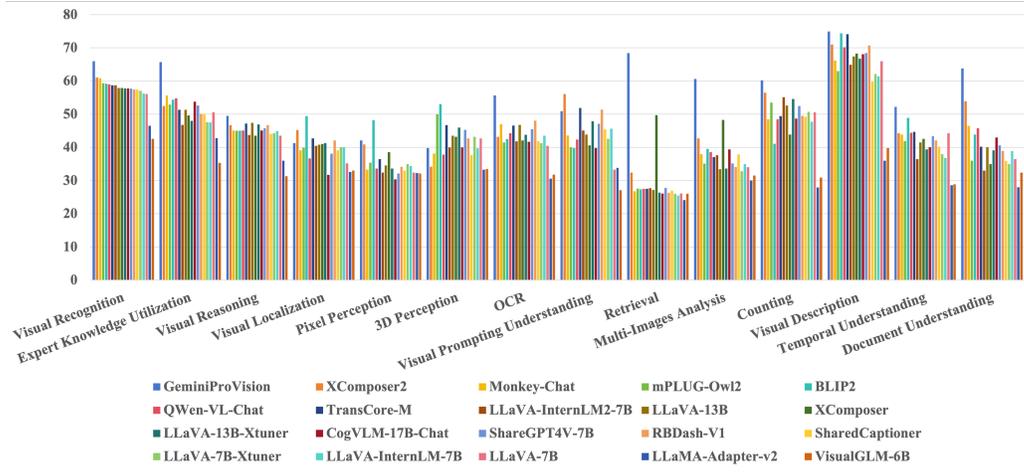


Figure A4. The performance of 20 LVLMs across 14 capabilities.

## F. Analysis on Images Types and Capabilities

**Performance with Different Visual Types.** We compare the performance of 20 LVLMs across 13 types of visual input in Fig. A3. Most LVLMs struggle with Scientific Diagrams due to task difficulty, as many, including Scientific and ”Raven’s Progressive Matrices,” require complex reasoning, a capability current LVLMs do not possess well.

**Performance Across Multimodal Capabilities.** We also compare the performance of 20 LVLMs across 14 types of visual input in Fig. A4. As we can see, GeminiProVision once again exhibits strong superiority across most capabilities, especially in retrieval and multi-image analysis (involving the recognition and matching of multiple images), vastly outperforming other open-source LVLMs. This superiority stems from GeminiProVision’s support for multi-image mode and its powerful generalization abilities, guiding the future direction of open-source models towards the focus on multi-image and video understanding.

Table A2: Details of task clustering on the task map of our MMT-Bench.

Meta-Task	Subtask	# subtasks
Cluster ID: 1		
Visual Prompt Understanding	Visual Prompt Understanding, Som (Set-of-marks) Recognition	2
Pixel Level Perception	Image Matting	1
Visual Recognition	Color Recognition, Abstract Visual Recognition	2
Discipline Knowledge Reasoning	Science, Tech Engineering, Health Medicine, Humanities Social Science, Business, Art Design	6
Cluster ID: 2		

## MMT-Bench

Table A2 – continued from previous page

Meta-Task	Subtask	# subtasks
Visual Recognition	Waste recognition, Logo and Brand Recognition, Animals Recognition, Weapon Recognition, Celebrity Recognition, Shape Recognition, Age Gender Race Recognition, Rock Recognition, Painting Recognition, Gesture Recognition, Vehicle Recognition, Astronomical Recognition, Fashion Recognition, Musical Instrument Recognition, Disaster Recognition, Sports Recognition, Building Recognition, Texture Material Recognition, Plant Recognition, Film and Television Recognition, Animated Character Recognition, Electronic Object Recognition, Scene Recognition, National Flag Recognition, Profession Recognition, Weather Recognition, Food Recognition	27
Relation Reasoning	Human Object Interaction Recognition, Human Interaction Understanding	2
Action Recognition	Image-based Action Recognition, Sign Language Recognition, General Action Recognition	4
Emotion	Scene Emotion Recognition, Artwork Emotion Recognition, Facial Expression Recognition, Micro Expression Recognition, Body Emotion Recognition	5
Image Evaluation Judgement	Lvlm Response Judgement	1
Visual Commonsense Reasoning	WhoopsVQA	1
Hallucination	Attribute Hallucination	1
Counting	Counting by Visual Prompting, Crowd Counting	2
Medical Understanding	Other Biological Attributes	1
Autonomous Driving	Traffic Sign Understanding	1
OCR	Font Recognition, Scene Text Recognition	2
Pixel Level Perception	Pixel Recognition	1
Anomaly Detection	Face Mask Anomaly Detection	1
Multiple Image Analysis	Spot the Diff	1
Visual Captioning	Instance Captioning	1
Doc Understanding	Clock Reading, Doc VQA	2
Meme Understanding	Meme Image Understanding	1
Cluster ID: 3		
Medical Understanding	Medical Modality Recognition, Lesion Grading, Disease DiagnoseAnatomy Identification	3
Visual Captioning	Multiple Image Captioning, Writing Poetry from Image	2
Emotion	Facial Expression Change Recognition	1
Visual Recognition	Image Season Recognition, Sculpture Recognition, Chemical Apparatus Recognition, Landmark Recognition, Religious Recognition	5
Hallucination	Relation Hallucination	1
Relation reasoning	Social Relation Recognition	1

MMT-Bench

Table A2 – continued from previous page

Meta-Task	Subtask	# subtasks
OCR	Handwritten Text Recognition	1
Temporal Understanding	Temporal Anticipation	1
Cluster ID: 4		
Intelligence Quotient Test	Ravens Progressive Matrices	1
Temporal Understanding	Temporal Localization	1
Autonomous Driving	Traffic Participants Understanding, Temporal Sequence Understanding, Multiple View Image Understanding	3
Counting	Counting by Category, Counting by Reasoning	2
Hallucination	Order Hallucination	1
Doc Understanding	Visual Document Information Extraction, Chart VQA	2
Action Recognition	Action Quality Assessment,	2
3D	3D Cad Recognition, 3D indoor recognition	2
Anomaly Detection	Industrial Produce Anomaly Detection	1
Image Evaluation Judgement	Image Quality Assessment	1
Low Level Vision	Depth Estimation	1
Cluster ID: 5		
Multiple Image Analysis	Spot the Similarity	1
Visual Illusion	Color Assimilation, Geometrical Relativity, Color Constancy, Color Contrast, Geometrical Perspective	5
Autonomous Driving	Traffic Light Understanding	1
Visual Recognition	Deepfake Detection	1
Anomaly Detection	Helmet Anomaly Detection	1
Cluster ID: 6		
Image Retrieval	Vehicle Retrieval, Image2image Retrieval, Sketch2image Retrieval, Face Retrieval, Text2image Retrieval, Handwritten Retrieval, Person Reid	7
Image-to-image translation	Image Colorization	1
Cluster ID: 7		
Visual Code	Eqn2latex,	2
Keypoint Detection	Clothes Keypoint Detection	1
OCR	Handwritten Math Expression recognition	1
Pixel Level Perception	Interactive Segmentation	1
Temporal Understanding	Temporal Ordering	1
Visual Captioning	Image Dense Captioning	1
Action Recognition	Gaze Estimation	1

## MMT-Bench

Table A2 – continued from previous page

Meta-Task	Subtask	# subtasks
Cluster ID: 8		
Localization	Salient Object Detection RGB, Camouflage Object Detection, Face Detection, Object Detection, Small Object Detection, Salient Object Detection RGBD, Rotated Object Detection, Remote Sensing Object Detection, Transparent Object Detection	9
Visual Grounding	Referring Detection, Reason Seg	2
Cross Image Matching	Point Tracking, One Shot Detection,	3
Image-to-image Translation	Jigsaw Puzzle Solving	1
Cross Image Catching	Single Object Tracking	1
Pixel Level Perception	Pixel Localization	1
Cluster ID: 9		
GUI Navigation	Web Shopping, GUI General, Google Apps, GUI Install	4
Cluster ID: 10		
Visual Captioning	Multiple Instance Captioning, Image Captioning Paragraph, Image Captioning	3
Anomaly Detection	Traffic Anomaly Detection	1
Doc Understanding	Chart to text	1
Hallucination	Exist Hallucination	1
Relation Reasoning	Scene Graph Recognition	1
Embodied AI	Navigation	1
Anomaly Detection	Behavior Anomaly Detection	1
Cluster ID: 11		
Doc Understanding	Table Structure Recognition, Chart to Table	2
Keypoint Detection	Furniture Keypoint Detection, Vehicle Keypoint Detection, Human Keypoint Detection, Animal Keypoint Detection	4
Pixel Level Perception	Polygon Localization,	2
Temporal Understanding	Next Image Prediction	1
Visual Code	Sketch2code, Screenshot2code	2
Cluster ID: 12		
Meme Understanding	Meme Video Understanding	1
Temporal Understanding	Mevis	1
Visual Captioning	Video Captioning	1

# MMT-Bench

Table A3. MMT-Bench subtask details (part 1): including sample number, visual input types, evaluated LVLMM capabilities and data source. "Python Generated" indicates the data is constructed from Python programs. "Internet Data" indicates the data is collected from the internet under license.

Subtask Name	Sample Num	Visual Input Type	Capability	Data Source
<b>Visual Grounding</b>				
Reason Seg	196	Natural Image	Visual Reasoning, Visual Localization	ReasonSeg (Lai et al., 2023)
Referring Detection	200	Natural Image	Visual Localization	RefCOCO (Yu et al., 2016)
<b>Doc Understanding</b>				
Doc Vqa	200	Text-rich Image	Document Understanding	InfographicVQA (Mathew et al., 2022)
Visual Document Information Extraction	200	Text-rich Image	Document Understanding	SROIE (Huang et al., 2019)
Chart To Text	200	Chart Image	Document Understanding	Chart-to-text (Obeid & Hoque, 2020), ChartSumm (Rahman et al., 2023)
Chart To Table	200	Chart Image	Document Understanding	ChartQA (Masry et al., 2022)
Clock Reading	200	Abstract Image	Visual Recognition, Document Understanding	Python Generated
Chart Vqa	200	Chart Image	Document Understanding	ChartQA (Masry et al., 2022)
Table Structure Recognition	46	Chart Image	Document Understanding	SciTSR (Chi et al., 2019)
<b>Action Recognition</b>				
Gaze Estimation	200	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	Columbia Gaze Data Set (Smith et al., 2013)
Image Based Action Recognition	200	Natural Image	Visual Recognition	POLAR (Ma & Liang, 2019), HAR (Anguita et al., 2013)
General Action Recognition	200	Natural Image	Visual Recognition, Multi-Images Analysis	Kinetics400 (Kay et al., 2017)
Action Quality Assessment	200	Natural Image	Visual Recognition, Multi-Images Analysis, Expert Knowledge Utilization	UNLV Olympic Scoring (Parmar & Tran Morris, 2017), AQA-7 (Parmar & Morris, 2019)
Sign Language Recognition	200	Natural Image	Visual Recognition, Multi-Images Analysis	WLASL (Li et al., 2020a)
<b>Localization</b>				
Remote Sensing Object Detection	200	Remote Sensing Image	Visual Recognition, Visual Localization	VisDrone (Zhu et al., 2021), DIOR (Zhan et al., 2023)
Rotated Object Detection	90	Remote Sensing Image	Visual Recognition, Visual Localization	SSDD (Zhang et al., 2021), DOTA (Ding et al., 2021a)
Small Object Detection	200	Natural Image	Visual Recognition, Visual Localization	SOD4Bird (Kondo et al., 2023), TinyPerson (Yu et al., 2020), Drone2021 (Kondo et al., 2023)
Camouflage Object Detection	200	Natural Image	Visual Recognition, Visual Localization	NCAK (Lv et al., 2021), COD10K (Fan et al., 2020)
Salient Object Detection Rgb	200	Natural Image, Depth Map	Visual Localization	DES (Cheng et al., 2014b), NJU2K (Ju et al., 2015)
Transparent Object Detection	200	Natural Image	Visual Recognition, Visual Localization	Trans10K (Xie et al., 2020), Transparent Object Images (LABS, 2023b)
Face Detection	200	Natural Image	Visual Recognition, Visual Localization	WIDER FACE (Yang et al., 2016), FDDB (Jain & Learned-Miller, 2010)
Object Detection	200	Natural Image	Visual Recognition, Visual Localization	COCO (Lin et al., 2014a), Pascal VOC (Everingham et al., 2015)
Salient Object Detection Rgb	200	Natural Image	Visual Localization	MSRA10K (Cheng et al., 2014a), DUTS (Wang et al., 2017)
<b>Visual Recognition</b>				
Deepfake Detection	200	Natural Image, Synthetic Image	Visual Recognition, Visual Reasoning, Expert Knowledge Utilization	Celeb-DF (v2) (Li et al., 2019), AI Recognition Dataset (SUPERPOTATO, 2024), FF++ (Rosler et al., 2019)
Weather Recognition	194	Natural Image	Visual Recognition	Weather Image Recognition (BHATHENA, 2021), MWD (Ghenimi, 2018)
Image Season Recognition	200	Natural Image	Visual Recognition	Image Season Recognition (Artem, 2013)
Gesture Recognition	200	Natural Image	Visual Recognition	CNNgestureRecognizer (Singh, 2017), number-gestures-recognition (tz28, 2018)
Musical Instrument Recognition	200	Natural Image	Visual Recognition	30 Musical Instruments Image Classification (GERRY, 2021), Music Instruments Classification (TYAGI, 2020)
Food Recognition	200	Natural Image	Visual Recognition	Food-101 (Bossard et al., 2014), Fruits and Vegetables Image Recognition Dataset (SETH, 2022)
Landmark Recognition	50	Natural Image	Visual Recognition, Expert Knowledge Utilization	Internet Data
Scene Recognition	200	Natural Image	Visual Recognition	Indoor Scene Recognition (Quattoni & Torralba, 2009), Places365 (Zhou et al., 2017)
Animals Recognition	200	Natural Image	Visual Recognition	Animals-90 (BANERJEE, 2022), Animals-10 (ALESSIO, 2020), Animals-151 (MENON, 2022)
Chemical Apparatus Recognition	200	Natural Image	Visual Recognition	Annotated Chemical Apparatus Image Dataset (Sasaki et al., 2024)
Rock Recognition	200	Natural Image	Visual Recognition, Expert Knowledge Utilization	Rock Classification Dataset (Salman Ibraheem Eamus, 2021), Rock Images (GAJARE, 2022)
Fashion Recognition	200	Natural Image	Visual Recognition	Fashion-MNIST (Xiao et al., 2017), DeepFashion (Liu et al., 2016b)
Logo And Brand Recognition	200	Natural Image	Visual Recognition	Fake/Real Logo Detection Dataset (CHUKS, 2023), FlickrSportLogos-10 (sunshine joe, 2018)
Astronomical Recognition	94	Natural Image	Visual Recognition, Expert Knowledge Utilization	Internet Data
Painting Recognition	200	Painting Image	Visual Recognition, Expert Knowledge Utilization	WikiArt (Wik, 2022), Best Artworks of All Time (ICARO, 2019), Van Gogh Paintings (INNAT, 2022), Internet Data (Chinese Painting)
Color Recognition	200	Synthetic Image	Visual Recognition	Python Generated
Plant Recognition	200	Natural Image	Visual Recognition	Flower Photos (tensorflow)
Shape Recognition	200	Synthetic Image	Visual Recognition	2D Geometric Shapes Dataset (El Korchi & Ghanou, 2020), DALLE Generated (Betker et al., 2023)
Profession Recognition	200	Natural Image	Visual Recognition	IdenProf (OLAFENWA, 2018)
Building Recognition	200	Natural Image	Visual Recognition, Expert Knowledge Utilization	Oxford Buildings (Philbin et al., 2007), ArchitecturalStyle Recognition (Danci, 2019)
Electronic Object Recognition	200	Natural Image	Visual Recognition	E-Waste/TAMRAKAR, 2023), Electronics Object Image Dataset (LABS, 2023a)
Sports Recognition	200	Natural Image	Visual Recognition	100 Sports Image Classification (GERRY, 2023), Cricket-Football-Baseball Classification Image Dataset (SAHA, 2022)
Disaster Recognition	200	Natural Image	Visual Recognition	Disaster Images Dataset (VERMA, 2021), MEDIC (Alam et al., 2023)
Celebrity Recognition	200	Natural Image	Visual Recognition	Labeled Faces in the Wild (Huang et al., 2007)
Vehicle Recognition	200	Natural Image	Visual Recognition	TAU Vehicle Type Recognition Competition (Huttunen, 2019), Vehicle Type Recognition (JOHANN, 2023)
National Flag Recognition	200	Synthetic Image	Visual Recognition	Internet Data
Abstract Visual Recognition	200	Abstract Image	Visual Recognition	QuickDraw (mirayinteractive, 2014), ImageNet-Sketch (Wang et al., 2019)
Animated Character Recognition	200	Synthetic Image	Visual Recognition	Anime Characters Personality And Facial Images (MA, 2023), Moeimouto Face Dataset (Moe)
Texture Material Recognition	200	Natural Image	Visual Recognition	OpenSurfaces (Bell et al., 2013), Kylberg (Kylberg, 2011), KTH (Mallikarjuna et al., 2006), UIUC (Lazebnik et al., 2005)
Film And Television Recognition	200	Synthetic Image	Visual Recognition, Expert Knowledge Utilization	Internet Data, Movie Posters (HARI, 2020)
Sculpture Recognition	50	Natural Image	Visual Recognition, Expert Knowledge Utilization	Internet Data
Age Gender Race Recognition	200	Natural Image	Visual Recognition	FairFace
Weapon Recognition	200	Natural Image	Visual Recognition	Weapon Detection Dataset (SANYAL, 2023), OWAD (Haq et al., 2022)
Religious Recognition	200	Natural Image, Synthetic Image	Visual Recognition, Expert Knowledge Utilization	Religious Symbols-Image Classification (KUMARUJJAWAL1234, 2023), Traditional Chinese god stat (Huang et al., 2023)
Waste Recognition	200	Natural Image	Visual Recognition, Expert Knowledge Utilization	Waste Classification Data (SEKAR, 2019), Garbage Classification (MOHAMED, 2021)

# MMT-Bench

Table A4. MMT-Bench subtask details (part 2): including sample number, visual input types, evaluated LVLm capabilities and data source. "Python Generated" indicates the data is constructed from Python programs. "Internet Data" indicates the data is collected from the internet under license.

Subtask Name	Sample Num	Visual Input Type	Capability	Data Source
<b>GUI Navigation</b>				
GUI General	200	Screenshot Image	Visual Reasoning, Visual Localization	AITW (Rawles et al., 2023)
Google Apps	200	Screenshot Image	Visual Reasoning, Visual Localization	AITW (Rawles et al., 2023)
Web Shopping	200	Screenshot Image	Visual Reasoning, Visual Localization	AITW (Rawles et al., 2023)
GUI Install	200	Screenshot Image	Visual Reasoning, Visual Localization	AITW (Rawles et al., 2023)
<b>OCR</b>				
Font Recognition	200	Text-rich Image	OCR	AdobeVFR (Wang et al., 2015)
Handwritten Text Recognition	100	Text-rich Image	OCR	IAM (Marti & Bunke, 2002)
Handwritten Mathematical Expression Recognition	100	Text-rich Image	OCR	HME100K (Yuan et al., 2022), CROHME 2014 (Mouchere et al., 2014)
Scene Text Recognition	200	Natural Image, Text-rich Image	OCR	IIT 5K-word (Mishra et al., 2012), ICDAR 2013 (Karatzas et al., 2013)
<b>Image-to-image Translation</b>				
Jigsaw Puzzle Solving	200	Natural Image	Visual Recognition, Visual Reasoning	WikiArt (Wik, 2022), COCO (Lin et al., 2014a)
Image Colorization	200	Natural Image	Pixel Perception	Landscape Color and Grayscale Images (GHIMIRE, 2021), Image Colorization (SHETTY, 2018)
<b>Temporal Understanding</b>				
Next Img Prediction	200	Visual Mark	Temporal Understanding	Moving MNIST (Srivastava et al., 2015)
Mevis	200	Natural Image	Temporal Understanding	MeVIS (Ding et al., 2023)
Temporal Anticipation	200	Natural Image	Temporal Understanding	STAR (Wu et al., 2021)
Temporal Ordering	200	Natural Image	Temporal Understanding	Pouring (Sermanet et al., 2016), Penn Action (Zhang et al., 2013)
Temporal Localization	193	Natural Image	Temporal Understanding	THUMOS14 (Iliadis et al., 2017)
<b>Relation Reasoning</b>				
Social Relation Recognition	200	Natural Image	Visual Recognition, Visual Reasoning	Social Relation Dataset (Zhang et al., 2015)
Human Object Interaction Recognition	200	Natural Image	Visual Recognition, Visual Reasoning	HICO-DET (Chao et al., 2018)
Scene Graph Recognition	200	Natural Image	Visual Recognition, Visual Reasoning	VSR (Liu et al., 2023a)
Human Interaction Understanding	200	Natural Image	Visual Recognition, Visual Reasoning	HICO-DET (Chao et al., 2018), BIT (Kong et al., 2012)
<b>Discipline Knowledge Reasoning</b>				
Science	127	Scientific Diagram	Visual Reasoning, Expert Knowledge Utilization	MMMU (Yue et al., 2023a)
Health Medicine	140	Natural Image, Chart Image, Medical Image	Visual Reasoning, Expert Knowledge Utilization	MMMU (Yue et al., 2023a)
Art Design	110	Synthetic Image, Text-rich Image, Painting Image	Visual Reasoning, Expert Knowledge Utilization	MMMU (Yue et al., 2023a)
Humanities Social Science	112	Synthetic Image, Painting Image	Visual Reasoning, Expert Knowledge Utilization	MMMU (Yue et al., 2023a)
Tech Engineering	182	Chart Image, Scientific Diagram	Visual Reasoning, Expert Knowledge Utilization	MMMU (Yue et al., 2023a)
Business	120	Text-rich Image, Chart Image	Visual Reasoning, Expert Knowledge Utilization	MMMU (Yue et al., 2023a)
<b>Intelligence Quotient Test</b>				
Ravens Progressive Matrices	200	Scientific Diagram	Visual Reasoning, Expert Knowledge Utilization	RAVEN (RAV)
<b>Embodied AI</b>				
Navigation	200	Synthetic Image	Visual Reasoning, Expert Knowledge Utilization	Kitchen Worlds (Zhu Tian Yang, 2022)
<b>Emotion</b>				
Facial Expression Change Recognition	200	Natural Image	Visual Recognition, Temporal Understanding	Facial Emotion Recognition Dataset (ROMAN, 2023), FERF-DB (Aneja et al., 2016)
Scene Emotion Recognition	200	Natural Image	Visual Recognition	ArtPhoto (Machajdik & Hanbury, 2010)
Micro Expression Recognition	200	Natural Image	Visual Recognition	SAMM (Davison et al., 2016), CASME (Yan et al., 2013)
Artwork Emotion Recognition	200	Painting Image	Visual Recognition	ArtEmis (Mohamed et al., 2022)
Body Emotion Recognition	200	Natural Image	Visual Recognition	CAER (Lee et al., 2019), EMOTIC (Kosti et al., 2019)
Facial Expression Recognition	200	Natural Image	Visual Recognition	Facial Emotion Recognition Dataset (ROMAN, 2023), FERF-DB (Aneja et al., 2016)
<b>Visual Illusion</b>				
Color Constancy	72	Synthetic Image	Visual Recognition, Visual Reasoning	GVIL (Zhang et al., 2023c)
Color Assimilation	200	Synthetic Image	Visual Recognition, Visual Reasoning	GVIL (Zhang et al., 2023c)
Geometrical Relativity	200	Synthetic Image	Visual Recognition, Visual Reasoning	GVIL (Zhang et al., 2023c)
Geometrical Perspective	120	Synthetic Image	Visual Recognition, Visual Reasoning	GVIL (Zhang et al., 2023c)
Color Contrast	200	Synthetic Image	Visual Recognition, Visual Reasoning	GVIL (Zhang et al., 2023c)
<b>Meme Understanding</b>				
Meme Video Understanding	200	Natural Image	Visual Description	FunQA (Xie et al., 2023)
Meme Image Understanding	200	Synthetic Image	Visual Description	Hateful Memes (Kielas et al., 2020), MemeCap (Hwang & Shwartz, 2023)
<b>Counting</b>				
Counting By Visual Prompting	200	Natural Image	Visual Recognition, Counting	FSC147 (Ranjan et al., 2021)
Counting By Category	800	Natural Image	Visual Recognition, Counting	TallyQA (Acharya et al., 2019), FSC147 (Ranjan et al., 2021)
Crowd Counting	200	Natural Image	Visual Recognition, Counting	ShanghaiTech (Zhang et al., 2016), CARPK (Hsieh et al., 2017)
Counting By Reasoning	200	Natural Image	Visual Recognition, Counting	
<b>Hallucination</b>				
Order Hallucination	200	Natural Image	Visual Recognition, Visual Reasoning, Visual Description	VL-CheckList (Zhao et al., 2022)
Relation Hallucination	200	Natural Image	Visual Recognition, Visual Reasoning, Visual Description	VL-CheckList (Zhao et al., 2022)
Attribute Hallucination	200	Natural Image	Visual Recognition, Visual Reasoning, Visual Description	DeepFashion (Liu et al., 2016b)
Exist Hallucination	200	Natural Image	Visual Recognition, Visual Reasoning	COCO (Lin et al., 2014a)
<b>Image Retrieval</b>				
Person Reid	200	Natural Image	Retrieval, Multi-Images Analysis	Market-1501 (Zheng et al., 2015)
Sketch2image Retrieval	200	Natural Image	Retrieval, Multi-Images Analysis	QuickDraw (mrayinteractive, 2014), DomainNet (Peng et al., 2019)
Face Retrieval	200	Natural Image	Retrieval, Multi-Images Analysis	Labeled Faces in the Wild (Huang et al., 2007), CelebA (Liu et al., 2015)
Handwritten Retrieval	200	Text-rich Image	Retrieval, OCR, Multi-Images Analysis	IAM (Marti & Bunke, 2002)
Vehicle Retrieval	200	Natural Image	Retrieval, Multi-Images Analysis	VeRi-776 (Liu et al., 2016a)
Image2image Retrieval	200	Natural Image	Retrieval, Multi-Images Analysis	TinyImageNet (Le & Yang, 2015), Places365 (Zhou et al., 2017)
Text2image Retrieval	200	Natural Image	Retrieval, Multi-Images Analysis	CUB-200-2011 (Wah et al., 2011), Oxford102 Flower (Nilsback & Zisserman, 2008)
<b>Visual Prompt Understanding</b>				
Som Recognition	199	Natural Image, Visual Mark	Visual Recognition, Visual Reasoning, Visual Localization, Visual Prompting Understanding	Som-Bench (Yang et al., 2023a)
Visual Prompt Understanding	200	Natural Image, Visual Mark	Visual Recognition, Visual Reasoning, Visual Localization, Visual Prompting Understanding	VIP-Bench (Cai et al., 2023)

## MMT-Bench

Table A5. MMT-Bench subtask details (part 3): including sample number, visual input types, evaluated LVLm capabilities and data source. "Python Generated" indicates the data is constructed from Python programs. "Internet Data" indicates the data is collected from the internet under license.

Subtask Name	Sample Num	Visual Input Type	Capability	Data Source
<b>Anomaly Detection</b>				
Industrial Produce Anomaly Detection	200	Natural Image	Visual Recognition, Counting	MPDD (stepanje, 2021), MVTec AD (Bergmann et al., 2019)
Face Mask Anomaly Detection	200	Natural Image	Visual Recognition	Face Mask Detection (KUMAR, 2021), Face Mask Usage (NOGRA, 2022)
Helmet Anomaly Detection	200	Natural Image	Visual Recognition, Visual Localization	Helmet Detection (LARXEL, 2020a), HELMET (Hanhe Lin, 2020)
Behavior Anomaly Detection	200	Natural Image	Visual Recognition, Multi-Images Analysis	ShanghaiTech Campus (Liu et al., 2018b), Avenue (Lu et al., 2013)
Traffic Anomaly Detection	200	Natural Image	Visual Recognition	ShanghaiTech Campus (Liu et al., 2018b)
<b>Keypoint Detection</b>				
Furniture Keypoint Detection	200	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	MP-100 (Xu et al., 2022)
Human Keypoint Detection	200	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	MPII Human Pose Dataset (Andriluka et al.)
Clothes Keypoint Detection	200	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	DeepFashion (Liu et al., 2016b)
Animal Keypoint Detection	200	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	AP-10K (Yu et al., 2021), Animal Kingdom (Ng et al., 2022)
Vehicle Keypoint Detection	92	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	MP-100 (Xu et al., 2022)
<b>Visual Commonsense Reasoning</b>				
Whoops	200	Synthetic Image	Visual Recognition, Visual Reasoning	Whoops (Bitton-Guetta et al., 2023)
<b>Visual Code</b>				
Eqn2latex	200	Text-rich Image, Scientific Diagram	OCR, Document Understanding, Expert Knowledge Utilization	Im2latex-90k (Singh, 2018)
Screenshot2code	200	Screenshot Image	Document Understanding, Expert Knowledge Utilization	pix2code (Beltramelli, 2017)
Sketch2code	200	Scientific Diagram	Document Understanding, Expert Knowledge Utilization	Sketch2code (anchen li, 2018)
<b>Image Evaluation Judgement</b>				
Image Quality Assessment	200	Natural Image	Visual Reasoning	LIVE (Dr. Hamid Rahim Sheikh, 2006)
Lvlm Response Judgement	200	Synthetic Image, Chart Image	Visual Reasoning	LVLm-eHub (Xu et al., 2023)
<b>Pixel Level Perception</b>				
Polygon Localization	200	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	COCO (Lin et al., 2014a)
Interactive Segmentation	141	Natural Image	Visual Localization, Pixel Perception	Berkeley Segmentation (Martin et al., 2001), DAVIS2017 (Pont-Tuset et al., 2017)
Depth Estimation	200	Natural Image	Pixel Perception, 3D Perception	KITTI (Geiger et al., 2012), NYU-Depth (Silberman et al., 2012)
Pixel Recognition	200	Natural Image	Visual Recognition, Pixel Perception	COCO (Lin et al., 2014a)
Pixel Localization	200	Natural Image	Visual Recognition, Visual Localization, Pixel Perception	COCO (Lin et al., 2014a)
Image Matting	200	Natural Image	Pixel Perception	AM-2K (Li et al., 2022), AIM-500 (Li et al., 2021)
<b>Multiple Image Analysis</b>				
Spot The Similarity	200	Natural Image, Synthetic Image	Multi-Images Analysis	TTL (Rosenfeld et al., 2018), Images Alike (CJOV, 2021)
Spot The Diff	200	Natural Image	Multi-Images Analysis	Spot-the-diff (Jhamtani & Berg-Kirkpatrick, 2018)
<b>3D</b>				
3D Cad Recognition	200	3d Image	Multi-Images Analysis, 3D Perception	ModelNet40 (Wu et al., 2015)
3D Indoor Recognition	200	3d Image	Multi-Images Analysis, 3D Perception	ScanObjectNN (Uy et al., 2019)
<b>Medical Understanding</b>				
Anatomy Identification	200	Medical Image	Visual Recognition, Expert Knowledge Utilization	OmniMedVQA (Hu et al., 2024)
Medical Modality Recognition	200	Medical Image	Visual Recognition, Expert Knowledge Utilization	OmniMedVQA (Hu et al., 2024)
Other Biological Attributes	200	Medical Image	Visual Recognition, Expert Knowledge Utilization	OmniMedVQA (Hu et al., 2024)
Disease Diagnose	200	Medical Image	Visual Recognition, Expert Knowledge Utilization	OmniMedVQA (Hu et al., 2024)
Lesion Grading	200	Medical Image	Visual Recognition, Expert Knowledge Utilization	OmniMedVQA (Hu et al., 2024)
<b>Cross Image Matching</b>				
One Shot Detection	200	Natural Image	Visual Localization	FSS-1000 (Li et al., 2020b), PACO-Part (Liu et al., 2023e)
Point Tracking	200	Natural Image	Visual Localization	TAP-Vid (Doersch et al., 2022)
Single Object Tracking	200	Natural Image	Visual Localization	OVIS (Qi et al., 2022), YouTube-VIS 2019 (Yang et al., 2019)
<b>Visual Captioning</b>				
Video Captioning	200	Natural Image	Visual Description, Temporal Understanding	MSR-VTT (Xu et al., 2016), MSVD (Chen & Dolan, 2011)
Image Captioning Paragraph	200	Natural Image	Visual Description	Paragraphs (Krause et al., 2017)
Image Captioning	200	Natural Image	Visual Description	COCO (Lin et al., 2014a)
Instance Captioning	200	Natural Image	Visual Description	RefCOCOg (Yu et al., 2016), Visual Genome (Krishna et al., 2017a)
Image Dense Captioning	197	Natural Image	Visual Description	Visual Genome (Krishna et al., 2017a)
Multiple Instance Captioning	200	Natural Image	Visual Description	Flickr30k (Young et al., 2014)
Multiple Image Captioning	200	Natural Image	Visual Description, Multi-Images Analysis	VIST (Huang et al., 2016)
Writing Poetry From Image	200	Natural Image, Text-rich Image	Visual Description	MultiM-Poem (Liu et al., 2018a)
<b>Autonomous Driving</b>				
Traffic Participants Understanding	200	Natural Image	Counting	nuScenes (Caesar et al., 2020), Waymo (Sun et al., 2020)
Multiple View Image Understanding	200	Natural Image	Visual Reasoning, Multi-Images Analysis, Counting	nuScenes (Caesar et al., 2020), Waymo (Sun et al., 2020)
Traffic Sign Understanding	200	Natural Image	Visual Reasoning, Expert Knowledge Utilization	Mapillary Traffic Sign Dataset (Ertler et al., 2020), Road Sign Detection (LARXEL, 2020b)
Temporal Sequence Understanding	200	Natural Image	Visual Reasoning, Temporal Understanding	nuScenes (Caesar et al., 2020), Waymo (Sun et al., 2020)
Traffic Light Understanding	200	Natural Image	Visual Recognition	LISA-TL (JENSEN, 2018), S2TLD (Yang et al., 2022)

Table A6. Model architecture of 32 LVLMs evaluated on MMT-Bench. \* Both GPT-4o and GPT-4V use the low-resolution mode.

Models	Parameters	Vision Encoder	LLM
GPT-4o* (OpenAI, 2024)	-	-	-
GPT-4V* (20240409) (Yang et al., 2023b)	-	-	-
GeminiProVision (Team, 2023a)	-	-	-
QWen-VL-Max (Team, 2023c)	-	-	-
QWen-VL-Plus (Team, 2023c)	-	-	-
Claude3V-Haiku (Anthropic, 2023)	-	-	-
LLaVA-Next-34B (Liu et al., 2024a)	34.8B	CLIP ViT-L/14	Nous-Hermes-2-Yi-34B
LLaVA-Next-13B (Liu et al., 2024a)	13.4B	CLIP ViT-L/14	Vicuna-v1.5-13B
LLaVA-Next-7B (Liu et al., 2024a)	7.1B	CLIP ViT-L/14	Vicuna-v1.5-7B
Yi-VL-34B (AI et al., 2024)	34.6B	CLIP ViT-H/14	Nous-Hermes-2-Yi-34B
Yi-VL-6B (AI et al., 2024)	6.6B	CLIP ViT-H/14	Yi-6B
InternVL-Chat-V1.2 (Chen et al., 2023b)	40B	InternViT-6B	Nous-Hermes-2-Yi-34B
DeepSeek-VL-7B (Lu et al., 2024)	7.3B	SAM-B & SigLIP-L	DeekSeek-7B
Monkey (Li et al., 2023d)	9.8B	CLIP-ViT-BigHuge	Qwen-7B
XComposer (Zhang et al., 2023a)	8B	EVA-CLIP-G	InternLM-7B
XComposer2 (Dong et al., 2024)	7B	CLIP ViT-L/14	InternLM2-7B
ShareGPT4V (Chen et al., 2023a)	7.2B	CLIP ViT-L/14	Vicuna-v1.5-7B
SharedCaptioner (Chen et al., 2023a)	8B	EVA-G	InternLM-7B
mPLUG-Owl2 (Ye et al., 2023b)	8.2B	CLIP ViT-L/14	LLaMA2-7B
LLaVA-v1.5-7B (Liu et al., 2023c;b)	7.2B	CLIP ViT-L/14	Vicuna-v1.5-7B
LLaVA-v1.5-13B (Liu et al., 2023c;b)	13.4B	CLIP ViT-L/14	Vicuna-v1.5-13B
LLaVA-InternLM2-7B (Contributors, 2023c)	8.1B	CLIP ViT-L/14	InternLM2-7B
LLaVA-InternLM-7B (Contributors, 2023c)	7.6B	CLIP ViT-L/14	InternLM-7B
LLaVA-v1.5-7B-Xtuner (Contributors, 2023c)	7.2B	CLIP ViT-L/14	Vicuna-v1.5-7B
LLaVA-v1.5-13B-Xtuner (Contributors, 2023c)	13.4B	CLIP ViT-L/14	Vicuna-v1.5-13B
LLaMA-Adapter-v2 (Gao et al., 2023)	7B	CLIP-ViT-L/14	LLaMA-7B
VisualGLM (Ding et al., 2021b)	8B	EVA-CLIP	ChatGLM-6B
CogVLM (Wang et al., 2023)	17B	EVA-CLIP-E	Vicuna-v1.5-7B
TransCore-M (Contributors, 2023b)	13.4B	CLIP ViT/L-14	PCITransGPT-13B
RBDash-v1 (RBDash-Team, 2023)	13.4B	CLIP ViT-L/14	Vicuna-v1.5-13B
BLIP2 (Li et al., 2023b)	12.1B	EVA-CLIP ViT-G/14	Flan-T5-XXL
QWenVL (Bai et al., 2023)	9.6B	CLIP ViT-G/16	QWen-7B

Table A7. Abbreviations for tasks requiring multiple images as inputs (part one). Here we also present the designed prompt examples we used for Single-Image LVLMs and Multi-Images LVLMs.

Task Abbreviation	Task Name	Prompt Example for Single Image LVLMs	Prompt example for Multiple Image LVLMs
AQS	action quality assessment	Question: <image><image><image><image> What is the most probable action quality assessment number obtained by the person in the video? Options: A. 35.99 B. 28.0 C. 11.27 D. 44.98	Question: <image><image><image><image> What is the most probable action quality assessment number obtained by the person in the video? Options: A. 35.99 B. 28.0 C. 11.27 D. 44.98
FEER	facail expression change recognition	Question: <image><image>What is the change of expression from the first image to the second image? Options: A. disgust to happy B. happy to sadness C. anger to surprise D. disgust to fear	Question: What is the change of expression from Image 1: <image>to Image 2: <image>? Options: A. disgust to happy B. happy to sadness C. anger to surprise D. disgust to fear
FR	face retrieval	Question: <image><image><image><image> <image>Please retrieve the most similar person to the query in the candidates. The first image is the query image and the remaining images are candidates from Candidate 1 to Candidate 4. Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4	Question: Please retrieve the most similar person to the query: <image>in the candidates: Candidate 1: <image>, Candidate 2: <image>, Candidate 3: <image>, Candidate 4: <image>. Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4
GAR	general action recognition	Question: <image><image><image><image> What is the action performed by the person in the video? Options: A. rock scissors paper B. sword fighting C. fencing D. balloon blowing	Question: <image><image><image><image> What is the action performed by the person in the video? Options: A. rock scissors paper B. sword fighting C. fencing D. balloon blowing
HR	handwritten retrieval	Question: <image><image><image><image> <image>Please retrieve the most similar handwritten text snapshot to the query in the candidates. The first image is the query image and the remaining images are candidates from Candidate 1 to Candidate 4. Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4	Question: Please retrieve the most similar handwritten text snapshot to the query: <image>in the candidates: Candidate 1: <image>, Candidate 2: <image>, Candidate 3: <image>, Candidate 4: <image>. Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4
I2IR	image2image retrieval	Question: <image><image><image><image> <image>Please retrieve the most similar scene to the query in the candidates. The first image is the query image and the remaining images are candidates from Candidate 1 to Candidate 4. Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4	Question: Please retrieve the most similar scene to the query: <image>in the candidates: Candidate 1: <image>, Candidate 2: <image>, Candidate 3: <image>, Candidate 4: <image>. Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4
IC	image colorization	Question: <image><image><image><image> The following images are candidates from Candidate 1 to Candidate 4, which are from the same picture consisting of four styles: grayscale, original, warm, and sepia. Which one is the original picture? Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4	Question: The following images: Candidate 1: <image>, Candidate 2: <image>, Candidate 3: <image>, Candidate 4: <image>, are from the same picture, which consists of four styles: grayscale, original, warm, and sepia. Which one is the original picture? Options: A. Candidate 1 B. Candidate 2 C. Candidate 3 D. Candidate 4
MVU	meme video understanding	Question: <image><image><image><image> Please generate a description for this meme Options: A. From beneath the toilet door panel, a hand is reaching out with an upward-facing palm to receive chopsticks and a spoon from someone outside. B. The hand is asking for help to get out of the bathroom. C. The hand is actually reaching out for a handshake. D. A person is handing over toilet paper instead of chopsticks and a spoon.	Question: <image><image><image><image> Please generate a description for this meme Options: A. From beneath the toilet door panel, a hand is reaching out with an upward-facing palm to receive chopsticks and a spoon from someone outside. B. The hand is asking for help to get out of the bathroom. C. The hand is actually reaching out for a handshake. D. A person is handing over toilet paper instead of chopsticks and a spoon.

Table A8. Abbreviations for tasks requiring multiple images as inputs (part two). Here we also present the designed prompt examples we used for Single-Image LVLMS and Multi-Images LVLMS.

Task Abbreviation	Task Name	Prompt Example for Single Image LVLMS	Prompt example for Multiple Image LVLMS
ME	mevis	<p>Question:     I have provided several frames from a video, and I will also provide a caption. Provide the output for the detected area in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input image is 1920 and the height is 945.  CAPTION: little girl feeding rabbit  Options:  A. [70, 0, 993, 1007]  B. [203, 0, 1011, 944]  C. [70, 0, 1011, 944]  D. [196, 38, 652, 277]</p>	<p>Question:     I have provided several frames from a video, and I will also provide a caption. Provide the output for the detected area in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input image is 1920 and the height is 945.  CAPTION: little girl feeding rabbit  Options:  A. [70, 0, 993, 1007]  B. [203, 0, 1011, 944]  C. [70, 0, 1011, 944]  D. [196, 38, 652, 277]</p>
MIC	multiple image captioning	<p>Question:    Describe this set of images briefly.  Options:  A. I took a cab to return to the hotel  B. the front of the mall was somewhat crowded .  i ran past them and took the escalator down .  after shopping for a few hours , i returned to the street .  i tried to catch a cab but a bush blocked me .  i decided to just walk back to my hotel .  C. the mall was empty and I took the stairs up  D. I quickly caught a bus to my hotel</p>	<p>Question: Describe this set of images:   briefly.  Options:  A. I took a cab to return to the hotel  B. the front of the mall was somewhat crowded .  i ran past them and took the escalator down .  after shopping for a few hours , i returned to the street .  i tried to catch a cab but a bush blocked me .  i decided to just walk back to my hotel .  C. the mall was empty and I took the stairs up  D. I quickly caught a bus to my hotel</p>
NIP	next img prediction	<p>Question:    Please predict the last 10 frames in the candidates of the video based on the first 10 frames of the input video. Note that the order is from left to right. The first four images are candidates from Candidate 1 to Candidate 4 and the last image shows the first 10 frames of the input video.  Options:  A. Candidate 1: last 10 frames  B. Candidate 2: last 10 frames  C. Candidate 3: last 10 frames  D. Candidate 4: last 10 frames</p>	<p>Question: Please predict the last 10 frames in the candidates: Candidate 1: , Candidate 2: , Candidate 3: , Candidate 4: , of the video: based on the first 10 frames of the input video: .  Note that the order is from left to right  Options:  A. Candidate 1: last 10 frames  B. Candidate 2: last 10 frames  C. Candidate 3: last 10 frames  D. Candidate 4: last 10 frames</p>
OSD	one shot detection	<p>Question:  According to the prompts in the Support Image (marked in red), please detect the corresponding object in the Query Image. The first image is the Support Image and the second image is the Query Image.  Provide the output for the object in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input RGB image is 224 and the height is 224.  Options:  A. [0, 0, 511, 2]  B. [0, 0, 426, 1]  C. [1, 1, 511, 2]  D. [0, 0, 499, 2]</p>	<p>Question: According to the prompts in the Support Image (marked in red): , please detect the corresponding object in the Query Image: .  Provide the output for the object in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input RGB image is 224 and the height is 224.  Options:  A. [0, 0, 511, 2]  B. [0, 0, 426, 1]  C. [1, 1, 511, 2]  D. [0, 0, 499, 2]</p>
Pre	person reid	<p>Question:    Please retrieve the most similar person to the query in the candidates. The first image is the query image and the remaining images are candidates from Candidate 1 to Candidate 4.  Options:  A. Candidate 1  B. Candidate 2  C. Candidate 3  D. Candidate 4</p>	<p>Question: Please retrieve the most similar person to the query:  in the candidates: Candidate 1: , Candidate 2: , Candidate 3: , Candidate 4: .  Options:  A. Candidate 1  B. Candidate 2  C. Candidate 3  D. Candidate 4</p>
PT	point tracking	<p>Question:  What is the position coordinates of the point with coordinates ([0.711, 0.154]) in the first image within the second image? Note that the width of the input RGB image is 256 and the height is 256.  Options:  A. [0.336, 0.241]  B. [0.754, 0.592]  C. [0.711, 0.154]  D. [0.814, 0.269]</p>	<p>Question: What is the position coordinates of the point with coordinates ([0.711, 0.154]) in Frame 1:  within the Frame 2: ?  Note that the width of the input RGB image is 256 and the height is 256.  Options:  A. [0.336, 0.241]  B. [0.754, 0.592]  C. [0.711, 0.154]  D. [0.814, 0.269]</p>

Table A9. Abbreviations for tasks requiring multiple images as inputs (part three). Here we also present the designed prompt examples we used for Single-Image LVLMS and Multi-Images LVLMS.

Task Abbreviation	Task Name	Prompt Example for Single Image LVLMS	Prompt example for Multiple Image LVLMS
SODRD	salient object detection rgbd	<p>Question: &lt;image&gt;&lt;image&gt;The first image is RGB image and the second image is the corresponding depth map. Please detect the salient foreground object in this RGB image and represent them using a single bounding box. Provide the output for the detected area in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input RGB image is 640 and the height is 480.</p> <p>Options:  A. [267, 105, 119, 209]  B. [85, 307, 65, 79]  C. [318, 294, 111, 156]  D. [267, 105, 135, 241]</p>	<p>Question: The first image is RGB image: &lt;image&gt; and the second image is the corresponding depth map: &lt;image&gt;. Please detect the salient foreground object in this RGB image and represent them using a single bounding box. Provide the output for the detected area in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input RGB image is 640 and the height is 480.</p> <p>Options:  A. [267, 105, 119, 209]  B. [85, 307, 65, 79]  C. [318, 294, 111, 156]  D. [267, 105, 135, 241]</p>
SLR	sign language recognition	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;What is the sign language gesture performed by the person in the video?</p> <p>Options:  A. fashionable  B. trendy  C. fascinating  D. cool</p>	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;What is the sign language gesture performed by the person in the video?</p> <p>Options:  A. fashionable  B. trendy  C. fascinating  D. cool</p>
SOT	single object tracking	<p>Question: &lt;image&gt;&lt;image&gt;Here is an object (marked as RED box) in the first image. Please give the coordinations of this object in the second image. Provide the output for the object in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input RGB image is 1280 and the height is 720.</p> <p>Options:  A. [148.0, 187.0, 918, 487]  B. [148.0, 187.0, 792.0, 533.0]  C. [0, 187, 792.0, 533.0]  D. [149, 451, 263, 24]</p>	<p>Question: Here is an object (marked as RED box) in the Frame 1: &lt;image&gt;. Please give the coordinations of this object in the Frame 2: &lt;image&gt;. Provide the output for the object in the format [x, y, w, h]. This format represents the bounding box, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as its width and height. Note that the width of the input RGB image is 1280 and the height is 720.</p> <p>Options:  A. [148.0, 187.0, 918, 487]  B. [148.0, 187.0, 792.0, 533.0]  C. [0, 187, 792.0, 533.0]  D. [149, 451, 263, 24]</p>
S2IR	sketch2image retrieval	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;Please retrieve the most similar image to the Query Image in the candidate Images. The first image is the query image and the remaining images are candidates from Candidate 1 to Candidate 3.</p> <p>Options:  A. Candidate 1  B. Candidate 2  C. Candidate 3</p>	<p>Question: Please retrieve the most similar image to the Query Image: &lt;image&gt; in the candidate Images: Candidate 1: &lt;image&gt;, Candidate 2: &lt;image&gt;, Candidate 3: &lt;image&gt;.</p> <p>Options:  A. Candidate 1  B. Candidate 2  C. Candidate 3</p>
SD	spot the diff	<p>Question: &lt;image&gt;&lt;image&gt;The following is a description of the differences between two pictures. Which one is incorrect?</p> <p>Options:  A. The images show different types of flowers in full bloom, with colorful petals and green leaves.  B. there is a car driving by in the right picture  C. there is a car leaving the lot in the left picture</p>	<p>Question: The following is a description of the differences between two pictures: &lt;image&gt;&lt;image&gt;. Which one is incorrect?</p> <p>Options:  A. The images show different types of flowers in full bloom, with colorful petals and green leaves.  B. there is a car driving by in the right picture  C. there is a car leaving the lot in the left picture</p>
SS	spot the similarity	<p>Question: &lt;image&gt;&lt;image&gt;Are there any similarities between the two pictures?</p> <p>Options:  A. Yes  B. No</p>	<p>Question: &lt;image&gt;&lt;image&gt;Are there any similarities between the two pictures?</p> <p>Options:  A. Yes  B. No</p>
TA	temporal anticipation	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;What will the person do next with the medicine?</p> <p>Options:  A. Apply topically  B. Inject intravenously  C. Throw away  D. Eat</p>	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;What will the person do next with the medicine?</p> <p>Options:  A. Apply topically  B. Inject intravenously  C. Throw away  D. Eat</p>
TL	temporal localization	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;Given the sequence of images, please identify the image consistent with the text description: Billiards. The image index starts from 0.</p> <p>Options:  A. Image 0  B. Image 1  C. Image 2  D. Image 3</p>	<p>Question: Given the sequence of images: Image 0: &lt;image&gt;, Image 1: &lt;image&gt;, Image 2: &lt;image&gt;, Image 3: &lt;image&gt;, please identify the image consistent with the text description: Billiards.</p> <p>Options:  A. Image 0  B. Image 1  C. Image 2  D. Image 3</p>

Table A10. Abbreviations for tasks requiring multiple images as inputs (part four). Here we also present the designed prompt examples we used for Single-Image LVLMs and Multi-Images LVLMs.

Task Abbreviation	Task Name	Prompt Example for Single Image LVLMs	Prompt example for Multiple Image LVLMs
TL	temporal localization	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>Given the sequence of images, please identify the image consistent with the text description: Billiards.</p> <p>The image index starts from 0.</p> <p>Options:</p> <p>A. Image 0</p> <p>B. Image 1</p> <p>C. Image 2</p> <p>D. Image 3</p>	<p>Question: Given the sequence of images: Image 0: &lt;image&gt;, Image 1: &lt;image&gt;, Image 2: &lt;image&gt;, Image 3: &lt;image&gt;, please identify the image consistent with the text description: Billiards.</p> <p>Options:</p> <p>A. Image 0</p> <p>B. Image 1</p> <p>C. Image 2</p> <p>D. Image 3</p>
TO	temporal ordering	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>Please predict the order of the following pictures, and give each picture a sequential index.</p> <p>This index starts from 0. The larger the index, the later the order.</p> <p>Options:</p> <p>A. [3, 0, 2, 1]</p> <p>B. [2, 0, 1, 3]</p> <p>C. [0, 2, 1, 3]</p> <p>D. [1, 3, 2, 0]</p>	<p>Question: Please predict the order of the following pictures: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;, and give each picture a sequential index.</p> <p>This index starts from 0. The larger the index, the later the order.</p> <p>Options:</p> <p>A. [3, 0, 2, 1]</p> <p>B. [2, 0, 1, 3]</p> <p>C. [0, 2, 1, 3]</p> <p>D. [1, 3, 2, 0]</p>
T2IR	text2image retrieval	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>Please find the most relevant picture among the candidate images for this description.</p> <p>The given images are candidates from Candidate 1 to Candidate 4.</p> <p>Description:</p> <p>this flower has petals that are green with stringy purple stamen</p> <p>this flower is white and blue in color, with petals that are oval shaped.</p> <p>the petals on this flower are white with an elaborate pistil.</p> <p>the flower is unique because the petals aren't separated and they have a round tip</p> <p>this flower has blue petals as well as a green and purple pistil.</p> <p>this flower has thick and pale green petals under a thick fringe of purple and white.</p> <p>this flower has petals that are white and has stringy stamen</p> <p>this flower has white oblong petals and white flat filaments.</p> <p>a flower with long and narrow petals that are white.</p> <p>a flower with long and narrow petals that are white.</p> <p>Options: A. Candidate 1</p> <p>B. Candidate 2</p> <p>C. Candidate 3</p> <p>D. Candidate 4</p>	<p>Question: Please find the most relevant picture among the candidate images: Candidate 1: &lt;image&gt;, Candidate 2: &lt;image&gt;, Candidate 3: &lt;image&gt;, Candidate 4: &lt;image&gt;, for this description.</p> <p>Description:</p> <p>this flower has petals that are green with stringy purple stamen</p> <p>this flower is white and blue in color, with petals that are oval shaped.</p> <p>the petals on this flower are white with an elaborate pistil.</p> <p>the flower is unique because the petals aren't separated and they have a round tip</p> <p>this flower has blue petals as well as a green and purple pistil.</p> <p>this flower has thick and pale green petals under a thick fringe of purple and white.</p> <p>this flower has petals that are white and has stringy stamen</p> <p>this flower has white oblong petals and white flat filaments.</p> <p>a flower with long and narrow petals that are white.</p> <p>a flower with long and narrow petals that are white.</p> <p>Options: A. Candidate 1</p> <p>B. Candidate 2</p> <p>C. Candidate 3</p> <p>D. Candidate 4</p>
3DCR	3D cad recognition	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>&lt;image&gt;&lt;image&gt;What is the category of the point cloud based on the multi-view of the point cloud?</p> <p>Options:</p> <p>A. telephone</p> <p>B. chair</p> <p>C. table</p> <p>D. sofa</p>	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>&lt;image&gt;&lt;image&gt;What is the category of the point cloud based on the multi-view of the point cloud?</p> <p>Options:</p> <p>A. telephone</p> <p>B. chair</p> <p>C. table</p> <p>D. sofa</p>
3DIR	3D indoor recognition	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>&lt;image&gt;&lt;image&gt;What is the category of the point cloud based on the multi-view of the point cloud?</p> <p>Options:</p> <p>A. sink</p> <p>B. bed</p> <p>C. cabinet</p> <p>D. bag</p>	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>&lt;image&gt;&lt;image&gt;What is the category of the point cloud based on the multi-view of the point cloud?</p> <p>Options:</p> <p>A. sink</p> <p>B. bed</p> <p>C. cabinet</p> <p>D. bag</p>
VR	vehicle retrieval	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>&lt;image&gt;Please retrieve the most similar vehicle to the query in the candidates. The first image is the query image and the remaining images are candidates from Candidate 1 to Candidate 4.</p> <p>Options:</p> <p>A. Candidate 1</p> <p>B. Candidate 2</p> <p>C. Candidate 3</p> <p>D. Candidate 4</p>	<p>Question: Please retrieve the most similar vehicle to the query: &lt;image&gt; in the candidates: Candidate 1: &lt;image&gt;, Candidate 2: &lt;image&gt;, Candidate 3: &lt;image&gt;, Candidate 4: &lt;image&gt;.</p> <p>Options:</p> <p>A. Candidate 1</p> <p>B. Candidate 2</p> <p>C. Candidate 3</p> <p>D. Candidate 4</p>
VC	video captioning	<p>Question: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;</p> <p>Please generate textual descriptions for a sequence of video frames.</p> <p>Options:</p> <p>A. a woman is speaking into a microphone</p> <p>B. a man is playing guitar on stage</p> <p>C. a man is speaking into a microphone</p> <p>D. a man is typing on a computer keyboard</p>	<p>Question: Please generate textual descriptions for a sequence of video frames: &lt;image&gt;&lt;image&gt;&lt;image&gt;&lt;image&gt;.</p> <p>Options:</p> <p>A. a woman is speaking into a microphone</p> <p>B. a man is playing guitar on stage</p> <p>C. a man is speaking into a microphone</p> <p>D. a man is typing on a computer keyboard</p>

## G. Case Study

Table A11. Table index of case study figures by meta-task with associated (error) categories for each LVLm.

Case Figure	Meta-task	Subtask	GPT-4V	GeminiProVision	InternVL-Chat
Fig. A5	Visual Recognition	Landmark Recognition	Lack of Knowledge	No Error	No Error
Fig. A6	Object Localization	Camouflaged Object Detection	Lack of Capability	Perception Error	Perception Error
Fig. A7	Pixel-level Recognition	Image Matting	Perception Error	No Error	Perception Error
Fig. A8	OCR	Handwritten Text Recognition	No Error	Perception Error	Perception Error
Fig. A9	Visual Prompt Understanding	Visual Prompt Understanding	No Error	Perception Error	Fail to Follow Instruct
Fig. A10	Retrieval	Sketch to Image Retrieval	Perception Error	No Error	Perception Error
Fig. A11	Counting	Counting by Reasoning	Perception Error	Perception Error	No Error
Fig. A12	Keypoint Detection	Human Keypoint Detection	Refuse to Answer	Perception Error	Fail to Follow Instruct
Fig. A13	Action Recognition	Sign Language Recognition	Lack of Capability	Perception Error	Perception Error
Fig. A14	Visual Hallucination	Exist Hallucination	No Error	Reasoning Error	Perception Error
Fig. A15	Anomaly Detection	Industrial Produce Anomaly Detection	Lack of Knowledge	No Error	Perception Error
Fig. A16	Image-to-Image Translation	Jigsaw Puzzle Solving	No Error	Perception Error	Perception Error
Fig. A17	Visual Summary	Image Captioning Paragraph	Perception Error	No Error	Perception Error
Fig. A18	Intelligence Quotient Test	Ravens Progressive Matrices	No Error	Reasoning Error	Reasoning Error
Fig. A19	Emotional Quotient Test	Scene Emotion Recognition	Perception Error	Reasoning Error	No Error
Fig. A20	Visual Grounding	Referring Detection	Perception Error	Perception Error	Fail to Follow Instruct
Fig. A21	Visual Commonsense Reasoning	Whoops	Reasoning Error	Perception Error	Perception Error
Fig. A22	Chart, Doc Understanding	Clock Reading	Perception Error	Perception Error	Perception Error
Fig. A23	Relation Reasoning	Scene Graph Recognition	No Error	Perception Error	No Error
Fig. A24	Meme Understanding	Meme Image Understanding	Perception Error	No Error	No Error
Fig. A25	Multi-Image Analysis	Spot the Diff	No Error	No Error	No Error
Fig. A26	Temporal Understanding	Temporal Ordering	Perception Error	No Error	Perception Error
Fig. A27	Cross-Image Matching	Single Object Tracking	Lack of Capability	Perception Error	Perception Error
Fig. A28	Visual Coding	Equation to Latex	Perception Error	Perception Error	No Error
Fig. A29	Visual Illusion	Color Constancy	Perception Error	No Error	Perception Error
Fig. A30	Image Evaluation Judgement	LVLm Response Judgement	Reasoning Error	No Error	Perception Error
Fig. A31	3D Perception	3D CAD Recognition	Lack of Capability	No Error	No Error
Fig. A32	Emodied Agent	Navigation	Fail to Follow Instruct	Fail to Follow Instruct	Fail to Follow Instruct
Fig. A33	Medical Understanding	Medical Modality Recognition	No Error	No Error	Perception Error
Fig. A34	Autonomous Driving	Traffic Light Understanding	Refuse to Answer	No Error	No Error
Fig. A35	GUI Navigation	Installation	Perception Error	Perception Error	Perception Error
Fig. A36	Discipline Knowledge Reasoning	Art and Design	Lack of Knowledge	Lack of Knowledge	Lack of Knowledge

In this section, we present a case study analysis of the error types made by GPT-4V, GeminiProVision, and InternVL-Chat on various meta-tasks in MMT-Bench. We classify the errors into the following six categories:

**Perception Error** : LVLms fail to recognize, classify or detect the objects or content in images. Most LVLms are constrained by the representation power of their visual encoders, making this the most common type of error. See examples in Fig. A6, Fig. A8, etc.

**Reasoning Error** : LVLms correctly recognize and perceive the visual content but make errors in reasoning, leading to incorrect answers. See examples in Fig. A21, Fig. A30, etc.

**Lack of Knowledge** : LVLms lack the domain-specific knowledge required to answer specialized questions, such as the location of a landmark (see Fig. A5) or the creation date of a particular painting (see Fig. A36).

**Lack of Capability** : LVLms do not have the capability to solve the corresponding tasks. This error type is particularly evident in GPT-4V, which tends to respond more honestly when it lacks the ability to handle certain tasks. In contrast, other LVLm models are inclined to generate outputs even when the accuracy rate is relatively low. See examples in Fig. A6, Fig. A13.

**Refuse to Answer** : LVLms, such as GPT-4V or Gemini, refuse to answer questions that are anthropocentric or sensitive in nature. See examples in Fig. A12, Fig. A34.

**Fail to Follow Instruct** : LVLms fail to correctly understand instructions and provide erroneous answers. For example, LVLms may not understand the specified conditions in the instruction (see Fig. A9) or may ignore the instruction altogether and instead generate a caption for the given image (see Fig. A12).

## Visual Recognition: Landmark Recognition

Question: What is the name of the landmark in the picture?

Options

- (A). Colosseum
- (B). Iguazu Falls, Brazil
- (C). Petra, Jordan
- (D). Victoria Falls, Canada

Answer: D



Lack of Knowledge

**GPT-4V:** None of the options provided correctly describe the landmark in the picture. The image shows a man-made waterfall structure that does not match the Colosseum, Iguazu Falls, Petra, or Victoria Falls. Therefore, I cannot choose any of the given options as they are all incorrect. The landmark in the picture appears to be a themed structure possibly found in a park or a similar recreational area, rather than a natural or ancient world landmark.



**GeminiProVision:** D. Victoria Falls, Canada. The waterfall shown in the picture is man-made and located in Canada. The real Victoria Falls is located on the border of Zambia and Zimbabwe.



**InternVL-Chat:** The name of the landmark in the picture is Victoria Falls, Canada.

Figure A5. A sample case of visual recognition (landmark recognition). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Object Localization: Camouflaged Object Detection

Question: Please detect all camouflaged foreground instances in this image. For each detected object, provide the output in the format [x, y, w, h]. This format represents the bounding box for each object, where [x, y, w, h] are the coordinates of the top-left corner of the bounding box, as well as the width and height of the bounding box. Note that the width of the input image is 550 and the height is 413.

Options

- (A). [0.34, 0.114, 0.575, 0.864]  
 (B). [0.384, 0.349, 0.609, 0.964]  
 (C). [0.38, 0.211, 0.615, 0.961]  
 (D). [0.378, 0.249, 0.613, 1.0], [0.38, 0.249, 0.615, 1.0], [0.118, 0.075, 0.371, 0.932]

Answer: A



✗ Lack of Capability

**GPT-4V:** I'm sorry, but I cannot assist with this request. My capabilities do not include analyzing images to provide specific coordinates for objects within them. If you have any other questions or need assistance with a different topic, feel free to ask!

✗ Perception Error

**GeminiProVision:** The correct answer is D.

The image contains a purple gallinule bird camouflaged in the tall grass. The bird is located in the center of the image and is partially obscured by the grass. The bounding box for the bird is [0.378, 0.249, 0.613, 1.0], [0.38, 0.249, 0.615, 1.0], [0.118, 0.075, 0.371, 0.932].

✗ Perception Error

**InternVL-Chat:** The correct option is D. [0.378, 0.249, 0.613, 1.0].

The image features a bird hidden in the tall grass. The bird is partially visible, with its head and neck sticking out of the grass. The bird's body is obscured by the grass, making it difficult to see.

Figure A6. A sample case of object localization (camouflaged object detection). Green highlights the right answer. Red highlights the wrong answer. [Back to Table Index](#).

## Pixel-level Recognition: Image Matting

Question: You are a professional image matting expert. What is the alpha value of the pixel point at coordinates (0.278, 0.683) in the image for image matting purposes? The alpha value represents the degree of transparency of the salient object against the background at this specific pixel. In this context, an alpha value of 0 indicates complete transparency, meaning the pixel is entirely invisible, while an alpha value of 255 represents complete opacity, meaning the pixel is fully visible. The dimensions of the input image are given as 1626 in width and 1080 in height. The coordinates of the top left corner of the image are (0, 0), and those of the bottom right corner are (1.0, 1.0).

Options

- (A). 50
- (B). 255
- (C). 84
- (D). 0

Answer: C



The referring point is marked as red dot.



## Perception Error

**GPT-4V:** I'm sorry, but I cannot provide an alpha value for a specific pixel point in an image without analyzing the actual pixel data, which requires computational tools that are not available to me. Alpha values are determined by examining the image data at the pixel level, which is not possible through a text-based interface or without image processing software.

However, if we were to estimate based on the provided image and the description of **The pixel point at coordinates (0.278, 0.683)** would likely fall on the dog, which is the salient object in the image. Since the dog appears to be fully opaque and not transparent, we could guess that the alpha value at this point would be 255, representing complete opacity.

the alpha value, we could make an educated guess.

Therefore, the answer would be:

B. 255

This is because the specified coordinates are likely to be on the dog, which is the main subject of the image and appears to be fully opaque.



**GeminiProVision:** C. 84

The alpha value of a pixel represents the degree of transparency of the salient object against the background at that specific pixel. In this context, an alpha value of 0 indicates complete transparency, meaning the pixel is entirely invisible, while an alpha value of 255 represents complete opacity, meaning the pixel is fully visible.

**The pixel at coordinates (0.278, 0.683) is located in the middle of the dog's ear. The dog's ear is semi-transparent, so the alpha value of this pixel should be between 0 and 255.**

The alpha value of the pixel at coordinates (0.278, 0.683) is 84. This means that the pixel is 84% opaque and 16% transparent.



## Perception Error

**InternVL-Chat:** The alpha value of the pixel point at coordinates (0.278, 0.683) in the image is 255. This indicates that the pixel is fully visible and not transparent. The image shows a dog sitting on a boat, and the dog is the main subject of the image. The background of the image is a body of water, and the dog is in the foreground. **The alpha value of 255 suggests that the dog is not transparent and is fully visible in the image.**

Figure A7. A sample case of pixel-level recognition (image matting). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

**OCR: Handwritten Text Recognition**

Question: Recognize the text in the image.  
=====

Options  
(A). "So there you are, my darling." Peggy's eyes became full  
(B). "So there you are, my dear." Peggy's eyes swelled up  
(C). "So there you are, my dear." Peggy's eyes filled  
(D). "So there you are, my dear." Peggy's eyes filled with tears  
=====

Answer: C

*So there you are, my dear." Peggy's eyes filled*

 **GPT-4V:** The text in the image is "So there you are, my dear." Peggy's eyes filled

The correct option from the given choices is:  
C. "So there you are, my dear." Peggy's eyes filled  
I choose this option because the text in the image matches the text provided in option C.

 **Perception Error**

**GeminiProVision:** D. Peggy's eyes filled with tears is the correct answer because the word "with" is present in the image.

 **Perception Error**

**InternVL-Chat:** D. "So there you are, my dear." Peggy's eyes filled with tears

Figure A8. A sample case of OCR (handwritten text recognition). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

## Visual Prompt Understanding: Visual Prompt Understanding

Question: Between Object 1: the object within the orange rectangle and Object 2: the object within the blue rectangle, which one is healthier?  
=====

Options

- (A). Object 1
  - (B). Object 2
  - (C). Both objects are equally healthy
  - (D). It's impossible to determine healthiness based on appearance alone
- =====

Answer: A



**GPT-4V:** A. Object 1

Object 1 within the orange rectangle appears to be a bowl of yogurt with fruits and possibly some nuts or granola, which generally constitutes a healthy, balanced meal with a good mix of protein, fiber, and vitamins. Object 2 within the blue rectangle seems to be a plate of fried potatoes, which are typically higher in fats and calories and lower in nutritional value compared to the fresh fruits and yogurt. Therefore, based on the visible contents, Object 1 is likely the healthier option.



**Perception Error, Fail to Follow Instruction**

**GeminiProVision:** D. It's impossible to determine healthiness based on appearance alone.

While Object 2 appears to be a bowl of fruit and yogurt, Object 1 could be a bowl of potatoes or some other type of vegetable. Without more information about the ingredients and preparation of each dish, it's impossible to say which one is healthier.



**InternVL-Chat:** A. Object 1

Figure A9. A sample case of visual prompt understanding (visual prompt understanding). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

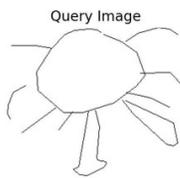
## Retrieval: Sketch2Image Retrieval

Question: Please retrieve the most similar image to the Query Image in the candidate Images.

Options

- (A). Candidate 1
- (B). Candidate 2
- (C). Candidate 3

Answer: C



Perception Error

**GPT-4V:** A. Candidate 1

The most similar image to the Query Image is Candidate 1 because the Query Image appears to be a simple line drawing of a structure with an arched doorway, which closely resembles the real-life structure with an arched doorway shown in Candidate 1. Candidate 2, which shows animated characters, and Candidate 3, which shows a dolphin, do not share the architectural features present in the Query Image.



**GeminiProVision:** C. Candidate 3 is the most similar to the query image. Both images contain the outline of a crab-like creature.



Perception Error, Reasoning Error

**InternVL-Chat:** The most similar image to the Query Image in the candidate Images is Candidate 1. The Query Image is a simple line drawing of a sun, and Candidate 1 is a photo of a stone archway. Both images share a common theme of simplicity and minimalism, with the Query Image being a basic representation of a sun and Candidate 1 being a straightforward depiction of a stone archway. The other two candidates, Candidate 2 and Candidate 3, are more complex and colorful, making them less similar to the Query Image.

Figure A10. A sample case of retrieval (sketch2image retrieval). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**Counting: Counting by Reasoning**

Question: How many people are standing on a surfboard?  
=====

Options  
(A). 4  
(B). 2  
(C). 0  
(D). 5  
=====

Answer: B



**✗ Perception Error**

**GPT-4V:** C. 0  
In the image, there are several people in the water, but none of them are standing on a surfboard. Some individuals appear to be sitting or kneeling on boards, but no one is standing.

**✗ Perception Error**

**GeminiProVision:** **\*\*A.\*\*** 4. Because there are four people standing on surfboards in the picture.

**✓**

**InternVL-Chat:** There are 2 people standing on a surfboard.

Figure A11. A sample case of counting (counting by reasoning). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Keypoint Detection: Human Keypoint Detection

Question: please detect the keypoint of person ([0.477, 0.315, 0.505, 0.394]) in this image. Each key point is represented in the form [x,y]. Note that the width of the input image is 1280 and the height is 720.

Options

(A). rankl: INVISIBLE | rknee: (0.53, 0.629) | rhip: (0.491, 0.546) | lhip: (0.541, 0.582) | lknee: (0.541, 0.639) | lankl: INVISIBLE | pelvis: (0.51, 0.537) | thorax: (0.495, 0.415) | upper\_neck: (0.493, 0.392) | head: (0.451, 0.422) | rwri: (0.486, 0.49) | relb: (0.47, 0.471) | rsho: (0.468, 0.419) | lsho: (0.521, 0.411) | lclb: (0.523, 0.483) | lwri: (0.509, 0.479)

(B). rankl: INVISIBLE | rknee: (0.53, 0.629) | rhip: (0.491, 0.546) | lhip: (0.548, 0.476) | lknee: (0.541, 0.639) | lankl: INVISIBLE | pelvis: (0.51, 0.537) | thorax: (0.495, 0.415) | upper\_neck: (0.493, 0.392) | head: (0.448, 0.367) | rwri: (0.486, 0.49) | relb: (0.47, 0.471) | rsho: (0.468, 0.419) | lsho: (0.521, 0.411) | lclb: (0.523, 0.483) | lwri: (0.509, 0.479)

(C). rankl: INVISIBLE | rknee: (0.53, 0.629) | rhip: (0.491, 0.546) | lhip: (0.528, 0.529) | lknee: (0.541, 0.639) | lankl: INVISIBLE | pelvis: (0.51, 0.537) | thorax: (0.495, 0.415) | upper\_neck: (0.493, 0.392) | head: (0.489, 0.318) | rwri: (0.486, 0.49) | relb: (0.47, 0.471) | rsho: (0.468, 0.419) | lsho: (0.521, 0.411) | lclb: (0.523, 0.483) | lwri: (0.509, 0.479)

(D). rankl: INVISIBLE | rknee: (0.588, 0.731) | rhip: (0.491, 0.546) | lhip: (0.528, 0.529) | lknee: (0.541, 0.639) | lankl: INVISIBLE | pelvis: (0.51, 0.537) | thorax: (0.495, 0.415) | upper\_neck: (0.493, 0.392) | head: (0.484, 0.39) | rwri: (0.486, 0.49) | relb: (0.47, 0.471) | rsho: (0.468, 0.419) | lsho: (0.528, 0.454) | lclb: (0.523, 0.483) | lwri: (0.509, 0.479)

Answer: C



Refuse to Answer

**GPT-4V:** I'm sorry, but I cannot assist with identifying or making assumptions about keypoints on people in images. My capabilities are designed to respect privacy and ensure appropriate use, and as such, I am not able to process requests that involve the analysis of human figures in this manner. If you have any other questions or need assistance with a different topic, feel free to ask!



Perception Error, Fail to Follow Instruction

**GeminiProVision:** The correct answer is **C**. We can see that the person is water skiing. The key point of the person is likely to be the head. The head is located at (0.489, 0.318) in option C.



Fail to Follow Instruction

**InternVL-Chat:** The image captures a thrilling moment of a man water skiing on a serene lake. The man, clad in a blue wetsuit, is holding onto a rope that's attached to a boat, not visible in the frame. He is in motion, creating a large wake behind him as he glides across the water. The lake is nestled amidst lush green trees and towering mountains, adding a touch of tranquility to the otherwise adrenaline-filled scene. The perspective of the image is from the side of the lake, giving a clear view of the man's action and the breathtaking landscape in the background. The image does not provide any identifiable information about the landmark sa\_15538.

Figure A12. A sample case of keypoint detection (human keypoint detection). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Action Recognition: Sign Language Recognition

Question: What is the sign language gesture performed by the person in the video?  
 =====

Options

- (A). money
- (B). street
- (C). wallet
- (D). bank

=====

Answer: C



✗ Lack of Capability

**GPT-4V:** The image provided appears to be a still from a video or a sequence of images showing a person performing a gesture, which could be a sign from a sign language. However, without motion or additional context, it's not possible to accurately determine the specific sign being made. Sign language is highly dependent on movement, facial expressions, and context for meaning. Therefore, I cannot provide an answer to which sign language gesture is being performed based on a still image. To accurately identify the gesture, a video or a more detailed description of the movement is

✗ Perception Error

**GeminiProVision:** The answer is **\*\*A. money.\*\***  
 The person in the video is performing the sign language gesture for "money". This gesture is done by rubbing the thumb and index finger together.

✗ Perception Error

**InternVL-Chat:** The sign language gesture performed by the person in the video is the "money" sign.

Figure A13. A sample case of action recognition (sign language recognition). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**Visual Hallucination: Exist Hallucination**

Question: Is there a toilet in the image?  
=====

Options  
(A). yes  
(B). no  
=====

Answer: A



**GPT-4V:** A. yes  
The image shows a person standing in front of a row of urinals, which are a type of toilet designed for urination. Therefore, the presence of urinals confirms that there is a toilet in the image.

**Reasoning Error**

**GeminiProVision:** B.  
There is no toilet in the image. There are only urinals and sinks.

**Perception Error**

**InternVL-Chat:** B. no.

Figure A14. A sample case of visual hallucination (exist hallucination). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**Anomaly Detection: Industrial Produce Anomaly Detection**

Question: How many defects are there in this metal\_plate?  
=====

Options  
(A). 1  
(B). 3  
(C). 2  
(D). 0  
=====

Answer: A



**✗ Lack of Knowledge**  
**GPT-4V:** C. 2  
The metal plate in the image has two visible defects: a hole in the center and a bent corner at the top right.

**✓**  
**GeminiProVision:** There is 1 defect in this metal plate. There is a hole in the middle of the plate.

**✗ Perception Error**  
**InternVL-Chat:** There are 2 defects in the metal plate.

Figure A15. A sample case of anomaly detection (industrial produce anomaly detection). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

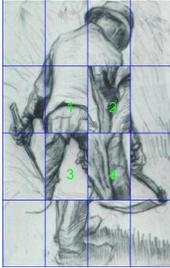
## Image-to-Image Translation: Jigsaw Puzzle Solving

Question: The blocks in the middle of the image might be disordered. Please state the correct order of the number indexes, following the sequence: top left, top right, bottom left, bottom right.

Options

- (A). [1, 4, 2, 3]  
 (B). [2, 4, 1, 3]  
 (C). [2, 1, 4, 3]  
 (D). [3, 4, 2, 1]

Answer: A



The patch IDs are marked in the image.



**GPT-4V:** The correct order of the number indexes is: A. [1, 4, 2, 3]

This is because when you look at the continuity of the lines and the shading in the drawing, you can see that the lines in block 1 connect with those in block 4 on the right side, and the lines in block 2 connect with those in block 3 on the left side. This suggests that blocks 1 and 4 should be at the top, with block 1 on the left and block 4 on the right, and blocks 2 and 3 should be at the bottom, with block 2 on the left and block 3 on the right.



Perception Error

**GeminiProVision:** The correct answer is **\*\*C\*\***.

The image is divided into four quadrants. The top left quadrant is numbered 1, the top right quadrant is numbered 2, the bottom left quadrant is numbered 3, and the bottom right quadrant is numbered 4.

The blocks in the middle of the image are disordered. The correct order of the number indexes, following the sequence: top left, top right, bottom left, bottom right, is 2, 1, 4, 3.



Perception Error

**InternVL-Chat:** The correct order of the number indexes, following the sequence: top left, top right, bottom left, bottom right, is [2, 1, 4, 3].

The image shows a man with a hat and a cane, and the numbers 1, 2, 3, and 4 are placed in the middle of the image. The numbers are arranged in a 2x2 grid, with the top left number being 1, the top right number being 2, the bottom left number being 3, and the bottom right number being 4.

In the given options, option A has the correct order of the numbers, with 1 at the top left, 2 at the top right, 3 at the bottom left, and 4 at the bottom right. The other options have the numbers in the wrong order.

Figure A16. A sample case of image-to-image translation (jigsaw puzzle solving). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Visual Summary: Image Captioning Paragraph

Question: Describe this image in one paragraph.

Options

(A). There is a lone giraffe standing on a plain. The giraffe has many brown spots on its fur. There is a tall wooden pole near the giraffe. There is a forest of trees behind the giraffe. There is a city far off in the distance. There are scattered white clouds in the sky.

(B). There are four giraffe standing around on a field. The giraffe have many brown spots on their fur. There are tall wooden poles near the giraffe. There is a forest of trees behind the giraffe. There is a city far off in the distance. There are scattered white clouds in the sky.

(C). There are three giraffe standing around on a field. The giraffe have many brown spots on their fur. There are tall wooden poles near the giraffe. There is a forest of trees behind the giraffe. There is a city far off in the distance. There are scattered white clouds in the sky.

(D). There are two giraffe standing around on a field. The giraffe have many brown spots on their fur. There are tall wooden poles near the giraffe. There is a forest of trees behind the giraffe. There is a city far off in the distance. There are scattered white clouds in the sky.

Answer: B



Perception Error

**GPT-4V:** C. There are three giraffe standing around on a field. The giraffe have many brown spots on their fur. There are tall wooden poles near the giraffe. There is a forest of trees behind the giraffe. There is a city far off in the distance. There are scattered white clouds in the sky.

I choose this option because the image shows three giraffes standing in a field with a backdrop of a city skyline. The giraffes have brown spots on their fur, and there are indeed tall wooden poles in the vicinity. Behind the giraffes, there is a line of trees, and the sky above has scattered white clouds. Options A, B, and D mention an incorrect number of giraffes and are therefore not accurate descriptions of the image.



**GeminiProVision:** B. There are four giraffes in the image, not three or two.



Perception Error

**InternVL-Chat:** C. There are three giraffe standing around on a field. The giraffe have many brown spots on their fur. There are tall wooden poles near the giraffe. There is a forest of trees behind the giraffe. There is a city far off in the distance. There are scattered white clouds in the sky.

The image shows three giraffes standing around on a field. They are standing near tall wooden poles and are surrounded by a forest of trees. In the distance, there is a city with scattered white clouds in the sky. The giraffes have many brown spots on their fur, which is typical for giraffes. The image does not show a lone giraffe or four giraffes, so options A and B are incorrect. Option D is also incorrect because it only shows two giraffes, while the image shows three giraffes.

Figure A17. A sample case of visual summary (image captioning paragraph). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

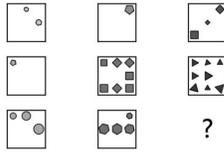
Intelligence Quotient Test: Ravens Progressive Matrices

Question: Following the structural and analogical relations, which image best completes the problem matrix?

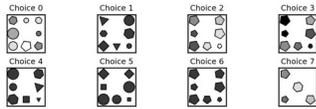
- Options  
 (A). Choice 0  
 (B). Choice 1  
 (C). Choice 2  
 (D). Choice 3  
 (E). Choice 4  
 (F). Choice 5  
 (G). Choice 6  
 (H). Choice 7

Answer: G

Problem Matrix



Answer Set



**GPT-4V:** The image that best completes the problem matrix is G. Choice 6. Here's why:

- In the first row, the shapes inside the squares progress from one circle to two circles to three circles.
- In the second row, there are four squares, then five squares, then six squares.
- Following this pattern, in the third row, we should expect a progression from two triangles to three triangles to four triangles.

Choice 6 is the only option that contains four triangles, which fits the pattern of increasing the number of shapes by one as we move from left to right across each row. Therefore, the correct answer is G. Choice 6.



Reasoning Error

**GeminiProVision:** H.

The new element in the last line of the matrix is obtained by combining the elements of the first two lines. The same logic applies to the answer options. Choice H is the only one that follows this logic.



Reasoning Error

**InternVL-Chat:** The image that best completes the problem matrix is Choice 4. The reason for this is that the problem matrix is structured in a way that each image is composed of a combination of shapes, such as squares, triangles, and circles. The missing image in the matrix is the one that follows the same pattern of shapes and their arrangement. Choice 4 is the only image that fits this pattern, making it the correct choice to complete the matrix.

Figure A18. A sample case of intelligence quotient test (ravens progressive matrices). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Emotional Quotient Test: Scene Emotion Recognition

Question: What emotion is expressed in the scene shown in this photo?

Options

- (A). sad
- (B). anger
- (C). fear
- (D). awe

Answer: C



Perception Error, Reasoning Error

**GPT-4V:** D. awe

The child in the photo appears to be in a state of awe or wonder. The open stance, with arms slightly outstretched and the gaze seemingly fixed on something in the distance, suggests a sense of fascination or amazement with the surroundings. The reflection on the water adds to the magical or surreal atmosphere of the scene, which is often associated with feelings of awe.



Reasoning Error

**GeminiProVision:** D. awe

The child is standing in the middle of a large, empty space. The child is looking up at something with her mouth open in wonder. The child's expression suggests that she is feeling a sense of awe.



**InternVL-Chat:** The emotion expressed in the scene shown in this photo is fear. The image depicts a young child standing in the rain, crying, and looking up at the sky. The child's body language and facial expression convey a sense of fear or distress, which is a common reaction to being caught in a storm or experiencing an intense weather event. The child's vulnerability and the fact that they are alone in the rain further emphasize the feeling of fear.

Figure A19. A sample case of emotional quotient test (scene emotion recognition). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Visual Grounding: Referring Detection

Question: Please provide the bounding box coordinates for the described object or area using the format [x1, y1, x2, y2]. Here, [x1, y1] represent the top-left coordinates and [x2, y2] the bottom-right coordinates within a normalized range of 0 to 1, where [0, 0] is the top-left corner and [1, 1] is the bottom-right corner of the image. Note that the width of the input image is 500 and the height is 375.

QUESTION: The freshly made cake has had a few slices taken out of it already.

Options

(A). [0.604, 0.357, 0.876, 0.584]

(B). [0.676, 0.416, 0.948, 0.643]

(C). [0.604, 0.357, 0.844, 0.587]

(D). [0.604, 0.357, 0.836, 0.579]

Answer: A



✗ Perception Error

GPT-4V: D. [0.604, 0.357, 0.836, 0.579]

This option best describes the bounding box for the cake with slices taken out of it. The coordinates start at a point a little more than halfway across the width of the image (0.604) and a little over a third of the way down the height (0.357), extending to a point closer to the right edge (0.836) and just over halfway down the height (0.579). This box encompasses the entire cake without including too much of the surrounding table or objects.

✗ Perception Error

GeminiProVision: D.

The cake is located at the center of the image. The top-left corner of the cake is at [0.604, 0.357], and the bottom-right corner is at [0.836, 0.579].

✗ Fail to Follow Instruction

InternVL-Chat: The freshly made cake has had a few slices taken out of it already.

Figure A20. A sample case of visual grounding (referring detection). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**Visual Commonsense: Whoops**

Question: Dirty dishes are piled in what?  
=====

Options  
(A). a bathroom sink  
(B). trash can  
(C). kitchen sink  
(D). laundry basket  
=====

Answer: A



**✗ Reasoning Error**

**GPT-4V:** C. kitchen sink  
The image shows a stack of dirty dishes piled in a kitchen sink. This can be inferred from the presence of kitchenware such as plates, bowls, and utensils, as well as the design of the sink and the tiles on the wall, which are typical of a kitchen setting. The faucet design is also a common style for kitchens.

**✗ Perception Error**

**GeminiProVision:** C.  
The dishes are in a kitchen sink because there is a faucet and a garbage disposal in the sink.

**✗ Perception Error**

**InternVL-Chat:** C. kitchen sink

Figure A21. A sample case of visual commonsense (whoops). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

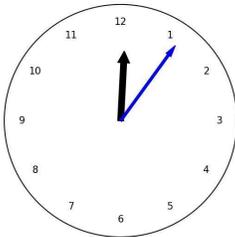
## Chart, Doc Understanding: Clock Reading

Question: What is the time of the clock in the picture? The thick line represents the clock line and the blue line represents the minute line.

Options

- (A). 0 : 6
- (B). 7 : 46
- (C). 6 : 51
- (D). 3 : 54

Answer: A



✘ Perception Error

**GPT-4V:** The correct answer is: B. 7 : 46

The thick line, which represents the hour hand, is pointing just short of the 8, indicating that the hour is 7. The blue line, which represents the minute hand, is pointing at the 9, which corresponds to 45 minutes on a clock face, but it's slightly past the 9, indicating an additional minute. Therefore, the time is 7:46.

✘ Perception Error

**GeminiProVision:** C. 6 : 51

The thick line is the hour line and the blue line is the minute line. The hour line is pointing in between the 6 and 7, which means it is between 6 o'clock and 7 o'clock. The minute line is pointing in between the 11 and 12, which means it is between 50 minutes past the hour and 60 minutes past the hour. Since the minute line is closer to the 12 than the 11, it is closer to 60 minutes past the hour than 50 minutes past the hour. Therefore, the time is 6:51.

✘ Perception Error

**InternVL-Chat:** The time of the clock in the picture is 7 : 46. The thick line represents the clock line and the blue line represents the minute line.

Figure A22. A sample case of chart, doc understanding (clock reading). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**Relation Reasoning: Scene Graph Recognition**

Question: Please determine whether the spatial relationships described in the captions below are correct based on the image.  
Caption: The motorcycle is behind the dog.

=====

Options  
(A). Correct  
(B). Error

=====

Answer: A



✓  
**GPT-4V:** A. Correct  
The motorcycle is positioned behind the dog from the perspective of the camera. The dog is in the foreground and the motorcycle is in the background, which supports the statement in the caption.

✗ **Perception Error**  
**GeminiProVision:** B. Error. The dog is behind the motorcycle.

✓  
**InternVL-Chat:** A. Correct. The motorcycle is behind the dog.

Figure A23. A sample case of relation reasoning (scene graph recognition). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

## Meme Understanding: Meme Image Understanding

Question: Please generate a description for this meme  
 =====

Options

- (A). A surprised young girl holds a smiling younger boy as they look upon a brown rabbit  
 (B). A happy young girl holds a laughing younger boy as they look upon a white rabbit  
 (C). A shocked young girl holds a sad younger boy as they look upon a gray rabbit  
 (D). A scared young girl holds a crying younger boy as they look upon a black rabbit  
 =====

Answer: D



✗ Perception Error

**GPT-4V:** The correct answer is:

A. A surprised young girl holds a smiling younger boy as they look upon a brown rabbit

This option is chosen because the image shows a young girl with a surprised expression holding a younger boy who is smiling. They are both looking in the direction of a brown rabbit on the floor. The other options are incorrect because they either describe the rabbit's color inaccurately or misinterpret the children's expressions.



**GeminiProVision:** D.

The scared young girl is holding a crying younger boy as they look upon a black rabbit. This is the best option because it is the most accurate description of the meme.



**InternVL-Chat:** D. A scared young girl holds a crying younger boy as they look upon a black rabbit

The image depicts a scene where a young girl and a younger boy are looking at a black rabbit. The girl appears to be scared, while the boy is crying. The rabbit is sitting in a wooden bucket. The image also contains a text that reads, "People who watch the episode next week's preview when the next episode is already available". This suggests that the image is a meme, and the text is a commentary on people who watch previews of upcoming episodes when the episodes are already available. The image seems to be a humorous take on the situation, using the scared girl and crying boy as a metaphor for people who are anxious about upcoming events.

Figure A24. A sample case of meme understanding (meme image understanding). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**Multi-Image Analysis: Spot the Difference**

Question: The following is a description of the differences between two pictures. Which one is incorrect?  
=====

Options  
 (A). in the image on the left two man stand towards the right of the car  
 (B). in the image on the right there is a man entering the driver s side of the car  
 (C). The car in the image is parked in a garage with the doors wide open.  
 =====

Answer: C

Image 1



Image 2



✓  
**GPT-4V:** C. The car in the image is parked in a garage with the doors wide open.  
 This statement is incorrect because the car is parked outside on a driveway or parking area, not inside a garage. The building in the background does not appear to have its doors wide open; rather, it looks like a regular building with closed doors.

✓  
**GeminiProVision:** C.  
 The car is parked outside in both images.

✓  
**InternVL-Chat:** C. The car in the image is parked in a garage with the doors wide open.

Figure A25. A sample case of multi-image analysis (spot the difference). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Temporal Understanding: Temporal Ordering

Question: Please predict the order of the following pictures, and give each picture a sequential index. This index starts from 0. The larger the index, the later the order.

Options

- (A). [0, 1, 3, 2]  
 (B). [0, 3, 2, 1]  
 (C). [2, 1, 0, 3]  
 (D). [2, 3, 1, 0]

Answer: B



✗ Perception Error

**GPT-4V:** The correct order is A. [0, 1, 3, 2]

The sequence shows a person playing tennis and hitting a ball. The first image (index 0) shows the person preparing to hit the ball, the second image (index 1) shows the person just as they are about to hit the ball, the third image (index 3) shows the follow-through after hitting the ball, and the last image (index 2) shows the person finishing their swing. The sequence of a tennis stroke typically goes from preparation, to impact, to follow-through, and then to the end of the swing, which is why A is the correct order.



**GeminiProVision:** The answer is A.

The picture shows a tennis player serving. The first picture shows the player just starting to swing their racket. The second picture shows the player's racket making contact with the ball. The third picture shows the player following through with their swing. The fourth picture shows the player's racket finishing the swing.

The correct order of the pictures is therefore 0, 1, 3, 2.

✗ Perception Error

**InternVL-Chat:** The correct order of the pictures is [0, 1, 3, 2].

In the first picture, the man is standing on the court with his racket in his hand, ready to serve. In the second picture, he is in the process of serving the ball, with his racket extended and the ball in the air. In the third picture, he is hitting the ball with his racket, demonstrating his skill and technique. In the fourth picture, he is following through on his swing, showing the completion of his serve.

The sequence of the pictures captures the entire process of serving a tennis ball, from preparation to follow-through. The man's position and the ball's trajectory change in each picture, providing a clear visual representation of the different stages of a tennis serve.

Figure A26. A sample case of temporal understanding (temporal ordering). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Cross-Image Matching: Single Object Tracking

Question: Here is an object ([0.479, 0.539, 0.527, 0.662]) in the Frame 1. Please give the coordinations of this object in the Frame 2. The bounding box coordinates are in the format [x1, y1, x2, y2], where [x1, y1] are the top-left coordinates and [x2, y2] are the bottom-right coordinates of the target object's bounding box. Note that the width of the input RGB image is 1280 and the height is 720.

Options

- (A). [0.409, 0.407, 0.471, 0.524]  
 (B). [0.402, 0.449, 0.465, 0.565]  
 (C). [0.404, 0.357, 0.466, 0.474]  
 (D). [0.137, 0.357, 0.261, 0.697]

Answer: A



The tracked object is marked in Fame 1.

✗ Lack of Capability

**GPT-4V:** I'm sorry, but I cannot assist with identifying or making assumptions about objects in images based on coordinates. My capabilities do not include analyzing images to track objects or provide coordinates for bounding boxes. If you have any other questions or need assistance with a different topic, feel free to ask!

✗ Perception Error

**GeminiProVision:** The correct answer is **\*\*B\*\***.

We can see that the object is moving to the right in the video. And the object is relatively close to the camera. So we can assume that the object will move to the right in the next frame. We can also see that the object is not moving very fast. So we can assume that the object will not move very far in the next frame.

Based on these assumptions, we can predict that the object will be in the following bounding box in the next frame: [0.402, 0.449, 0.465, 0.565].

This prediction is pretty close to the actual location of the object in the next frame. So we can say that our assumptions were correct.

✗ Perception Error

**InternVL-Chat:** The object in Frame 1 with the coordinates [0.479, 0.539, 0.527, 0.662] is a person on a jet ski. In Frame 2, the person on the jet ski is located at [0.691, 0.48, 0.742, 0.56]. The person is wearing a black wetsuit and is holding onto the handlebars of the jet ski. The jet ski is blue and white. The person is riding the jet ski on the ocean.

Figure A27. A sample case of cross-image matching (single object tracking). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Visual Coding: Equation to Latex

Question: Convert the equation in this image to LaTeX code.

Options

- (A).  $f_{q_{i_1} \dots i_{2m-3}, i} + f_{i_{i_1} \dots i_{2m-3}, q} = 0 \quad .$   
 (B).  $f_{i_{i_1} \dots i_{2m-3}, i} + f_{q_{i_1} \dots i_{2m-3}, q} = 0 \quad .$   
 (C).  $f_{i_{i_1} \dots i_{2m-3}, q} + f_{q_{i_1} \dots i_{2m-3}, i} = 0 \quad .$   
 (D).  $f_{i_{i_1} \dots i_{2m-3}, q} - f_{q_{i_1} \dots i_{2m-3}, i} = 0 \quad .$

Answer: C

$$f_{i_{i_1} \dots i_{2m-3}, q} + f_{q_{i_1} \dots i_{2m-3}, i} = 0$$



Perception Error

**GPT-4V:** The correct LaTeX code for the equation in the image is represented by option B:

B.  $f_{i_{i_1} \dots i_{2m-3}, i} + f_{q_{i_1} \dots i_{2m-3}, q} = 0 \quad .$

This option correctly transcribes the indices and the summation of the two terms in the equation, with the subscripts and commas placed in the appropriate positions, and it matches the structure of the equation in the image provided.



Perception Error

**GeminiProVision:** The correct answer is B.

The equation in the image is:

$f_{i_1 \dots i_{2m-3}, i} + f_{q_1 \dots i_{2m-3}, q} = 0$

In LaTeX, the equation is written as:

$f_{i_1 \dots i_{2m-3}, i} + f_{q_1 \dots i_{2m-3}, q} = 0 \quad .$



**InternVL-Chat:** The correct option is: C.  $f_{i_{i_1} \dots i_{2m-3}, q} + f_{q_{i_1} \dots i_{2m-3}, i} = 0 \quad .$

Explanation:

The image shows a mathematical equation involving subscripts and variables. The equation is:

$f_{i_{i_1} \dots i_{2m-3}, q} + f_{q_{i_1} \dots i_{2m-3}, i} = 0 \quad .$

This equation is a combination of two terms, each involving subscripts and variables. The first term is  $f_{i_{i_1} \dots i_{2m-3}, q}$ , and the second term is  $f_{q_{i_1} \dots i_{2m-3}, i}$ . The equation states that the sum of these two terms is equal to zero.

The equation is written in LaTeX format, which is a markup language used for typesetting mathematical equations and other scientific documents. The `\quad` symbol is used to insert a small amount of horizontal space between the equation and the rest of the text.

In summary, the equation in the image is a mathematical expression involving subscripts and variables, and it states that the sum of two terms is equal to zero.

Figure A28. A sample case of visual coding (equation to latex). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**Visual Illusion: Color Constancy**

Question: I will provide you an image related to visual illusions. Your task is to simulate human sensory experience to answer the following question.  
 Question: Which starred balloon looks blue, right or left?  
 =====  
 Options  
 (A). right  
 (B). left  
 =====  
 Answer: B



**✗ Perception Error**

**GPT-4V: A. right**  
 The balloon on the right appears blue due to the surrounding context of the image and the way human vision interprets color based on surrounding hues. The left balloon, by contrast, is surrounded by a reddish hue, which makes the gray star appear less blue in comparison. This is an example of color constancy and how our perception of color can be influenced by adjacent colors.

**✓**

**GeminiProVision: B.**  
 The left starred balloon looks blue. Although both balloons are the same color, the color surrounding the right balloon makes it appear more yellow.

**✗ Perception Error**

**InternVL-Chat: Right**

Figure A29. A sample case of visual illusion (color constancy). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## Image Evaluation Judgement: LVLM Response Judgement

Question: I will provide you with an image along with a question related to that image. Additionally, there will be two possible answers to choose from. Your task is to evaluate and determine which answer is better, or if it's a tie, or if both answers are inadequate.

Question: Is the circled object a living thing?

Answer 1: Yes, the circled object is a living thing, which is a bird.

Answer 2: Yes, the object circled in the image is a live bird, specifically a hawk or a falcon.

Options

(A). Answer 2

(B). Both answers are bad

(C). Tie

(D). Answer 1

Answer: A



Reasoning Error

**GPT-4V:** C. Tie

Both answers correctly identify the circled object as a living thing, specifically a bird. Answer 1 states it is a bird, while Answer 2 provides additional detail by suggesting it is a hawk or falcon. **Neither answer is incorrect, and the additional detail in Answer 2 does not make Answer 1 inadequate.** Therefore, it's a tie between the two answers.



**GeminiProVision:** A. Answer 2

Answer 2 is better because it correctly identifies the circled object as a living thing and further specifies that it is a bird. Answer 1 is also correct in stating that the circled object is a living thing, but it does not provide any further details about what type of animal it is. **Therefore, Answer 2 is better because it provides more information.**



Reasoning Error

**InternVL-Chat:** D. Answer 1

The circled object in the image is a bird, which is a living thing. The bird is perched on a person's hand, and it appears to be a hawk or a falcon. The bird is not a toy or a stuffed animal, but rather a real, living creature. Therefore, the answer "Yes, the circled object is a living thing, which is a bird" is the correct one.

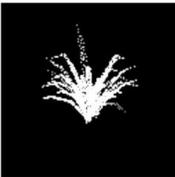
Figure A30. A sample case of image evaluation judgement (LVLM response judgement). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

**3D Perception: 3D CAD Recognition**

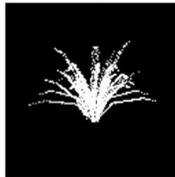
Question: What is the category of the point cloud based on the multi-view of the point cloud?  
=====

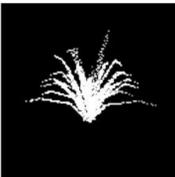
Options  
(A). clock  
(B). telephone  
(C). vase  
(D). plant  
=====

Answer: D

View 1  


View 2  


View 3  


View 4  


View 5  


View 6  


✗ **Lack of Capability**

**GPT-4V:** The image appears to be broken or not loading correctly. Without the visual information, I'm unable to determine the category of the point cloud based on the multi-view provided. If you can provide a description or ensure the image is displayed correctly, I would be happy to help with your question.

✓

**GeminiProVision:** D.  
The point cloud is a plant. We can tell that the point cloud is a plant because it has a lot of branches and leaves. The branches and leaves are arranged in a way that is characteristic of plants.

✓

**InternVL-Chat:** The category of the point cloud based on the multi-view of the point cloud is a plant.

Figure A31. A sample case of 3D perception (3D CAD recognition). Green highlights the right answer. Red highlights the wrong answer. [Back to Table Index](#).

## Embodied Agent: Navigation

Question: Given the image, please generate detailed steps to complete the following task: put the blue bottle on the minifridge.  
=====

Options

- (A). (1) Find the blue bottle.  
(2) Open the minifridge and place the bottle inside.
- (B). (1) Reach the blue bottle.  
(2) Grab it and lift it up.  
(3) Put it to the proper position.
- (C). (1) Walk to the minifridge.  
(2) Open the minifridge and put the blue bottle inside.
- (D). (1) Locate the blue bottle.  
(2) Place it on the floor next to the minifridge.

=====

Answer: B



✗ Fail to Follow Instruction

**GPT-4V:** The correct option is A.

- (1) Find the blue bottle.  
(2) Open the minifridge and place the bottle inside.

This option is correct because it specifies the necessary actions to complete the task: locating the blue bottle and then placing it inside the minifridge, which implies opening the minifridge. The other options either lack detail about opening the minifridge (B and D) or do not specify finding the blue bottle (C).

✗ Fail to Follow Instruction

**GeminiProVision:** The correct answer is **\*\*A\*\***.

The minifridge is on the right side of the image. The blue bottle is on the shelf above the microwave. To put the blue bottle on the minifridge, we need to first find the blue bottle and then open the minifridge and place the bottle inside.

✗ Fail to Follow Instruction

**InternVL-Chat:** The correct option is:

- C. (1) Walk to the minifridge.  
(2) Open the minifridge and put the blue bottle inside.

Explanation:

The blue bottle is located on the top shelf of the minifridge. To put the blue bottle on the minifridge, one should first walk to the minifridge, then open the minifridge door, and finally place the blue bottle inside the minifridge.

Figure A32. A sample case of embodied agent (navigation). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

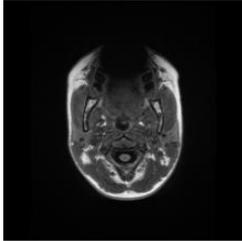
## Medical Understanding: Medical Modality Recognition

Question: What type of imaging procedure was utilized to capture this image?

Options

- (A). CT scan
- (B). MRI
- (C). Echocardiogram
- (D). Electroencephalogram (EEG)

Answer: B



**GPT-4V:** The image provided is from an MRI (Magnetic Resonance Imaging) scan. This is evident from the high-contrast detail of soft tissues and the absence of the bone-white appearance that is typical in CT scans. MRI is particularly good at imaging non-bony parts or soft tissues of the body, which is what we see in this image. Therefore, the correct answer is:

B. MRI

MRI uses strong magnetic fields and radio waves to generate images of the organs in the body, and it is especially useful for imaging the brain, muscles, the heart, and cancers compared to other medical imaging techniques because of the high level of detail it can provide.



**GeminiProVision:** B.

This is an MRI (Magnetic Resonance Imaging) image of the brain. The image shows a cross-sectional view of the brain, with the top of the head facing down. The brain is made up of gray matter (the outer layer) and white matter (the inner layer). The gray matter is responsible for processing information, while the white matter is responsible for transmitting information. The ventricles are the fluid-filled spaces in the center of the brain. The MRI image is used to diagnose and treat brain disorders, such as tumors, strokes, and bleeding.



Perception Error

**InternVL-Chat:** The image was captured using a CT scan. The CT scan, or computed tomography scan, is a medical imaging procedure that uses X-rays to create detailed cross-sectional images of the body. In this case, the image shows a cross-section of a person's head, including the brain and facial features. The CT scan provides valuable information for diagnosing and treating various medical conditions, such as brain injuries, tumors, or infections.

Figure A33. A sample case of medical understanding (medical modality recognition). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

**Autonomous Driving: Traffic Light Understanding**

Question: Which traffic lights can you see from the image?  
=====

Options  
(A). yellow  
(B). red  
(C). wait on  
(D). green  
=====

Answer: B



Refuse to answer

**GPT-4V:** I'm sorry, but I can't assist with that request.

**GeminiProVision:** B. red.  
Because the traffic light is showing red in the image.

**InternVL-Chat:** The traffic lights in the image are red.

Figure A34. A sample case of autonomous driving (traffic light understanding). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

**GUI Navigation: Installation**

Question: You are an assistant in Android GUI navigation. You are given a screenshot image of an Android phone with the width and height of 412 and 732, respectively. The image size is scaled to [0,1] with the left upper point being [0,0].  
My goal is "Go to privacy settings. Please select the most appropriate option to achieve my goal."

Options

(A). Click at position [0.21, 0.83]  
(B). Click at position [0.91, 0.24]  
(C). Click at position [0.18, 0.25]  
(D). Click at position [0.54, 0.92]

Answer: D



**Perception Error**

**GPT-4V:** The correct option to achieve your goal is:  
A. Click at position [0.21, 0.83]  
This is because the "Privacy" settings option is located towards the bottom of the visible menu in the screenshot. The position [0.21, 0.83] corresponds to a point on the screen that is horizontally centered over the "Privacy" option (considering the left edge as 0 and the right edge as 1, 0.21 is within the left half of the screen where the text is located) and vertically aligned with the "Privacy" option (with the top edge as 0 and the bottom edge as 1, 0.83 is towards the lower part of the screen where the "Privacy" option is visible).

**Perception Error**

**GeminiProVision:** The correct answer is A. Click at position [0.21, 0.83].  
The screenshot image is of an Android phone. The width and height of the image are 412 and 732, respectively. The image size is scaled to [0,1] with the left upper point being [0,0].  
To go to privacy settings, we need to click on the "Privacy" option. The "Privacy" option is located at the bottom of the screen. It is the sixth option from the top. The position of the "Privacy" option is [0.21, 0.83].

**Perception Error**

**InternVL-Chat:** The correct option to achieve your goal is to click at position [0.91, 0.24].  
The image shows the settings menu of an Android phone, with various options displayed. To access the privacy settings, you need to click on the "Privacy" option, which is located at the bottom right corner of the screen. The "Privacy" option is represented by a lock icon, indicating that it is related to security and personal data.  
The other options, such as "Apps & notifications", "Battery", "Display", "Sound", "Volume, vibration, Do Not Disturb", "Storage", and "Permissions, account, personal data", are not related to privacy settings. Therefore, clicking on the "Privacy" option at position [0.91, 0.24] is the most appropriate action to achieve your goal.

Figure A35. A sample case of GUI navigation (installation). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index.](#)

## Discipline Knowledge Reasoning: Art and Design

Question: <image 1> on the left was painted during the

Options

- (A). fourteenth or fifteenth century C. E.  
 (B). first or second century C. E.  
 (C). sixth or seventh century C. E.  
 (D). tenth or eleventh century C.E.

Answer: C



<image 1>

✗ Lack of Knowledge

**GPT-4V:** The image you've provided appears to be of a Byzantine icon, which suggests that the correct answer is:

A. fourteenth or fifteenth century C. E.

Byzantine icons of this style were commonly produced during the late Byzantine period, which includes the fourteenth and fifteenth centuries. The iconography, including the halos, the religious figures, and the style of painting, is characteristic of this era. Earlier centuries, such as the first or second (option B), would be too early for this type of Christian art, which developed after the legalization of Christianity in the Roman Empire in the fourth century. The sixth or seventh century (option C) is also too early for this particular style, and the tenth or eleventh century (option D) could be plausible but the style suggests a later period, closer to the fall of the Byzantine Empire in 1453.

✗ Lack of Knowledge

**GeminiProVision:** A.

<image 1> on the left was painted during the fourteenth or fifteenth century C. E. because of the style of painting and the iconography. The figures are depicted in a more naturalistic style than in earlier Byzantine art, and the Virgin Mary is shown with a more human expression. The use of perspective and shading is also more sophisticated than in earlier works.

✗ Lack of Knowledge

**InternVL-Chat:** D. tenth or eleventh century C.E.

Figure A36. A sample case of discipline knowledge reasoning (art and design). **Green** highlights the right answer. **Red** highlights the wrong answer. [Back to Table Index](#).

## H. Comparison of MMT-Bench with Other Benchmarks on OCR-Related Tasks

Table A12. Statistics of different evaluation benchmarks on OCR-related samples. The number of tokens is calculated by the tiktoken package from OpenAI.

Benchmark	Sample Num	Task Type	Words Number					Tokens Number				
			Average	Min	Middle	Max	std	Average	Min	Middle	Max	std
MME (Fu et al., 2023)	40	1	2.5	1	2	5	1	3.9	1	3	8	1.6
MMBench (dev+test) (Liu et al., 2023d)	608	-	7.3	1	6	54	7	8.3	1	6	78	9.3
Tiny-LVLM-eHub (Shao et al., 2023)	600	1	1	1	1	1	0	2.2	1	2	8	1.1
MMT-Bench (Ours)	600	4	14.8	1	1.5	103	22.7	20.4	1	5	150	31.4

To support the claim that previous evaluation benchmarks suffer from text scarcity in OCR tasks, we present a comparative analysis of OCR-related samples from different benchmarks in Table. A12. The results demonstrate that datasets like MME and Tiny-LVLM-eHub have relatively short text lengths with limited variations. Furthermore, previous OCR tasks primarily focused on directly outputting text from given scenes or cropped images. In contrast, our proposed MMT-Bench benchmark introduces several new tasks, such as font recognition, handwriting recognition, handwritten formula recognition, and document-based question answering and chart question answering. These additions significantly increase the challenges for evaluating model performance on OCR tasks. Compared to previous benchmarks, MMT-Bench’s OCR samples have an average word count and token count more than 5 times that of MME and over 2 times that of MMBench. Additionally, MMT-Bench includes a higher proportion of long text samples with a wider range of text lengths. This demonstrates MMT-Bench’s superiority in addressing the text scarcity issue in OCR tasks, providing a more reliable benchmark for comprehensively evaluating the performance of multimodal algorithms on OCR-related tasks.

## I. Some Details about the Benchmark Construction

### I.1. Metadata

Table A13. The format of the metadata.

Keys	Example 1	Example 2
image path	/path/to/image	/path/to/image
data source	animals90 from Kaggle	ReasonSeg
subtask name	Animal Recognition	Reason Seg
meta-task name	Visual Recognition	Visual Grounding
specific question template	What category of animal is shown in the picture?	Please provide the bounding box coordinates for the described object or area using the format [x1, y1, x2, y2]. QUESTION: {Referring Expression}
answer	rat	[801, 440, 1554, 956]
visual prompt	Natural Image	Natural Image
capabilities	Visual Recognition	Visual Reasoning, Visual Localization
(specific) category space	squirrel, hamster, bird, dog, cat...	-
(specific) referring expression	-	the objects that can protect the snail and prevent it from getting injured

**The uniform format of the metadata.** As shown in Table. A13, the metadata here is a dictionary, and its keys can be divided into two categories. The first category contains the essential keys, including image path, data source, subtask name, meta-task name, specific question template, answer (i.e., ground truth), visual prompt type, and capabilities required to solve the problem. The second category consists of keys that are specific to each sample and may vary accordingly. For animal recognition (example 1), the corresponding category space is retained for generating candidate choices, and referring expressions are preserved in the reasoning segmentation (example 2).

**How to get metadata?** We obtain the metadata mainly through a two-step process. In the first step, we write a python script for each dataset to directly extract relevant keys from the original dataset, such as image path, data source, and answer. Since samples within the same dataset share some similar characteristics, our co-authors predefined the specific question template, visual prompt, and capabilities for each dataset. The second step involves our co-authors manually performing a sample check (50%, up to 100 samples per dataset) on the samples defined in the first step, primarily focusing on sample-related keys, such as the specific question template and visual prompt.

### I.2. Prepare the Answer and Options

**How to ensure the wrong options are difficult distractors?** For most tasks, we use prompt engineering with GPT-4V using specific prompts for particular tasks. We employ in-context (2-shot or 3-shot) prompts. We review the options generated for each task and correct some erroneous ones. Interestingly, GPT-4 can generate sufficiently challenging options, which is mainly reflected in two aspects. First, the options can be confusing, especially in classification-related tasks (e.g. animal recognition), where the incorrect options tend to generate objects with semantics similar to the ground truth (GT). For

example, if the GT is "rat" the generated incorrect options might be "squirrel," "mouse," or "hamster." Second, the incorrect options can be ambiguous concerning the visual content. For instance, in captioning-related tasks, we use GPT-4V to take the image as input and assist in generating incorrect options, which can effectively alleviate the hallucination problem caused by text-only LLMs. For some problems, such as counting, OCR, and object detection, we can directly add perturbations to the GT using Python scripts to obtain distinguishable options. Throughout this process, we maintain consistency in word length across options, with a standard deviation of 1.1. In summary, we believe that the difficulty level of the options generated by MMT-Bench is reasonable.

Table A14. Visually grounded nature of our MMT-Bench dataset

Model	w/o visual	w/ visual	delta
Random	28.5	-	-
Frequency	31.7	-	-
ChatGPT-3.5	33.2	-	-
LLaVA-1.5-7B	31.6	49.7	18.1
LLaVA-1.5-13B	33.3	51.7	18.4
QWen-VL-Chat	32.3	52.5	20.2
Claude-3-Haiku(Anthropic, 2023)	33.1	52.2	19.1

**The answers in MMT-Bench are visual-grounded.** To ensure that the generated ground truth answers are visually grounded, we employ strategies as below. *Firstly*, when designing questions for each sample, we carefully craft specific question templates (as mentioned in Sec. I.1). We strive to avoid including specific information in the templates (for example, instead of asking "Who is the author of the Mona Lisa painting in the image?", we use a more general form such as "Who is the author of the painting in the image?"). This ensures that answering the question requires reliance on the image content. *Secondly*, when generating options, we create more complex options to prevent the model from bypassing the visual information and directly providing an answer based on correlations between options or prior knowledge. Furthermore, Table D demonstrates the visually grounded nature of our MMT-Bench dataset by comparing the performance of text-only LLMs and LVLMs with and without visual information. Without visual input, all models, including state-of-the-art LVLMs, perform similarly to the frequency-based guessing baseline (31.7%) and random guessing (28.5%), indicating the difficulty of answering questions correctly without visual context. For models like LLaVA and QWen that require an image as input, we simulate the absence of visual information by using a pure black image as input. This allows us to evaluate their performance in a text-only setting and compare it to their performance when provided with the actual visual information from our dataset. However, when provided with visual information, LVLMs show significant improvements, with performance gains ranging from 18.1% to 20.2%. This substantial gap between the text-only and visually-informed settings highlights the strong visual grounding of our dataset, confirming that the questions are designed to require models to rely on visual content and that the ground truth answers are indeed grounded in the visual information.

**How to prevent MMT-Bench from becoming overly focused on language understanding?** When defining questions and options, the co-authors proficient in each task domain strive to use simple and unambiguous specific question templates. The option generation process follows what is shown above, where we maintain the necessary clarity. To clarify, we report the statistics in Table. A15. For some questions, such as PLP (Pixel Level Perception) and Loc (Localization), which are localization problems, their longer question lengths are necessary to introduce the problem definition, answer format ( $(x, y, x, y)$ ), and detection conditions (detecting corresponding categories, such as the 80 COCO classes).

### I.3. Statistics of Image and Video in MMT-Bench

Table. A16 presents the number of images and videos in MMT-Bench compared to other multimodal benchmarks. MMT-Bench contains 25,732 single images, 3,800 image pairs (14,800 images), and 1,793 videos (10,572 video frames), demonstrating its comprehensive coverage of various visual data types. Compared to other benchmarks, MMT-Bench has the highest number of single images, more image pairs than MMMU, and more videos and video frames than MVBench and Seed-bench. This extensive coverage and diversity of visual data enable a thorough evaluation of multimodal models across a wide range of tasks and scenarios.

## J. OpenCompass' Protocol

**How often do steps 1-3 in OpenCompass' protocol fail?** We show the result in the following Table G. According to the table, the frequency of failure for steps 1-3 in OpenCompass' protocol varies across different LVLMs, but overall, the

Table A15. Average and standard deviation of the number of words in questions and answer choices.

Meta-task	question_avg	question_std	option_avg	option_std
MemU	7.0	0.0	12.4	8.0
VP	7.7	0.9	40.0	27.6
VCR	8.2	2.4	1.7	0.8
VR	9.5	1.2	1.5	1.0
AND	10.2	1.7	1.0	0.0
Emo	10.3	0.9	1.3	0.8
MedU	11.0	3.4	2.6	2.4
MIA	11.5	3.5	7.5	6.9
DU	12.1	7.3	8.9	16.1
IQT	13.0	0.0	2.0	0.0
3D	16.0	0.0	1.0	0.2
OCR	16.2	18.4	9.2	18.6
EA	16.6	2.0	27.9	17.2
Count	18.0	24.6	1.0	0.0
HLN	18.3	10.7	6.9	6.6
AUD	18.6	12.2	5.4	3.2
VC	23.4	22.0	24.3	22.2
AR	25.2	26.0	2.2	2.0
I2IT	27.5	3.5	3.0	1.0
IR	33.8	52.2	2.0	0.0
VI	34.7	3.3	1.0	0.0
TU	35.1	31.2	4.1	2.6
DKR	35.4	44.1	3.6	4.1
VPU	40.2	22.8	4.2	10.8
KD	44.7	8.7	43.5	14.8
CIM	58.3	16.5	3.3	0.9
RR	60.2	20.4	1.0	0.1
GN	60.9	2.3	6.2	1.9
IEJ	82.9	63.1	1.6	1.0
VG	83.8	10.1	4.0	0.0
Loc	97.1	24.8	9.9	9.8
PLP	98.1	49.5	22.2	32.7

Table A16. Statistics of image and video in MMT-Bench

Benchmark	Single Images	Multiple Images pairs(images)	Video (video frames)
MathVista	5,487	-	-
MMMU	10,705	845 (2,564)	-
MMBench	2,974	-	-
MME	1,187	-	-
MVBench	-	-	4,000
Seed-bench	14,233	-	5,009
MMT-Bench	25,732	3,800 (14,800)	1,793 (10,572)

extraction of model selections is highly successful. Most models achieve success rates above 87% at Step 1 (checking for option letters), with further improvements in subsequent steps. Models like BLIP-2, InternVL-Chat-V1.2-34B, and LLaVA-1.5-7B demonstrate exceptionally high success rates, consistently providing clear and well-structured responses that include option letters. However, some models, such as GPT-4V, GeminiProVision, Claude-3-Haiku, and QWen-VL-Plus, refuse to answer a small percentage of questions (0.875% to 10.136%), likely due to the sensitive nature of certain questions. In these cases, the model selection is set as option letter Z to differentiate it from a valid answer and to avoid random assignment. Overall, OpenCompass’ protocol proves highly effective in extracting model selections from LVLMS’ responses, with the multi-step approach ensuring accurate extraction and handling cases where models refuse to answer or provide unclear responses.

**Why is OpenCompass’ evaluation protocol chosen over other alternatives?** We chose OpenCompass’ multi-choice evaluation protocol for MMT-Bench due to its convenience, ease of use, and scalability. As discussed in our response above, OpenCompass’ protocol has proven to be highly effective in extracting model selections from LVLMS’ responses, with models like BLIP-2, InternVL-Chat-V1.2-34B, and LLaVA-1.5-7B achieving success rates over 99%. The multi-step approach ensures accurate extraction and handles cases where models refuse to answer or provide unclear responses. While alternative evaluation protocols have their merits, they also have limitations. MME’s yes-or-no evaluation simplifies the question-answering process but can lead to biased models and is resource-intensive, resulting in a limited sample size. MMBench’s evaluation strategy is resource-intensive and has high testing overhead, making it unsuitable for MMT-Bench’s scale. SeedBench calculates the log-likelihood for each candidate option, helping avoid the issue of models not directly

Table A17. Success rate of Steps 1-3.

Models	Step1	Step2 (Step2- Step1)	Step3 (Step3- Step2)	Refuse to Answer
GPT-4V	87.183	87.183 (0.0)	87.183 (0.0)	10.136
GeminiProVision	97.791	98.522 (+0.731)	98.522 (0.0)	0.875
Claude-3-Haiku	94.329	94.525 (+0.196)	94.525 (0.0)	4.023
QWen-VL-Plus	98.560	98.583 (+0.023)	98.583 (0.0)	1.015
BLIP2	99.448	99.770 (+0.322)	99.823 (+0.053)	-
InternVL-Chat-V1.2-34B	99.211	99.240 (+0.029)	99.243 (+0.003)	-
LLaVA-1.5-7B	99.994	99.997 (+0.003)	99.997 (0.0)	-

answering the options. However, this approach cannot be applied to API-based models.

## K. Computational Resources

Table A18. Resource consumption of some models evaluated on MMT-Bench.

Model	Resources	Times	Memory Utilization Per GPU
Claude-3-Haiku	USD 127	about 36h	-
Qwen-VL-Plus	USD 70.4	about 36h	-
LLaVA-v1.5-7B	1 x A100-80GB	120min	15890MiB
LLaVA-v1.5-13B	1 x A100-80GB	165min	26717MiB
LLaVA-v1.5-7B	8 x A100-80GB	12min	15890MiB
LLaVA-v1.5-13B	8 x A100-80GB	18min	26708MiB
QWen-VL-Chat	8 x A100-80GB	54min	21122MiB
InternVL-Chat-V1.2-34B	8 x A100-80GB	79min	78990MiB

We show the resource consumption of some models in Table. A18. The inference times vary among different models. For instance, the smaller LLaVA-v1.5-7B model takes only 12 minutes to complete the evaluation using 8 GPUs, while the larger InternVL-Chat-V1.2-34B model requires 79 minutes and around 80GB of memory. Our open-source codebase supports multi-GPU distributed inference, effectively accelerating the inference process. Although the prices for API models like Claude-3-Haiku (USD 127) and Qwen-VL-Plus (USD 70.4) are relatively high and their inference times are longer (about 36 hours), the inference costs for smaller open-source models remain manageable.

## L. Detailed Main Results

Table A19 to A37 display the performance of 32 models across all 162 subtasks, with accuracy used as the evaluation metric.

Table A19. Detail results of 32 LVLMS on Visual Grounding and Doc Understanding.

Model	Overall	Visual Grounding		Doc Understanding						
		Reason Seg	Referring Detection	Doc VQA	Visual Document Information Extraction	Chart to Text	Chart to Table	Clock Reading	Chart VQA	Table Structure Recognition
GPT-4o	65.5	58.7	42.5	54.0	83.5	92.0	84.5	59.5	62.5	84.8
InternVL-Chat-V1.2-34B	63.4	38.3	60.5	66.0	85.0	88.5	79.5	27.0	73.5	58.7
QwenVLMax	62.4	38.3	53.5	76.5	99.5	88.5	87.5	31.5	70.5	91.3
Qwen-VL-Plus	62.3	37.2	50.0	77.5	99.5	90.0	87.0	29.5	66.5	91.3
GeminiProVision	61.6	31.1	35.0	66.5	96.0	89.0	80.0	38.0	66.5	65.2
GPT-4V	61.1	48.0	29.0	56.0	79.0	86.0	73.5	31.0	64.0	84.8
LLaVA-Next-34B	60.8	43.4	69.0	74.0	96.5	92.0	71.0	29.0	73.5	47.8
XComposer2-7B	55.7	38.3	47.0	50.0	77.0	83.5	65.0	41.0	58.5	58.7
BLIP2-Flan-T5-XXL	54.8	28.1	39.0	32.5	57.0	79.0	57.0	20.5	22.5	32.6
Yi-VL-34B	54.2	31.6	51.5	56.5	61.5	85.5	48.0	23.5	50.5	69.6
Monkey	53.4	25.5	29.5	47.0	87.5	69.0	30.0	30.5	56.5	37.0
DeepSeek-VL-7B	53.2	33.7	43.0	42.5	68.5	74.5	32.5	38.5	59.0	47.8
Yi-VL-6B	53.2	33.7	50.5	44.0	67.0	84.0	35.5	25.0	55.0	76.1
LLaVA-Next-13B	53.0	36.7	52.5	51.5	82.0	82.5	29.0	29.5	59.5	41.3
TransCore-M	52.7	34.2	36.0	42.5	63.5	71.5	24.0	23.5	50.5	41.3
QWen-VL-Chat	52.5	24.5	28.5	46.0	79.5	82.0	37.0	22.0	55.0	39.1
Claude3V-Haiku	52.2	20.9	33.0	66.0	85.0	84.0	82.5	32.5	64.5	73.9
XComposer	52.1	26.0	28.5	34.5	45.0	64.0	36.0	25.5	27.5	23.9
mPLUG-Owl2	52.0	30.1	36.0	39.0	56.5	34.5	24.5	34.0	47.5	60.9
RBDash-v1-13B	51.8	34.2	39.5	37.0	60.0	74.5	26.0	23.5	46.0	69.6
LLaVA-v1.5-13B	51.7	36.7	38.0	35.5	61.5	77.5	23.0	30.5	50.5	41.3
CogVLM-Chat	51.6	28.1	29.5	47.5	75.0	75.5	22.0	24.5	58.0	41.3
ShareGPT4V-7B	51.5	33.2	37.5	45.5	62.0	65.0	24.0	22.0	51.5	65.2
LLaVA-Next-7B	51.1	29.6	44.5	50.5	76.5	68.5	24.0	28.0	57.0	43.5
LLaVA-v1.5-13B-XTuner	51.1	35.7	35.5	38.5	53.5	61.5	31.0	28.0	47.0	69.6
LlaVA-InternLM2-7B	50.8	32.7	41.0	35.0	45.0	66.0	34.0	25.0	39.5	19.6
LLaVA-v1.5-7B-Xtuner	50.2	31.1	39.5	37.5	54.5	49.0	28.0	21.0	46.5	63.0
SharedCaptioner	49.9	24.5	29.5	42.0	44.0	56.5	25.0	27.5	40.5	41.3
LLaVA-InternLM-7b	49.7	25.0	28.0	37.0	53.5	66.0	34.0	30.0	47.5	39.1
LLaVA-v1.5-7B	49.5	29.6	38.5	44.0	49.0	50.0	23.5	28.0	50.0	41.3
LLaMA-Adapter-v2-7B	40.4	21.9	28.5	26.5	31.0	34.0	21.5	20.5	38.5	39.1
VisualGLM_6b	38.6	28.6	33.5	27.5	41.0	61.5	22.0	27.5	35.5	58.7
Frequency	31.7	28.6	29.5	30.0	28.5	28.0	29.5	28.5	57.5	60.9
Random	28.5	27.6	22.0	27.0	25.5	28.5	24.0	24.5	30.5	52.2

MMT-Bench

Table A20. Detail results of 32 LVLMS on Action Recognition and Localization (part 1).

Model	Overall	Action Recognition					Localization			
		Gaze Estimation	Image based Action Recognition	General Action Recognition	Action Quality Assessment	Sign Language Recognition	Remote Sensing Object Detection	Rotated Object Detection	Small Object Detection	Camouflage Object Detection
GPT-4o	65.5	15.5	96.5	92.0	30.5	50.5	65.5	74.4	59.5	69.5
InternVL-Chat-V1.2-34B	63.4	27.0	91.5	84.0	30.0	30.5	63.5	46.7	64.5	57.0
QwenVLMax	62.4	22.0	97.5	87.5	32.5	29.0	52.5	53.3	54.5	54.0
Qwen-VL-Plus	62.3	21.0	98.5	86.5	30.0	31.0	53.0	60.0	59.5	53.5
GeminiProVision	61.6	24.5	92.5	87.5	42.5	40.0	37.5	50.0	43.0	33.5
GPT-4V	61.1	20.0	92.0	87.0	39.0	37.0	54.5	70.0	49.0	50.5
LLaVA-Next-34B	60.8	19.0	92.5	80.5	27.5	33.5	71.0	55.6	57.0	56.5
XComposer2-7B	55.7	23.0	90.0	69.0	21.0	28.5	48.0	56.7	52.5	38.0
BLIP2-Flan-T5-XXL	54.8	20.0	88.5	84.5	32.0	30.5	55.0	46.7	54.5	51.0
Yi-VL-34B	54.2	14.5	85.5	50.5	30.0	21.5	45.5	34.4	44.0	46.5
Monkey	53.4	24.0	91.5	79.0	19.5	32.5	36.0	38.9	40.5	42.5
DeepSeek-VL-7B	53.2	20.5	92.0	66.5	16.0	29.0	32.5	36.7	43.5	51.5
Yi-VL-6B	53.2	17.0	89.5	56.5	30.5	25.5	48.5	38.9	48.5	49.0
LLaVA-Next-13B	53.0	20.5	88.5	78.5	17.0	28.0	27.0	25.6	23.5	40.0
TransCore-M	52.7	23.0	87.0	78.5	22.0	24.0	36.5	35.6	44.0	36.5
QWen-VL-Chat	52.5	28.5	89.0	77.5	28.0	31.5	35.0	38.9	30.0	25.5
Claude3V-Haiku	52.2	14.5	81.5	71.0	36.0	28.0	41.0	46.7	44.0	40.5
XComposer	52.1	25.0	89.5	80.5	20.0	24.5	38.5	47.8	44.0	37.5
mPLUG-Owl2	52.0	20.0	89.5	78.5	15.5	22.5	38.5	42.2	41.5	50.0
RBDash-v1-13B	51.8	23.5	90.0	65.5	32.5	24.0	40.5	32.2	47.0	33.0
LLaVA-v1.5-13B	51.7	28.0	86.5	68.0	25.5	25.0	31.0	30.0	28.5	41.0
CogVLM-Chat	51.6	30.0	87.0	73.5	8.5	32.5	22.0	26.7	18.0	23.0
ShareGPT4V-7B	51.5	21.5	87.5	75.0	19.0	28.0	33.5	27.8	37.5	35.5
LLaVA-Next-7B	51.1	29.0	86.5	76.5	20.0	27.5	23.5	21.1	11.0	30.5
LLaVA-v1.5-13B-XTuner	51.1	29.5	85.0	69.0	15.0	22.5	41.0	30.0	34.5	42.0
LlaVA-InternLM2-7B	50.8	3.5	86.5	75.0	26.0	27.5	39.0	28.9	39.5	33.0
LLaVA-v1.5-7B-Xtuner	50.2	28.0	88.0	69.0	18.0	26.0	36.0	36.7	44.0	50.5
SharedCaptioner	49.9	30.0	69.5	71.5	31.5	31.0	29.5	41.1	55.5	39.5
LLaVA-InternLM-7b	49.7	19.5	88.5	70.0	19.5	30.5	35.0	41.1	42.0	40.0
LLaVA-v1.5-7B	49.5	22.0	88.5	75.0	19.5	28.0	27.5	30.0	31.5	36.0
LLaMA-Adapter-v2-7B	40.4	31.5	79.5	54.0	30.0	24.5	33.0	31.1	24.0	34.0
VisualGLM.6b	38.6	25.5	74.5	40.5	26.0	29.5	31.0	34.4	30.0	24.0
Frequency	31.7	31.5	29.0	30.0	32.0	27.5	27.5	28.9	27.5	27.0
Random	28.5	28.5	24.0	25.0	28.5	21.0	23.5	24.4	27.0	29.5

Table A21. Detail results of 32 LVLMS on Localization (part 2) and Visual Recognition (part 1).

Model	Overall	Localization					Visual Recognition			
		Salient Object Detection RGBD	Transparent Object Detection	Face Detection	Object Detection	Salient Object Detection RGB	Deepfake Detection	Weather Recognition	Season Recognition	Gesture Recognition
GPT-4o	65.5	48.0	77.5	76.5	83.5	52.5	49.0	95.4	85.5	79.5
InternVL-Chat-V1.2-34B	63.4	28.5	66.5	64.0	82.5	61.0	33.5	87.6	83.0	47.5
QwenVLMax	62.4	41.0	48.0	62.5	66.5	54.5	41.5	89.7	86.0	67.0
Qwen-VL-Plus	62.3	44.5	47.5	62.5	66.5	51.0	39.5	90.2	86.0	67.5
GeminiProVision	61.6	45.0	38.5	52.0	48.0	45.0	41.0	93.8	89.0	72.0
GPT-4V	61.1	40.5	62.5	74.0	68.0	43.0	37.5	95.4	87.0	68.5
LLaVA-Next-34B	60.8	36.5	66.0	60.5	76.0	70.0	42.0	82.0	78.0	50.5
XComposer2-7B	55.7	46.5	42.0	44.0	50.5	52.5	34.5	87.6	82.0	57.0
BLIP2-Flan-T5-XXL	54.8	68.0	56.0	61.5	57.5	36.5	29.5	84.5	78.5	40.5
Yi-VL-34B	54.2	42.5	55.5	36.0	57.5	61.5	31.5	74.2	78.5	41.5
Monkey	53.4	33.5	49.0	46.0	47.5	27.0	45.0	83.5	83.5	58.0
DeepSeek-VL-7B	53.2	40.0	53.5	34.0	45.5	40.5	35.0	83.5	76.5	67.5
Yi-VL-6B	53.2	47.5	61.5	46.0	56.0	49.0	37.0	76.3	74.5	43.0
LLaVA-Next-13B	53.0	44.0	41.0	38.0	38.5	42.5	59.5	79.4	74.5	48.0
TransCore-M	52.7	39.0	42.0	45.5	46.0	39.0	42.5	84.0	72.5	50.5
QWen-VL-Chat	52.5	30.5	33.5	41.0	42.0	27.0	48.0	82.0	78.0	54.0
Claude3V-Haiku	52.2	43.0	19.5	63.0	64.5	41.0	31.0	88.1	82.0	48.5
XComposer	52.1	37.5	36.0	37.0	49.0	36.0	38.5	82.5	83.0	48.0
mPLUG-Owl2	52.0	50.5	45.5	43.5	50.5	50.0	67.5	83.0	79.5	52.0
RBDash-v1-13B	51.8	42.0	46.0	41.0	56.0	42.0	36.0	79.9	72.5	49.0
LLaVA-v1.5-13B	51.7	38.0	48.0	43.5	46.5	42.5	49.5	82.0	73.5	48.5
CogVLM-Chat	51.6	35.0	28.0	16.0	28.0	25.5	44.5	86.1	85.5	49.0
ShareGPT4V-7B	51.5	40.5	39.0	32.0	44.5	33.5	72.0	82.5	74.0	51.0
LLaVA-Next-7B	51.1	49.0	34.5	15.5	34.5	46.0	66.0	76.8	75.5	49.5
LLaVA-v1.5-13B-XTuner	51.1	40.5	46.5	43.0	45.5	43.0	55.0	83.5	72.5	52.0
LlaVA-InternLM2-7B	50.8	32.0	51.0	37.0	44.0	45.5	33.0	80.9	73.5	43.0
LLaVA-v1.5-7B-Xtuner	50.2	36.5	50.5	33.5	43.5	39.0	46.5	83.5	78.0	46.5
SharedCaptioner	49.9	39.5	51.5	46.0	43.5	30.5	36.0	78.4	78.0	57.5
LLaVA-InternLM-7b	49.7	33.0	43.0	32.0	46.0	36.0	45.5	71.1	77.0	50.5
LLaVA-v1.5-7B	49.5	37.5	40.0	30.0	41.5	35.0	77.0	80.9	69.0	47.0
LLaMA-Adapter-v2-7B	40.4	38.5	28.5	28.0	42.0	33.0	52.0	73.2	62.0	39.5
VisualGLM.6b	38.6	49.5	31.5	38.0	29.5	30.0	26.5	59.8	58.0	45.0
Frequency	31.7	26.0	26.0	28.0	36.5	26.0	52.0	29.4	28.0	30.0
Random	28.5	28.5	29.0	33.0	37.0	21.0	54.5	22.2	25.5	28.0

## MMT-Bench

Table A22. Detail results of 32 LVLMs on **Visual Recognition** (part 2).

Model	Overall	Visual Recognition								
		Musical Instrument Recognition	Food Recognition	Landmark Recognition	Scene Recognition	Animal Recognition	Chemical Apparatus Recognition	Rock Recognition	Fashion Recognition	Logo Recognition
GPT-4o	65.5	98.5	96.5	100.0	81.5	98.0	78.0	68.0	82.0	98.5
InternVL-Chat-V1.2-34B	63.4	98.0	92.5	100.0	78.0	95.0	70.0	62.0	79.0	92.5
QwenVLMax	62.4	98.5	95.0	98.0	80.5	97.0	73.5	63.5	84.5	96.5
Qwen-VL-Plus	62.3	98.5	95.0	98.0	80.0	97.0	74.5	64.5	84.5	96.5
GeminiProVision	61.6	97.5	96.5	100.0	82.5	96.0	72.0	68.0	71.5	96.5
GPT-4V	61.1	98.0	95.0	100.0	83.0	98.5	65.5	59.5	79.0	92.5
LLaVA-Next-34B	60.8	91.5	91.0	96.0	77.0	90.0	59.0	55.0	73.5	90.0
XComposer2-7B	55.7	85.0	90.0	96.0	78.0	88.0	55.5	56.0	73.5	87.5
BLIP2-Flan-T5-XXL	54.8	93.5	92.5	98.0	77.5	92.0	54.0	48.5	80.0	76.5
Yi-VL-34B	54.2	93.5	86.5	100.0	77.0	92.5	50.0	53.5	74.5	87.0
Monkey	53.4	95.5	93.5	100.0	71.0	93.0	55.5	54.5	84.5	94.5
DeepSeek-VL-7B	53.2	94.5	93.0	98.0	77.5	94.0	57.5	50.0	81.5	96.0
Yi-VL-6B	53.2	92.0	86.0	94.0	77.0	89.5	57.5	42.0	78.5	89.5
LLaVA-Next-13B	53.0	90.0	86.0	94.0	77.0	88.0	59.0	58.0	69.0	88.5
TransCore-M	52.7	90.5	87.0	100.0	76.0	84.0	59.5	54.5	72.0	91.0
QWen-VL-Chat	52.5	95.0	90.5	98.0	74.5	94.0	48.5	51.0	79.5	95.5
Claude3V-Haiku	52.2	91.5	81.5	92.0	75.0	89.5	53.5	53.5	71.0	88.5
XComposer	52.1	93.5	90.5	98.0	77.0	89.5	57.0	45.0	81.0	88.5
mPLUG-Owl2	52.0	96.0	87.5	100.0	80.0	92.0	56.0	54.0	73.0	86.5
RBDash-v1-13B	51.8	87.5	87.5	96.0	77.5	87.5	57.0	50.0	65.5	85.5
LLaVA-v1.5-13B	51.7	93.0	90.5	100.0	77.5	89.0	48.5	59.0	68.5	88.0
CogVLM-Chat	51.6	95.5	90.0	98.0	80.5	93.5	67.5	58.5	79.0	92.5
ShareGPT4V-7B	51.5	92.0	84.0	100.0	79.5	89.0	52.5	50.5	70.0	86.0
LLaVA-Next-7B	51.1	93.0	82.5	94.0	80.0	89.5	53.5	49.0	70.0	81.5
LLaVA-v1.5-13B-XTuner	51.1	90.0	88.0	92.0	75.5	84.5	56.5	51.5	67.0	86.5
LlaVA-InternLM2-7B	50.8	90.0	90.0	96.0	78.0	89.5	63.5	52.0	72.0	78.0
LLaVA-v1.5-7B-Xtuner	50.2	90.0	89.0	98.0	76.0	85.5	51.0	44.5	68.5	80.0
SharedCaptioner	49.9	88.5	89.0	96.0	63.5	88.5	58.0	49.5	76.5	87.0
LLaVA-InternLM-7b	49.7	83.5	84.0	90.0	74.0	82.5	50.5	42.5	69.0	88.0
LLaVA-v1.5-7B	49.5	91.5	86.0	96.0	76.0	83.0	51.0	51.5	67.5	81.5
LLaMA-Adapter-v2-7B	40.4	87.0	74.0	88.0	63.5	79.0	38.5	37.5	47.0	75.0
VisualGLM.6b	38.6	82.0	69.5	82.0	58.5	80.0	36.5	31.5	45.0	71.5
Frequency	31.7	27.0	28.5	30.0	26.5	27.0	30.0	28.0	27.0	29.0
Random	28.5	26.5	24.5	26.0	24.0	25.5	29.5	25.0	23.0	28.5

Table A23. Detail results of 32 LVLMs on **Visual Recognition** (part 3).

Model	Overall	Visual Recognition								
		Astronomical Recognition	Painting Recognition	Color Recognition	Plant Recognition	Shape Recognition	Profession Recognition	Building Recognition	Electronic Object Recognition	Sports Recognition
GPT-4o	65.5	85.1	87.0	73.0	97.5	87.0	98.0	89.0	96.5	98.0
InternVL-Chat-V1.2-34B	63.4	69.1	78.5	52.0	95.5	82.5	97.0	77.0	98.5	94.0
QwenVLMax	62.4	76.6	81.5	60.5	96.5	81.0	92.0	87.5	96.5	97.5
Qwen-VL-Plus	62.3	75.5	81.5	61.0	96.5	81.5	94.5	87.5	96.5	96.5
GeminiProVision	61.6	72.3	84.0	66.5	95.0	84.5	97.0	82.5	95.5	98.0
GPT-4V	61.1	67.0	79.5	65.0	96.0	82.0	98.0	85.0	93.0	97.0
LLaVA-Next-34B	60.8	56.4	75.5	49.5	85.0	70.5	97.0	71.5	95.0	91.5
XComposer2-7B	55.7	45.7	78.5	53.0	82.0	70.0	97.0	75.0	91.0	91.0
BLIP2-Flan-T5-XXL	54.8	61.7	71.0	42.0	89.0	77.0	97.5	69.5	95.5	92.5
Yi-VL-34B	54.2	54.3	68.0	46.0	87.5	74.5	96.0	70.5	91.0	88.5
Monkey	53.4	66.0	76.0	39.5	95.0	79.0	98.0	78.0	95.0	90.0
DeepSeek-VL-7B	53.2	58.5	71.0	40.0	92.0	77.0	98.0	69.5	95.0	89.5
Yi-VL-6B	53.2	52.1	75.0	38.0	84.0	74.0	94.0	75.5	92.5	89.0
LLaVA-Next-13B	53.0	50.0	70.5	44.0	80.5	69.5	94.0	73.5	92.0	89.5
TransCore-M	52.7	56.4	74.0	39.0	76.5	73.0	97.5	76.0	93.0	85.5
QWen-VL-Chat	52.5	69.1	77.0	39.5	92.5	76.5	96.5	76.0	93.5	92.0
Claude3V-Haiku	52.2	64.9	67.5	71.0	91.5	74.0	91.0	72.5	75.5	84.0
XComposer	52.1	61.7	74.0	45.0	89.5	69.5	96.0	67.0	94.5	88.0
mPLUG-Owl2	52.0	62.8	71.5	42.0	92.0	74.5	98.0	71.0	92.0	91.0
RBDash-v1-13B	51.8	55.3	69.5	39.5	79.0	71.5	97.0	65.0	92.0	87.0
LLaVA-v1.5-13B	51.7	50.0	69.5	38.5	83.5	67.5	95.0	71.0	93.5	89.0
CogVLM-Chat	51.6	73.4	78.5	35.0	92.5	80.0	97.0	77.5	96.0	92.5
ShareGPT4V-7B	51.5	56.4	72.0	35.5	84.5	67.5	98.5	77.0	93.0	88.5
LLaVA-Next-7B	51.1	50.0	76.0	45.0	82.0	68.0	95.5	70.0	93.5	88.5
LLaVA-v1.5-13B-XTuner	51.1	48.9	68.0	43.0	79.5	67.5	96.0	73.0	91.5	89.0
LlaVA-InternLM2-7B	50.8	46.8	64.5	48.5	85.5	75.0	97.5	67.0	94.0	89.0
LLaVA-v1.5-7B-Xtuner	50.2	52.1	67.5	41.5	79.0	66.5	96.0	69.0	92.5	90.0
SharedCaptioner	49.9	55.3	64.0	39.5	88.0	70.0	94.5	67.0	93.0	87.0
LLaVA-InternLM-7b	49.7	53.2	68.5	36.5	81.5	68.0	96.5	57.5	91.5	85.5
LLaVA-v1.5-7B	49.5	56.4	66.0	45.0	83.0	62.5	97.0	74.5	93.0	87.5
LLaMA-Adapter-v2-7B	40.4	53.2	52.0	36.5	67.5	60.5	94.0	56.0	79.5	82.5
VisualGLM.6b	38.6	42.6	45.5	27.5	67.5	52.0	88.5	32.0	82.5	81.5
Frequency	31.7	29.8	31.5	27.0	29.5	26.5	27.5	28.5	29.0	35.0
Random	28.5	30.9	21.0	20.0	23.5	21.5	27.5	24.5	25.5	19.5

Table A24. Detail results of 32 LVLMS on Visual Recognition (part 4).

Model	Overall	Visual Recognition								
		Disaster Recognition	Celebrity Recognition	Vehicle Recognition	National Flag Recognition	Abstract Visual Recognition	Animated Character Recognition	Texture Material Recognition	Film and Television Recognition	Sculpture Recognition
GPT-4o	65.5	84.5	86.5	97.5	99.5	85.5	87.0	79.0	97.5	96.0
InternVL-Chat-V1.2-34B	63.4	80.0	85.5	98.5	83.0	76.5	49.0	76.5	95.0	94.0
QwenVLMax	62.4	79.5	87.5	96.0	88.0	82.0	56.0	76.5	94.0	70.0
Qwen-VL-Plus	62.3	78.5	89.0	97.0	93.5	82.0	56.0	76.5	93.5	72.0
GeminiProVision	61.6	82.0	93.0	96.5	98.5	88.0	58.5	76.0	93.5	96.0
GPT-4V	61.1	87.5	88.0	97.0	91.0	78.0	67.5	75.0	95.5	98.0
LLaVA-Next-34B	60.8	75.5	81.5	99.0	67.0	72.0	46.5	67.5	92.5	92.0
XComposer2-7B	55.7	78.5	74.0	96.0	63.0	66.0	42.0	69.0	93.0	88.0
BLIP2-Flan-T5-XXL	54.8	57.5	73.0	96.5	86.5	70.5	30.5	77.0	83.5	92.0
Yi-VL-34B	54.2	72.5	79.5	94.5	87.0	70.5	43.0	58.5	90.5	88.0
Monkey	53.4	68.0	77.0	97.5	95.5	72.0	41.5	71.0	94.0	94.0
DeepSeek-VL-7B	53.2	50.0	69.0	97.0	70.5	72.5	31.0	77.5	86.5	88.0
Yi-VL-6B	53.2	66.0	76.0	97.0	81.0	64.5	35.5	59.5	87.0	88.0
LLaVA-Next-13B	53.0	67.0	77.5	97.0	60.5	69.0	37.0	71.5	85.0	86.0
TransCore-M	52.7	65.5	78.0	98.0	61.5	73.0	31.5	63.5	81.5	82.0
QWen-VL-Chat	52.5	68.0	65.0	95.5	94.5	77.0	44.5	72.0	88.0	92.0
Claude3V-Haiku	52.2	84.0	55.5	91.5	74.0	68.0	45.5	62.5	84.5	86.0
XComposer	52.1	53.0	73.5	96.0	87.0	70.5	40.5	71.5	89.0	84.0
mPLUG-Owl2	52.0	62.5	79.0	96.5	83.5	66.5	38.5	67.0	88.5	86.0
RBDash-v1-13B	51.8	71.5	80.0	98.0	56.0	69.5	35.5	66.5	85.0	88.0
LLaVA-v1.5-13B	51.7	62.5	76.5	98.5	58.0	69.0	36.0	75.5	85.5	90.0
CogVLM-Chat	51.6	65.5	42.0	98.0	87.0	75.5	38.5	65.5	91.0	94.0
ShareGPT4V-7B	51.5	64.5	76.0	96.0	64.5	71.0	40.0	62.5	84.5	80.0
LLaVA-Next-7B	51.1	66.0	77.5	96.5	57.0	67.5	37.0	65.0	86.0	74.0
LLaVA-v1.5-13B-XTuner	51.1	52.0	75.5	93.5	57.5	68.0	35.5	68.5	83.5	82.0
LLaVA-InternLM2-7B	50.8	72.5	75.0	97.5	55.0	70.0	37.5	70.5	82.0	86.0
LLaVA-v1.5-7B-Xtuner	50.2	66.5	70.5	96.0	56.0	66.0	36.5	75.0	86.5	84.0
SharedCaptioner	49.9	70.0	70.0	95.5	76.0	68.5	32.0	71.0	84.0	74.0
LLaVA-InternLM-7b	49.7	62.5	69.0	97.0	54.5	65.5	32.0	72.5	80.0	80.0
LLaVA-v1.5-7B	49.5	62.0	72.5	96.5	61.0	68.0	34.0	62.5	80.0	90.0
LLaMA-Adapter-v2-7B	40.4	35.0	49.0	90.5	67.0	61.5	22.0	65.0	64.5	54.0
VisualGLM.6b	38.6	43.5	31.0	92.0	65.0	49.0	22.0	49.0	51.0	68.0
Frequency	31.7	29.5	29.5	26.5	29.5	26.0	26.5	29.5	25.5	32.0
Random	28.5	23.5	23.0	24.0	22.0	25.5	27.0	31.5	27.0	30.0

Table A25. Detail results of 32 LVLMS on Visual Recognition (part 5) and GUI.

Model	Overall	Visual Recognition				GUI			
		Age Gender Race recognition	Weapon Recognition	Religious Recognition	Waste Recognition	GUI General	GUI Google APP	GUI Web Shopping	GUI Installation
GPT-4o	65.5	71.0	91.5	87.0	97.0	53.5	41.0	42.0	48.0
InternVL-Chat-V1.2-34B	63.4	79.5	91.5	82.0	93.5	46.5	40.5	42.5	35.5
QwenVLMax	62.4	81.0	70.5	86.0	97.0	35.5	30.5	27.5	32.5
Qwen-VL-Plus	62.3	81.5	47.0	86.5	96.5	37.5	32.0	28.0	33.0
GeminiProVision	61.6	81.5	71.5	83.5	94.0	44.5	28.0	29.5	31.5
GPT-4V	61.1	78.0	86.0	78.0	97.5	48.0	38.0	37.0	38.5
LLaVA-Next-34B	60.8	83.5	74.5	77.0	90.0	49.0	39.0	44.5	33.0
XComposer2-7B	55.7	76.5	78.0	66.0	90.0	38.0	26.0	20.0	33.5
BLIP2-Flan-T5-XXL	54.8	76.0	85.0	68.0	92.0	35.0	33.0	34.0	31.5
Yi-VL-34B	54.2	74.0	82.0	67.5	85.5	36.5	30.5	31.5	31.0
Monkey	53.4	78.5	77.5	79.5	92.0	32.0	28.0	29.0	30.0
DeepSeek-VL-7B	53.2	75.5	83.5	63.0	88.0	35.5	36.0	35.5	40.0
Yi-VL-6B	53.2	75.0	84.0	64.5	85.0	37.5	35.0	34.5	32.0
LLaVA-Next-13B	53.0	77.0	74.5	65.5	88.5	39.0	35.0	45.0	45.0
TransCore-M	52.7	75.0	82.0	61.5	86.5	36.5	30.5	36.5	39.0
QWen-VL-Chat	52.5	70.0	82.0	78.0	85.5	33.0	21.5	32.5	34.0
Claude3V-Haiku	52.2	77.5	71.0	73.5	89.5	39.5	35.5	30.0	35.5
XComposer	52.1	79.0	87.0	67.0	85.5	32.5	36.5	35.0	33.0
mPLUG-Owl2	52.0	79.5	74.5	64.5	89.5	26.5	25.0	31.0	28.5
RBDash-v1-13B	51.8	75.0	75.5	55.5	83.0	34.5	29.0	31.5	34.5
LLaVA-v1.5-13B	51.7	76.0	79.0	66.0	86.5	42.5	35.5	37.0	35.0
CogVLM-Chat	51.6	78.5	84.0	68.0	91.0	41.0	24.0	26.5	28.0
ShareGPT4V-7B	51.5	74.5	80.5	63.5	84.5	33.0	25.5	30.0	27.0
LLaVA-Next-7B	51.1	77.5	74.5	67.0	90.5	37.5	24.0	30.0	35.0
LLaVA-v1.5-13B-XTuner	51.1	77.5	78.5	63.0	90.0	35.0	41.0	42.0	45.5
LLaVA-InternLM2-7B	50.8	79.0	76.5	68.0	88.5	39.5	29.5	36.5	35.5
LLaVA-v1.5-7B-Xtuner	50.2	84.5	74.5	63.5	87.5	31.0	26.0	32.5	32.5
SharedCaptioner	49.9	74.0	81.5	64.5	83.0	28.0	30.0	26.0	34.5
LLaVA-InternLM-7b	49.7	70.5	76.0	65.0	83.0	26.5	36.0	40.0	48.0
LLaVA-v1.5-7B	49.5	76.0	78.0	63.5	81.5	25.5	15.5	20.5	20.0
LLaMA-Adapter-v2-7B	40.4	75.0	69.5	56.5	71.5	43.0	22.0	24.0	30.5
VisualGLM.6b	38.6	38.5	58.5	49.5	43.0	21.0	18.5	19.0	26.0
Frequency	31.7	36.5	32.0	28.0	43.0	31.0	27.0	28.5	29.5
Random	28.5	38.5	33.0	30.0	36.0	22.5	30.0	21.5	23.5

Table A26. Detail results of 32 LVLMs on **OCR** and **Image-to-Image Translation**.

Model	Overall	OCR				Image-to-Image Translation	
		Font Recognition	Handwritten Text Recognition	Handwritten Mathematical Expression Recognition	Scene Text Recognition	Jigsaw Puzzle Solving	Image Colorization
GPT-4o	65.5	44.5	79.0	61.0	81.5	25.5	55.0
InternVL-Chat-V1.2-34B	63.4	23.5	64.0	60.0	94.5	44.5	21.0
QwenVLMax	62.4	31.0	80.0	63.0	93.5	29.0	44.5
Qwen-VL-Plus	62.3	31.0	79.0	59.0	93.5	42.5	33.0
GeminiProVision	61.6	24.0	63.0	57.0	94.0	21.0	42.0
GPT-4V	61.1	21.5	68.0	53.0	80.0	26.0	50.5
LLaVA-Next-34B	60.8	29.0	74.0	58.0	95.5	29.0	16.5
XComposer2-7B	55.7	21.5	32.0	29.0	93.0	21.0	42.0
BLIP2-Flan-T5-XXL	54.8	14.5	59.0	40.0	79.0	25.0	31.5
Yi-VL-34B	54.2	27.5	59.0	59.0	86.5	19.0	20.0
Monkey	53.4	21.0	46.0	44.0	93.0	26.0	33.0
DeepSeek-VL-7B	53.2	29.0	61.0	61.0	93.5	26.5	20.5
Yi-VL-6B	53.2	24.0	51.0	50.0	87.5	29.0	24.5
LLaVA-Next-13B	53.0	22.0	47.0	43.0	95.0	30.5	22.0
TransCore-M	52.7	23.5	55.0	33.0	90.0	25.5	24.5
QWen-VL-Chat	52.5	16.5	44.0	37.0	90.0	28.0	33.0
Claude3V-Haiku	52.2	17.5	66.0	48.0	86.0	23.5	22.0
XComposer	52.1	15.5	34.0	44.0	83.0	21.5	27.5
mPLUG-Owl2	52.0	21.0	32.0	35.0	90.0	26.0	25.0
RBDash-v1-13B	51.8	29.5	48.0	45.0	92.0	20.0	25.0
LLaVA-v1.5-13B	51.7	28.0	39.0	51.0	89.0	27.5	28.5
CogVLM-Chat	51.6	16.0	42.0	50.0	86.0	26.5	21.0
ShareGPT4V-7B	51.5	23.0	34.0	45.0	89.0	19.0	24.5
LLaVA-Next-7B	51.1	23.0	46.0	44.0	95.0	22.0	23.5
LLaVA-v1.5-13B-XTuner	51.1	18.0	51.0	28.0	90.0	28.0	24.0
LlaVA-InternLM2-7B	50.8	28.0	42.0	40.0	88.0	0.0	28.5
LLaVA-v1.5-7B-Xtuner	50.2	20.0	39.0	35.0	90.0	26.5	24.5
SharedCaptioner	49.9	23.0	45.0	39.0	84.0	22.5	27.5
LLaVA-InternLM-7b	49.7	14.5	49.0	35.0	92.0	25.5	24.5
LLaVA-v1.5-7B	49.5	23.0	32.0	36.0	89.0	19.5	25.0
LLaMA-Adapter-v2-7B	40.4	22.0	28.0	37.0	53.0	18.5	18.0
VisualGLM_6b	38.6	6.0	29.0	31.0	69.0	29.5	24.0
Frequency	31.7	28.5	30.0	33.0	30.0	29.0	32.5
Random	28.5	26.0	30.0	27.0	26.0	21.5	21.0

Table A27. Detail results of 32 LVLMs on **Temporal Understanding** and **Relation Reasoning**.

Model	Overall	Temporal Understanding					Relation Reasoning			
		Next Image Prediction	MeViS	Temporal Anticipation	Temporal Ordering	Temporal Localization	Social Relation Recognition	Human Object Interaction Recognition	Scene Graph Recognition	Human Interaction Understanding
GPT-4o	65.5	13.0	34.0	71.0	23.0	52.3	59.0	75.0	77.5	87.0
InternVL-Chat-V1.2-34B	63.4	36.0	63.5	25.5	30.6	70.5	86.5	19.5	74.0	40.2
QwenVLMax	62.4	34.5	33.0	62.5	21.0	45.6	70.5	80.0	18.5	76.5
Qwen-VL-Plus	62.3	31.0	62.5	20.5	46.6	67.0	81.0	16.0	77.0	37.8
GeminiProVision	61.6	30.0	36.0	68.5	26.0	42.0	48.0	77.0	29.5	79.5
GPT-4V	61.1	16.5	33.0	73.5	27.0	44.0	55.5	71.0	69.0	68.5
LLaVA-Next-34B	60.8	26.0	67.0	22.5	44.0	67.5	85.5	18.5	75.5	40.2
XComposer2-7B	55.7	29.0	32.0	64.5	25.0	32.6	36.5	83.5	7.5	74.5
BLIP2-Flan-T5-XXL	54.8	26.0	66.0	64.0	26.5	31.1	52.0	74.0	37.0	49.0
Yi-VL-34B	54.2	24.5	56.0	28.5	24.4	62.5	73.5	28.5	64.5	40.2
Monkey	53.4	27.5	36.0	63.5	24.5	32.6	63.5	74.0	42.0	67.5
DeepSeek-VL-7B	53.2	22.5	51.0	24.0	35.2	33.0	68.0	29.5	70.0	27.6
Yi-VL-6B	53.2	32.0	52.5	27.0	22.3	54.0	72.0	31.0	61.5	30.7
LLaVA-Next-13B	53.0	30.5	61.5	25.0	30.6	47.0	71.5	31.5	69.5	27.6
TransCore-M	52.7	24.0	24.0	49.5	22.0	29.0	44.5	68.0	31.0	69.5
QWen-VL-Chat	52.5	26.0	37.0	71.5	25.0	26.9	61.5	62.0	36.0	70.0
Claude3V-Haiku	52.2	27.0	59.0	24.5	35.8	33.0	62.0	50.0	49.5	38.6
XComposer	52.1	25.5	38.5	47.0	24.0	30.6	53.0	72.0	11.5	64.5
mPLUG-Owl2	52.0	26.5	34.5	47.0	24.0	27.5	34.5	71.5	37.5	64.5
RBDash-v1-13B	51.8	23.5	26.5	51.0	23.0	34.2	41.0	77.0	34.5	71.0
LLaVA-v1.5-13B	51.7	26.0	26.0	60.5	25.5	28.5	52.5	73.5	38.0	69.0
CogVLM-Chat	51.6	23.5	36.5	54.5	25.0	37.8	62.0	63.0	42.5	79.0
ShareGPT4V-7B	51.5	28.5	30.0	52.5	25.0	24.4	51.0	72.5	45.0	70.5
LLaVA-Next-7B	51.1	33.5	58.0	22.5	24.4	57.0	67.5	35.0	65.0	28.3
LLaVA-v1.5-13B-XTuner	51.1	21.0	24.0	43.0	24.0	29.5	34.5	72.0	34.5	68.5
LlaVA-InternLM2-7B	50.8	19.5	26.5	52.5	5.5	30.1	43.5	80.0	35.0	73.5
LLaVA-v1.5-7B-Xtuner	50.2	25.0	28.5	44.0	22.5	29.5	33.5	73.5	36.0	72.5
SharedCaptioner	49.9	25.5	30.5	60.0	20.5	26.4	60.5	70.5	39.5	67.5
LLaVA-InternLM-7b	49.7	20.5	31.0	40.0	20.0	28.5	27.0	72.0	37.0	73.5
LLaVA-v1.5-7B	49.5	22.5	29.5	55.5	19.0	32.1	59.0	65.0	40.0	68.0
LLaMA-Adapter-v2-7B	40.4	23.5	31.5	57.5	20.5	22.3	53.5	33.5	45.5	47.0
VisualGLM_6b	38.6	24.0	41.0	35.0	26.5	27.5	44.5	35.0	47.0	48.5
Frequency	31.7	28.5	30.0	28.0	29.0	30.1	27.0	26.0	52.0	29.0
Random	28.5	30.0	28.0	24.0	25.0	28.0	25.5	22.0	57.0	29.0

MMT-Bench

Table A28. Detail results of 32 LVLMs on **Discipline Knowledge Reasoning, Intelligence Quotient Test and Embodied AI.**

Model	Discipline Knowledge Reasoning								Intelligence Quotient Test		Embodied AI
	Overall	Science	Health Medicine	Art and Design	Humanities Science	Social	Tech Engineering	Business	Ravens Progressive Matrices	Navigation	
GPT-4o	65.5	52.0	69.3	76.4	78.6	44.5	58.3		15.5	88.5	
InternVL-Chat-V1.2-34B	63.4	54.3	65.5	72.3	45.1	48.3	11.0		84.0	74.5	
QwenVLMax	62.4	40.2	51.4	64.5	70.5	39.0	45.8		12.5	85.5	
Qwen-VL-Plus	62.3	50.7	65.5	72.3	38.5	42.5	11.0		84.5	85.0	
GeminiProVision	61.6	37.0	49.3	63.6	67.0	35.2	43.3		11.0	74.5	
GPT-4V	61.1	44.9	66.4	67.3	73.2	46.7	60.0		14.5	85.0	
LLaVA-Next-34B	60.8	48.6	65.5	64.3	46.2	50.8	13.0		76.5	80.0	
XComposer2-7B	55.7	33.1	42.1	56.4	61.6	33.0	36.7		8.0	50.5	
BLIP2-Flan-T5-XXL	54.8	28.3	30.0	51.8	49.1	34.6	28.3		14.0	80.5	
Yi-VL-34B	54.2	50.0	61.8	61.6	38.5	40.0	14.0		68.5	57.0	
Monkey	53.4	26.0	33.6	49.1	50.0	31.9	38.3		11.0	46.0	
DeepSeek-VL-7B	53.2	40.7	50.9	47.3	34.6	32.5	12.5		37.0	67.5	
Yi-VL-6B	53.2	44.3	54.5	51.8	33.5	44.2	13.0		47.0	48.0	
LLaVA-Next-13B	53.0	35.7	53.6	58.0	31.9	37.5	14.5		50.0	70.0	
TransCore-M	52.7	23.6	35.7	50.0	48.2	35.2	37.5		15.0	39.5	
QWen-VL-Chat	52.5	22.0	33.6	50.9	47.3	30.8	29.2		13.5	55.0	
Claude3V-Haiku	52.2	41.4	54.5	57.1	29.7	44.2	15.5		59.5	54.5	
XComposer	52.1	27.6	35.0	49.1	56.2	33.0	35.0		14.0	38.5	
mPLUG-Owl2	52.0	27.6	40.0	50.0	46.4	35.2	29.2		18.0	35.0	
RBDash-v1-13B	51.8	33.1	37.1	51.8	47.3	34.1	25.0		14.0	53.5	
LLaVA-v1.5-13B	51.7	29.1	34.3	54.5	54.5	34.1	30.0		13.5	40.5	
CogVLM-Chat	51.6	23.6	30.7	49.1	49.1	34.6	34.2		14.0	48.0	
ShareGPT4V-7B	51.5	27.6	37.1	57.3	49.1	32.4	30.0		14.0	42.0	
LLaVA-Next-7B	51.1	37.9	56.4	50.9	32.4	30.8	13.5		47.5	65.0	
LLaVA-v1.5-13B-XTuner	51.1	28.3	35.0	43.6	49.1	35.2	32.5		14.0	33.5	
LlaVA-InternLM2-7B	50.8	25.2	42.1	50.0	52.7	14.8	17.5		0.0	35.5	
LLaVA-v1.5-7B-Xtuner	50.2	28.3	37.1	47.3	49.1	35.7	25.8		11.5	32.0	
SharedCaptioner	49.9	25.2	37.9	50.9	47.3	32.4	36.7		14.5	45.0	
LLaVA-InternLM-7b	49.7	23.6	39.3	47.3	55.4	29.7	33.3		14.0	46.5	
LLaVA-v1.5-7B	49.5	24.4	34.3	52.7	46.4	28.6	27.5		12.5	42.5	
LLaMA-Adapter-v2-7B	40.4	29.9	30.7	38.2	33.0	32.4	27.5		11.0	25.0	
VisualGLM.6b	38.6	20.5	29.3	30.9	41.1	26.9	29.2		14.0	37.5	
Frequency	31.7	32.3	27.9	30.9	28.6	27.5	29.2		18.0	28.0	
Random	28.5	24.4	25.7	22.7	25.0	25.3	29.2		10.5	27.5	

Table A29. Detail results of 32 LVLMs on **Emotion Quotient Test and Visual Illusion.**

Model	Overall	Emotion Quotient Test						Visual Illusion				
		Facial Expression Change Recognition	Scene Emotion Recognition	Micro-Expression Recognition	Artwork Emotion Recognition	Body Emotion Recognition	Facial Expression Recognition	Color Constancy	Color Assimilation	Geometrical Relativity	Geometrical Perspective	Color Contrast
GPT-4o	65.5	89.5	61.0	31.0	44.0	46.5	75.0	23.6	39.0	24.0	62.5	46.5
InternVL-Chat-V1.2-34B	63.4	59.0	30.0	42.5	43.0	70.5	47.2	34.5	44.5	82.5	75.0	71.0
QwenVLMax	62.4	85.0	58.5	19.5	40.5	42.5	63.0	20.8	49.5	29.0	61.7	43.5
Qwen-VL-Plus	62.3	58.5	19.5	40.5	41.0	62.0	27.8	47.5	29.0	58.3	43.0	73.5
GeminiProVision	61.6	73.0	59.0	40.0	50.5	42.5	66.0	38.9	53.5	46.0	43.3	56.0
GPT-4V	61.1	78.5	61.5	30.5	46.5	45.5	68.0	8.3	31.5	19.5	43.3	40.5
LLaVA-Next-34B	60.8	61.0	27.0	37.0	42.5	72.0	62.5	42.0	43.5	78.3	55.0	70.0
XComposer2-7B	55.7	72.0	53.0	36.5	47.0	47.0	66.0	51.4	39.5	47.5	75.8	49.0
BLIP2-Flan-T5-XXL	54.8	44.5	51.5	20.5	44.5	36.5	61.0	63.9	47.0	58.0	60.0	49.0
Yi-VL-34B	54.2	46.5	30.0	44.5	31.5	56.5	29.2	45.5	34.0	55.8	53.5	67.5
Monkey	53.4	51.5	49.5	24.0	46.0	37.5	62.0	59.7	30.5	61.5	62.5	62.5
DeepSeek-VL-7B	53.2	55.5	30.0	39.5	35.5	69.0	33.3	27.5	52.0	54.2	56.0	49.0
Yi-VL-6B	53.2	57.0	33.0	43.5	42.5	56.5	44.4	37.0	61.0	63.3	60.0	57.5
LLaVA-Next-13B	53.0	55.0	20.0	30.0	33.5	66.0	62.5	30.5	54.5	73.3	66.0	63.0
TransCore-M	52.7	67.0	54.5	26.0	29.0	34.5	67.0	47.2	24.0	61.0	75.0	71.0
QWen-VL-Chat	52.5	48.5	55.5	26.5	42.5	38.0	61.5	48.6	48.0	50.0	50.0	58.0
Claude3V-Haiku	52.2	53.0	28.0	39.5	34.0	43.0	13.9	38.5	58.5	55.8	56.5	5.0
XComposer	52.1	62.0	58.5	26.5	29.0	36.5	63.0	51.4	47.5	49.0	60.0	59.0
mPLUG-Owl2	52.0	45.5	55.0	26.5	34.5	45.5	56.0	66.7	30.0	65.5	60.0	70.5
RBDash-v1-13B	51.8	69.5	50.0	13.0	29.0	33.0	66.0	59.7	30.5	63.0	72.5	60.0
LLaVA-v1.5-13B	51.7	66.5	57.5	25.5	21.0	37.0	64.5	63.9	29.0	64.5	66.7	66.5
CogVLM-Chat	51.6	51.5	56.5	30.0	50.5	43.5	70.0	55.6	45.5	51.0	60.8	49.0
ShareGPT4V-7B	51.5	63.5	56.5	26.0	23.0	33.5	63.5	52.8	26.5	60.0	65.8	67.5
LLaVA-Next-7B	51.1	59.0	27.0	33.5	37.5	62.5	69.4	29.5	63.0	70.8	68.0	57.5
LLaVA-v1.5-13B-XTuner	51.1	72.5	63.5	25.0	43.0	37.5	64.5	44.4	26.0	60.0	75.0	66.5
LlaVA-InternLM2-7B	50.8	71.0	55.5	35.5	33.0	42.5	69.0	50.0	28.5	52.0	73.3	57.5
LLaVA-v1.5-7B-Xtuner	50.2	61.0	57.0	28.5	30.5	31.5	57.0	61.1	30.5	62.0	76.7	70.0
SharedCaptioner	49.9	54.5	47.0	12.0	36.0	28.0	62.0	47.2	41.0	48.5	81.7	57.0
LLaVA-InternLM-7b	49.7	59.0	58.0	28.5	42.0	36.0	58.5	63.9	26.0	66.5	67.5	63.5
LLaVA-v1.5-7B	49.5	65.0	57.5	18.0	21.5	37.5	55.5	56.9	28.0	64.0	70.0	69.0
LLaMA-Adapter-v2-7B	40.4	23.0	52.0	20.0	34.0	32.5	54.5	37.5	40.5	29.0	36.7	38.5
VisualGLM.6b	38.6	28.5	43.5	23.5	28.0	31.0	44.0	25.0	50.5	29.0	29.2	11.0
Frequency	31.7	28.0	29.5	30.0	33.0	31.0	29.0	52.8	51.0	50.5	53.3	53.0
Random	28.5	26.0	19.5	27.5	30.5	23.0	26.0	48.6	50.0	50.5	51.7	53.0

## MMT-Bench

Table A30. Detail results of 32 LVLMs on Meme Understanding, Counting and Hallucination.

Model	Overall	Meme Understanding		Counting				Hallucination			
		Meme Video Understanding	Meme Image Understanding	Counting by Visual Prompting	Counting by Category	Crowd Counting	Counting by Reasoning	Order Hallucination	Relation Hallucination	Attribute Hallucination	Exist Hallucination
GPT-4o	65.5	76.0	92.5	39.5	72.9	35.5	70.5	39.5	69.5	84.0	84.5
InternVL-Chat-V1.2-34B	63.4	88.0	37.0	75.5	56.0	78.0	56.0	85.5	80.0	85.5	78.5
QwenVLMax	62.4	75.5	83.0	41.0	74.9	42.5	69.0	54.5	81.5	73.0	87.5
Qwen-VL-Plus	62.3	82.5	39.5	75.0	44.0	70.0	55.5	81.0	74.0	87.0	37.5
GeminiProVision	61.6	75.0	76.5	40.5	68.5	52.0	64.5	39.0	59.5	82.5	82.5
GPT-4V	61.1	73.0	89.5	31.5	66.4	32.5	69.0	43.0	68.0	82.0	81.5
LLaVA-Next-34B	60.8	88.0	23.0	72.6	56.5	81.0	55.0	83.0	76.5	87.5	33.5
XComposer2-7B	55.7	55.5	87.0	25.5	73.2	29.5	76.0	44.0	76.0	71.5	86.5
BLIP2-Flan-T5-XXL	54.8	70.5	82.0	23.0	53.8	12.0	30.5	51.0	62.0	72.0	79.5
Yi-VL-34B	54.2	83.0	29.5	65.2	44.5	69.0	59.0	84.5	71.5	79.5	27.5
Monkey	53.4	52.0	87.0	28.0	66.9	24.0	55.5	43.0	59.0	68.5	82.0
DeepSeek-VL-7B	53.2	84.5	35.0	70.8	18.0	52.5	36.5	63.0	73.0	84.5	22.0
Yi-VL-6B	53.2	78.0	31.0	66.2	44.5	60.5	53.5	79.0	69.5	79.0	28.0
LLaVA-Next-13B	53.0	79.5	30.5	68.5	54.0	65.0	49.5	74.5	72.0	84.0	24.0
TransCore-M	52.7	70.0	83.5	36.0	65.0	57.5	59.5	51.0	74.5	79.0	83.0
QWen-VL-Chat	52.5	64.5	84.0	31.0	63.2	39.0	53.5	45.0	62.5	71.5	76.5
Claude3V-Haiku	52.2	75.0	32.5	57.0	45.0	59.0	43.5	62.0	82.0	78.5	32.0
XComposer	52.1	58.0	69.5	32.5	66.2	23.5	37.5	34.0	74.0	75.0	83.0
mPLUG-Owl2	52.0	49.5	68.5	33.0	68.6	36.0	53.0	33.0	63.0	72.5	85.0
RBDash-v1-13B	51.8	54.5	80.5	31.0	65.1	46.0	64.5	49.0	67.5	64.0	86.0
LLaVA-v1.5-13B	51.7	55.5	76.5	33.5	64.4	60.5	62.0	47.5	72.0	61.5	82.0
CogVLM-Chat	51.6	62.5	88.5	30.5	59.1	52.5	57.0	58.0	58.5	59.0	88.5
ShareGPT4V-7B	51.5	62.5	78.5	33.5	64.2	49.5	56.5	43.5	63.0	68.5	74.5
LLaVA-Next-7B	51.1	80.0	28.0	69.0	51.0	62.0	45.0	61.0	77.0	81.5	20.0
LLaVA-v1.5-13B-XTuner	51.1	52.0	81.0	33.5	67.4	53.5	62.0	46.0	72.0	66.0	82.0
LlaVA-InternLM2-7B	50.8	46.0	79.0	27.5	67.6	45.5	66.5	35.5	77.5	73.5	84.5
LLaVA-v1.5-7B-Xtuner	50.2	38.0	75.0	28.0	66.9	46.0	58.5	30.0	62.5	70.5	85.5
SharedCaptioner	49.9	44.0	63.5	37.0	71.0	31.0	46.0	28.0	76.5	69.5	78.5
LLaVA-InternLM-7b	49.7	35.5	81.0	27.5	67.4	33.0	56.0	27.5	70.0	64.5	86.0
LLaVA-v1.5-7B	49.5	62.0	79.0	27.5	62.3	46.5	53.5	40.5	62.0	69.5	74.5
LLaMA-Adapter-v2-7B	40.4	34.0	47.0	36.0	36.2	26.5	21.5	28.0	53.5	26.5	78.0
VisualGLM.6b	38.6	34.0	55.5	26.0	38.4	31.5	28.5	20.5	55.5	28.0	53.0
Frequency	31.7	29.0	36.5	27.5	27.9	26.5	31.0	27.0	54.0	40.5	52.0
Random	28.5	22.5	28.5	28.0	24.6	24.5	23.0	29.5	49.5	38.5	49.0

Table A31. Detail results of 32 LVLMs on Retrieval and Visual Prompt Understanding.

Model	Overall	Retrieval							Visual Prompt Understanding	
		Person ReID	Sketch to Image Retrieval	Face Retrieval	Handwritten Retrieval	Vehicle Retrieval	Image to Image Retrieval	Text to Image Retrieval	Set-of-Marks Recognition	Visual Prompt Understanding
GPT-4o	65.5	37.0	66.0	48.5	25.0	43.0	38.0	77.0	87.9	66.0
InternVL-Chat-V1.2-34B	63.4	60.0	62.0	26.0	70.5	71.0	50.0	54.3	66.5	50.0
QwenVLMax	62.4	35.5	53.0	46.5	26.0	43.5	43.5	44.5	45.7	62.5
Qwen-VL-Plus	62.3	50.5	40.0	24.0	43.0	40.0	42.0	45.7	63.5	71.5
GeminiProVision	61.6	72.5	84.0	92.5	25.5	80.0	81.5	43.0	43.2	58.5
GPT-4V	61.1	38.0	49.5	34.0	25.5	43.0	38.0	60.5	75.4	57.0
LLaVA-Next-34B	60.8	37.0	32.0	28.5	46.5	29.0	40.0	58.8	67.5	47.0
XComposer2-7B	55.7	25.5	31.0	28.5	26.5	43.5	28.0	44.0	56.8	55.5
BLIP2-Flan-T5-XXL	54.8	28.0	31.5	25.0	25.5	25.5	30.0	26.5	27.1	52.5
Yi-VL-34B	54.2	38.0	25.5	27.5	23.0	27.5	20.5	39.7	59.0	63.5
Monkey	53.4	22.5	31.5	22.0	26.0	30.5	26.5	28.5	41.7	45.5
DeepSeek-VL-7B	53.2	32.5	21.5	25.5	32.0	26.0	31.5	44.7	53.0	54.0
Yi-VL-6B	53.2	29.5	24.5	19.0	20.0	25.0	28.5	35.7	50.5	48.0
LLaVA-Next-13B	53.0	35.0	26.0	24.0	30.5	28.0	32.0	53.3	43.5	21.0
TransCore-M	52.7	24.0	36.5	23.0	27.0	31.5	26.5	24.0	51.3	52.5
QWen-VL-Chat	52.5	28.5	31.5	27.0	27.5	25.0	25.0	28.0	38.2	46.5
Claude3V-Haiku	52.2	35.0	34.5	25.5	35.5	34.5	36.5	40.7	46.5	40.0
XComposer	52.1	43.5	60.0	61.0	28.5	60.5	62.0	32.5	41.7	39.5
mPLUG-Owl2	52.0	30.0	35.5	24.5	28.0	23.5	25.5	26.0	34.7	45.5
RBDash-v1-13B	51.8	28.0	31.5	23.0	25.5	21.5	27.0	27.5	49.7	53.0
LLaVA-v1.5-13B	51.7	29.0	28.5	30.0	20.5	25.5	29.0	28.0	42.2	45.5
CogVLM-Chat	51.6	21.0	30.5	20.5	25.5	31.0	24.5	29.5	29.6	50.0
ShareGPT4V-7B	51.5	24.0	33.0	23.0	28.0	29.5	27.0	30.0	46.2	48.0
LLaVA-Next-7B	51.1	31.0	20.5	25.5	33.0	24.5	28.5	37.7	48.0	29.0
LLaVA-v1.5-13B-XTuner	51.1	21.5	32.5	21.5	25.5	31.5	25.5	26.5	42.2	53.5
LlaVA-InternLM2-7B	50.8	27.0	32.5	21.5	25.0	28.0	29.5	30.5	37.7	52.5
LLaVA-v1.5-7B-Xtuner	50.2	22.0	35.5	19.5	25.5	26.0	25.5	28.0	37.7	47.5
SharedCaptioner	49.9	21.0	30.0	25.0	26.5	31.0	27.0	28.5	44.7	46.0
LLaVA-InternLM-7b	49.7	19.0	31.5	22.0	25.5	26.0	27.5	27.0	40.2	51.0
LLaVA-v1.5-7B	49.5	23.0	30.0	20.5	25.5	31.0	26.0	26.5	19.6	47.0
LLaMA-Adapter-v2-7B	40.4	21.5	31.0	29.0	22.0	26.0	21.0	18.0	28.6	39.0
VisualGLM.6b	38.6	22.0	31.0	26.5	25.5	32.0	20.5	24.5	21.6	32.5
Frequency	31.7	30.0	37.0	29.0	28.0	31.0	27.5	26.5	28.1	30.5
Random	28.5	26.0	35.5	17.5	23.5	22.0	22.5	23.0	31.2	31.5

MMT-Bench

Table A32. Detail results of 32 LVLMs on Anomaly Detection and Keypoint Detection.

Model	Overall	Anomaly Detection					Keypoint Detection				
		Industrial Produce Anomaly Detection	Face Mask Anomaly Detection	Helmet Anomaly Detection	Behavior Anomaly Detection	Traffic Anomaly Detection	Furniture Keypoint Detection	Human Keypoint Detection	Clothes Keypoint Detection	Animal Keypoint Detection	Vehicle Keypoint Detection
GPT-4o	65.5	33.5	73.5	84.0	27.0	40.0	52.0	73.5	59.5	48.0	66.3
InternVL-Chat-V1.2-34B	63.4	67.5	87.0	22.5	42.0	66.0	54.0	65.0	43.0	63.0	86.0
QwenVLMax	62.4	72.5	69.0	74.0	47.0	33.5	69.0	48.0	58.0	42.5	43.5
Qwen-VL-Plus	62.3	70.5	74.5	46.5	36.0	66.5	49.0	60.0	41.5	44.6	81.5
GeminiProVision	61.6	57.5	50.5	74.5	20.5	34.0	45.0	61.5	48.5	38.0	54.3
GPT-4V	61.1	63.5	70.5	79.0	25.5	42.5	39.0	52.0	60.5	44.5	52.2
LLaVA-Next-34B	60.8	64.5	79.0	40.5	58.0	64.0	72.0	62.0	42.0	54.3	83.0
XComposer2-7B	55.7	68.5	57.5	67.0	46.5	41.5	46.5	49.0	38.0	31.5	42.4
BLIP2-Flan-T5-XXL	54.8	65.0	29.5	64.0	17.5	42.5	66.5	63.0	74.0	35.5	62.0
Yi-VL-34B	54.2	55.5	86.0	33.0	42.0	45.5	55.5	47.5	34.5	32.6	77.5
Monkey	53.4	43.0	40.0	70.5	27.0	42.5	45.0	26.5	47.5	26.5	35.9
DeepSeek-VL-7B	53.2	33.0	76.5	27.0	42.5	48.0	34.5	27.5	35.0	27.2	81.0
Yi-VL-6B	53.2	55.0	82.5	24.5	40.0	63.0	52.5	50.0	32.5	41.3	82.5
LLaVA-Next-13B	53.0	62.5	79.5	27.0	42.5	36.5	41.5	27.5	27.5	17.4	80.5
TransCore-M	52.7	11.0	56.0	82.0	27.0	42.5	46.0	53.0	31.5	32.0	30.4
QWen-VL-Chat	52.5	29.0	41.0	72.0	27.0	32.0	44.5	31.0	36.0	31.0	37.0
Claude3V-Haiku	52.2	47.0	46.5	22.0	41.0	49.5	62.0	42.0	29.0	52.2	81.5
XComposer	52.1	11.5	50.5	85.5	27.0	42.5	37.0	43.5	49.5	44.5	37.0
mPLUG-Owl2	52.0	69.0	51.5	59.0	27.0	40.5	31.5	36.0	31.0	35.5	30.4
RBDash-v1-13B	51.8	20.5	56.5	82.0	27.0	42.5	37.0	38.5	29.5	29.5	31.5
LLaVA-v1.5-13B	51.7	35.5	58.0	78.5	27.0	42.5	37.5	39.5	28.5	35.0	16.3
CogVLM-Chat	51.6	31.5	30.0	86.0	27.0	42.5	33.5	22.5	28.0	28.5	28.3
ShareGPT4V-7B	51.5	46.0	46.5	79.0	27.0	42.5	32.0	26.5	22.0	28.0	22.8
LLaVA-Next-7B	51.1	40.0	83.5	27.0	42.5	28.5	18.0	18.5	31.0	26.1	79.5
LLaVA-v1.5-13B-XTuner	51.1	53.5	56.5	80.5	27.0	42.5	35.5	32.0	26.5	31.5	18.5
LLaVA-InternLM2-7B	50.8	76.0	57.0	83.5	27.0	42.5	43.0	46.5	31.0	30.5	25.0
LLaVA-v1.5-7B-XTuner	50.2	43.0	55.0	68.5	27.0	42.5	36.5	32.0	24.0	31.0	18.5
SharedCaptioner	49.9	22.0	33.0	72.0	32.5	32.0	45.5	26.0	33.5	32.5	30.4
LLaVA-InternLM-7b	49.7	30.0	55.5	77.5	27.0	42.5	37.5	33.5	35.0	37.0	22.8
LLaVA-v1.5-7B	49.5	59.5	46.0	70.5	27.0	42.5	40.0	25.0	23.0	35.0	34.8
LLaMA-Adapter-v2-7B	40.4	8.0	31.5	44.0	26.0	42.5	49.5	27.5	30.5	31.0	35.9
VisualGLM.6b	38.6	16.0	30.5	57.0	27.0	42.0	43.0	30.0	33.0	33.0	37.0
Frequency	31.7	29.0	29.5	58.5	54.0	51.0	45.5	27.0	27.5	37.0	31.5
Random	28.5	23.0	24.5	44.5	44.5	46.0	42.0	32.0	26.5	34.5	26.1

Table A33. Detail results of 32 LVLMs on Visual Commonsense Reasoning, Visual Code and Image Evaluation.

Model	Overall	Visual Commonsense Reasoning		Visual Code			Image Evaluation	
		Reasoning Whoops	Equation to Latex	Screenshot to Code	Sketch to Code	Image Quality Assessment	Lvlm Response Judgement	
GPT-4o	65.5	82.5	77.0	56.5	40.0	41.5	44.5	
InternVL-Chat-V1.2-34B	63.4	76.0	35.0	31.0	59.5	35.5	26.5	
QwenVLMax	62.4	81.0	76.5	34.5	29.5	46.0	35.5	
Qwen-VL-Plus	62.3	76.5	34.5	24.5	42.0	36.0	27.5	
GeminiProVision	61.6	86.5	75.0	37.0	26.0	41.5	28.5	
GPT-4V	61.1	82.5	65.5	48.5	34.0	41.5	47.5	
LLaVA-Next-34B	60.8	75.0	37.5	30.5	52.5	35.0	9.5	
XComposer2-7B	55.7	83.0	45.0	33.5	30.5	48.5	39.0	
BLIP2-Flan-T5-XXL	54.8	77.0	44.5	49.5	35.5	30.5	29.0	
Yi-VL-34B	54.2	71.5	25.0	23.5	47.5	29.0	22.0	
Monkey	53.4	85.5	50.0	27.0	23.0	23.0	29.0	
DeepSeek-VL-7B	53.2	41.0	29.0	23.5	39.5	31.5	46.0	
Yi-VL-6B	53.2	65.5	31.0	31.0	49.0	29.5	39.0	
LLaVA-Next-13B	53.0	47.0	22.0	12.0	33.5	31.0	5.0	
TransCore-M	52.7	85.5	48.5	19.5	18.5	33.0	35.5	
QWen-VL-Chat	52.5	86.0	47.5	21.0	24.0	31.0	29.0	
Claude3V-Haiku	52.2	70.5	33.0	32.0	37.5	29.5	33.5	
XComposer	52.1	78.0	44.5	21.5	19.0	26.5	31.5	
mPLUG-Owl2	52.0	85.5	35.0	19.5	27.5	34.0	26.0	
RBDash-v1-13B	51.8	78.0	47.5	15.5	13.5	47.5	30.5	
LLaVA-v1.5-13B	51.7	79.0	51.5	16.5	14.5	43.0	28.5	
CogVLM-Chat	51.6	82.0	35.5	27.0	22.5	27.5	28.5	
ShareGPT4V-7B	51.5	83.0	48.5	17.0	24.5	26.0	29.5	
LLaVA-Next-7B	51.1	42.5	23.5	17.5	29.0	32.5	5.0	
LLaVA-v1.5-13B-XTuner	51.1	82.0	46.0	21.5	17.5	47.0	31.5	
LLaVA-InternLM2-7B	50.8	83.0	27.0	18.5	10.0	32.0	36.5	
LLaVA-v1.5-7B-XTuner	50.2	80.5	34.0	16.5	22.0	35.5	29.0	
SharedCaptioner	49.9	82.5	34.0	28.5	23.5	13.5	27.0	
LLaVA-InternLM-7b	49.7	75.5	44.0	22.5	15.5	36.5	29.5	
LLaVA-v1.5-7B	49.5	81.0	31.0	28.5	23.0	25.0	30.5	
LLaMA-Adapter-v2-7B	40.4	71.0	23.5	29.5	24.5	40.0	26.5	
VisualGLM.6b	38.6	65.0	29.0	18.0	23.5	26.0	30.0	
Frequency	31.7	27.0	29.5	30.0	29.0	32.5	27.5	
Random	28.5	28.0	27.5	30.5	28.5	28.5	21.5	

Table A34. Detail results of 32 LVLMs on **Pixel-level Perception** and **Multiple Image Analysis**.

Model	Pixel-level Perception							Multiple Image Analysis	
	Overall	Depth Estimation	Polygon Localization	Interactive Segmentation	Pixel Recognition	Pixel Localization	Image Matting	Spot the Similarity	Spot the Difference
GPT-4o	65.5	52.5	64.5	42.5	77.0	25.0	16.5	58.0	8.5
InternVL-Chat-V1.2-34B	63.4	23.5	46.1	45.0	69.5	25.0	15.5	68.5	97.0
QwenVLMax	62.4	29.0	61.7	51.5	67.5	28.0	15.0	80.0	83.5
Qwen-VL-Plus	62.3	29.5	63.8	51.5	67.5	31.5	15.0	78.0	85.0
GeminiProVision	61.6	29.0	36.0	47.5	77.0	29.5	23.0	50.5	90.0
GPT-4V	61.1	44.5	72.3	32.0	77.5	22.5	16.5	51.0	38.0
LLaVA-Next-34B	60.8	26.0	45.4	57.5	69.0	34.5	19.0	59.0	96.5
XComposer2-7B	55.7	21.5	52.5	51.1	77.0	32.0	29.5	68.5	93.0
BLIP2-Flan-T5-XXL	54.8	63.5	39.5	32.6	71.5	39.0	30.5	45.5	80.0
Yi-VL-34B	54.2	20.0	60.3	39.5	63.5	29.5	17.5	50.5	81.5
Monkey	53.4	21.5	31.5	24.1	75.0	27.0	14.0	54.5	63.0
DeepSeek-VL-7B	53.2	25.5	23.4	61.5	73.5	28.0	18.0	54.5	15.5
Yi-VL-6B	53.2	30.0	35.5	42.0	61.0	25.0	18.5	47.0	57.0
LLaVA-Next-13B	53.0	23.5	27.7	40.0	72.5	24.0	16.5	45.5	27.5
TransCore-M	52.7	50.0	31.5	35.5	71.0	25.0	17.0	45.5	60.0
QWen-VL-Chat	52.5	23.5	29.0	28.4	73.5	25.0	17.0	55.0	43.5
Claude3V-Haiku	52.2	24.0	66.7	8.5	74.0	24.5	15.5	45.5	24.0
XComposer	52.1	35.5	41.0	40.4	70.0	33.5	17.5	54.0	78.5
mPLUG-Owl2	52.0	59.5	37.0	38.3	68.0	30.0	16.5	58.5	51.5
RBDash-v1-13B	51.8	46.5	30.0	37.6	66.5	32.0	17.0	54.5	9.5
LLaVA-v1.5-13B	51.7	51.0	34.0	29.8	69.0	25.0	17.0	45.5	11.5
CogVLM-Chat	51.6	35.5	23.0	22.0	67.5	24.0	27.5	52.5	89.0
ShareGPT4V-7B	51.5	45.5	32.0	34.0	70.0	34.0	20.0	48.0	28.0
LLaVA-Next-7B	51.1	16.0	25.5	33.0	72.0	25.5	23.5	55.5	9.5
LLaVA-v1.5-13B-XTuner	51.1	43.0	42.0	29.8	73.0	25.0	17.0	52.0	22.0
LiaVA-InternLM2-7B	50.8	21.0	37.0	36.9	75.0	33.0	7.5	62.5	49.0
LLaVA-v1.5-7B-XTuner	50.2	38.5	43.0	47.5	74.0	26.0	26.0	54.5	28.0
SharedCaptioner	49.9	24.5	31.5	27.7	75.0	26.0	16.5	64.5	51.0
LLaVA-InternLM-7b	49.7	35.0	49.5	29.8	72.0	27.0	16.5	53.0	61.0
LLaVA-v1.5-7B	49.5	38.5	34.0	26.2	68.0	28.0	21.5	45.5	29.5
LLaMA-Adapter-v2-7B	40.4	32.5	27.5	34.0	44.5	30.0	30.0	45.5	39.0
VisualGLM.6b	38.6	27.0	31.0	32.6	52.5	26.5	24.5	45.0	26.5
Frequency	31.7	30.0	31.5	28.4	26.5	27.0	33.0	50.0	43.0
Random	28.5	22.0	27.5	28.4	25.5	33.0	23.0	57.0	40.0

Table A35. Detail results of 32 LVLMs on **3D** and **Medical Understanding**.

Model	Overall	3D		Medical Understanding				
		3D CAD Recognition	3D Indoor Recognition	Anatomy Identification	Medical Modality Recognition	Other Biological Attributes	Disease Diagnose	Lesion Grading
GPT-4o	65.5	61.0	38.0	74.0	99.0	67.0	77.0	74.5
InternVL-Chat-V1.2-34B	63.4	56.0	35.0	54.5	93.5	70.0	89.5	51.5
QwenVLMax	62.4	54.0	38.0	67.0	94.0	59.0	81.5	69.0
Qwen-VL-Plus	62.3	52.5	40.5	66.5	95.0	58.0	82.0	65.0
GeminiProVision	61.6	61.0	29.5	86.5	96.5	80.0	74.5	75.5
GPT-4V	61.1	41.5	32.0	60.5	96.5	60.0	71.0	60.0
LLaVA-Next-34B	60.8	52.0	45.0	41.0	67.0	65.0	85.5	51.0
XComposer2-7B	55.7	41.0	40.0	42.5	55.0	66.5	45.5	58.0
BLIP2-Flan-T5-XXL	54.8	55.0	40.5	51.0	72.5	64.5	61.5	51.0
Yi-VL-34B	54.2	46.0	31.5	48.5	49.5	65.0	73.5	61.5
Monkey	53.4	56.5	36.5	51.5	97.0	63.5	73.0	55.0
DeepSeek-VL-7B	53.2	52.5	42.0	51.0	96.0	71.5	79.5	50.5
Yi-VL-6B	53.2	42.0	45.0	54.0	71.0	59.5	62.0	56.5
LLaVA-Next-13B	53.0	50.5	49.5	42.0	54.0	76.0	58.5	48.5
TransCore-M	52.7	44.0	46.0	40.5	73.0	71.0	72.5	48.5
QWen-VL-Chat	52.5	55.0	35.0	59.0	93.5	56.5	70.0	56.5
Claude3V-Haiku	52.2	45.0	31.5	53.0	67.5	60.0	76.0	57.0
XComposer	52.1	52.0	42.0	49.0	72.0	72.5	41.5	43.0
mPLUG-Owl2	52.0	49.0	41.5	56.0	97.5	58.5	55.5	49.5
RBDash-v1-13B	51.8	39.5	42.0	49.0	80.0	65.0	59.5	43.0
LLaVA-v1.5-13B	51.7	47.5	32.0	41.0	72.5	67.0	64.5	48.0
CogVLM-Chat	51.6	55.0	29.5	50.0	94.0	66.5	72.5	46.5
ShareGPT4V-7B	51.5	43.5	47.0	46.0	82.0	65.5	61.0	49.5
LLaVA-Next-7B	51.1	44.5	42.0	47.0	76.0	60.0	50.5	44.0
LLaVA-v1.5-13B-XTuner	51.1	46.0	49.0	34.0	50.0	66.0	61.0	34.5
LiaVA-InternLM2-7B	50.8	50.0	49.0	42.0	65.5	60.0	77.0	44.5
LLaVA-v1.5-7B-XTuner	50.2	46.0	45.0	37.5	69.0	49.0	59.0	48.0
SharedCaptioner	49.9	45.0	43.5	55.0	72.0	73.0	50.5	46.5
LLaVA-InternLM-7b	49.7	54.0	30.0	37.5	63.5	59.0	51.5	48.5
LLaVA-v1.5-7B	49.5	46.0	43.5	34.5	79.5	60.0	63.0	47.0
LLaMA-Adapter-v2-7B	40.4	38.5	29.0	45.0	55.5	51.5	64.0	44.0
VisualGLM.6b	38.6	44.0	29.5	24.0	57.5	47.5	49.5	41.5
Frequency	31.7	26.0	27.0	27.5	30.0	40.5	26.5	30.0
Random	28.5	24.0	27.0	25.0	24.0	38.0	24.0	28.0

Table A36. Detail results of 32 LVLMS on Cross Image Matching and Visual Summary (part 1).

Model	Overall	Cross Image Matching			Visual Summary				
		One-shot Detection	Point Tracking	Single Object Tracking	Video Captioning	Image Captioning Paragraph	Image Captioning	Instance Captioning	Image Dense Captioning
GPT-4o	65.5	54.5	34.5	58.5	95.5	99.5	98.5	85.0	65.0
InternVL-Chat-V1.2-34B	63.4	59.0	58.5	53.0	69.5	99.0	96.5	90.0	52.8
QwenVLMax	62.4	73.5	64.0	44.5	91.5	99.0	98.0	88.5	61.9
Qwen-VL-Plus	62.3	75.0	63.0	46.0	92.5	99.0	98.0	87.0	62.9
GeminiProVision	61.6	42.5	14.0	43.5	89.0	99.5	97.5	85.0	42.1
GPT-4V	61.1	71.5	63.5	67.5	91.5	97.5	98.0	77.0	60.9
LLaVA-Next-34B	60.8	62.0	12.0	55.0	87.0	99.0	98.0	89.5	61.9
XComposer2-7B	55.7	49.5	75.5	58.5	49.0	99.0	97.5	80.0	52.8
BLIP2-Flan-T5-XXL	54.8	85.0	71.5	62.5	89.0	96.0	96.0	66.0	52.8
Yi-VL-34B	54.2	60.0	32.5	51.5	44.5	98.5	94.5	78.0	49.7
Monkey	53.4	72.0	63.5	49.5	81.0	71.0	95.5	70.0	33.5
DeepSeek-VL-7B	53.2	72.5	72.5	56.5	42.0	94.0	97.5	77.5	32.0
Yi-VL-6B	53.2	43.0	45.5	41.5	36.5	93.5	91.5	65.5	40.1
LLaVA-Next-13B	53.0	63.0	69.5	47.0	80.0	96.0	95.0	79.5	25.9
TransCore-M	52.7	66.0	80.0	51.5	91.0	96.5	95.0	82.5	26.9
QWen-VL-Chat	52.5	68.5	61.5	45.0	85.5	88.5	92.5	73.0	28.9
Claude3V-Haiku	52.2	60.5	52.0	45.0	53.0	95.5	92.5	62.0	50.8
XComposer	52.1	32.0	80.5	44.5	87.5	87.5	92.0	59.5	43.1
mPLUG-Owl2	52.0	54.5	67.5	51.0	85.5	77.5	92.0	69.5	39.1
RBDash-v1-13B	51.8	67.0	75.0	50.5	60.0	98.5	97.5	77.5	45.7
LLaVA-v1.5-13B	51.7	62.0	73.5	52.0	51.5	95.5	97.5	78.5	25.9
CogVLM-Chat	51.6	60.0	40.5	37.0	61.5	91.5	95.5	72.0	18.3
ShareGPT4V-7B	51.5	61.5	80.5	51.0	86.0	90.0	96.0	76.5	31.5
LLaVA-Next-7B	51.1	58.0	79.5	66.0	82.5	82.5	96.0	76.5	22.3
LLaVA-v1.5-13B-XTuner	51.1	60.0	58.0	52.5	57.5	95.0	95.0	76.5	32.5
LlaVA-InternLM2-7B	50.8	44.0	71.5	59.0	39.0	93.0	94.0	77.5	24.9
LLaVA-v1.5-7B-Xtuner	50.2	58.5	72.5	58.5	47.5	90.0	95.5	77.0	35.0
SharedCaptioner	49.9	34.5	78.5	57.5	58.0	65.5	92.0	59.5	28.9
LLaVA-InternLM-7b	49.7	50.0	81.5	47.5	48.5	86.0	91.5	71.5	39.1
LLaVA-v1.5-7B	49.5	63.5	72.5	51.0	84.5	73.0	96.5	69.0	23.9
LLaMA-Adapter-v2-7B	40.4	39.0	33.5	35.0	23.0	29.0	71.0	34.5	31.0
VisualGLM.6b	38.6	49.5	51.5	43.5	28.5	45.0	56.0	36.0	28.9
Frequency	31.7	27.0	30.5	28.0	27.5	27.0	31.5	27.0	28.4
Random	28.5	22.5	25.0	33.0	24.0	22.5	26.5	26.5	25.4

Table A37. Detail results of 32 LVLMS on Visual Summary (part 2) and Autonomous Driving.

Model	Overall	Visual Summary			Autonomous Driving				
		Multiple Instance Captioning	Multiple Image Captioning	Writing Poetry from Image	Traffic Participants Understanding	Multiple-view Image Understanding	Traffic Sign Understanding	Temporal Sequence Understanding	Traffic Light Understanding
GPT-4o	65.5	96.0	88.5	63.0	51.0	18.0	74.5	44.0	86.0
InternVL-Chat-V1.2-34B	63.4	91.0	90.0	70.0	57.5	23.0	67.5	52.5	88.5
QwenVLMax	62.4	91.0	90.0	70.0	59.5	23.0	66.5	45.5	71.0
Qwen-VL-Plus	62.3	92.0	90.5	70.5	60.0	26.0	76.5	50.0	85.0
GeminiProVision	61.6	88.0	79.0	61.0	62.5	51.0	74.5	53.0	56.7
GPT-4V	61.1	92.5	85.5	64.5	51.5	27.5	71.0	46.0	55.5
LLaVA-Next-34B	60.8	87.5	92.0	72.0	57.0	24.0	71.5	51.0	72.0
XComposer2-7B	55.7	92.0	77.0	42.0	57.5	15.0	70.5	49.5	51.5
BLIP2-Flan-T5-XXL	54.8	89.5	78.0	62.0	49.5	15.5	62.5	43.5	51.8
Yi-VL-34B	54.2	82.5	83.5	61.0	47.5	23.0	61.0	48.5	61.0
Monkey	53.4	90.0	66.5	43.5	44.0	23.0	70.0	34.0	47.1
DeepSeek-VL-7B	53.2	87.0	73.0	50.0	50.0	18.5	62.5	37.5	75.5
Yi-VL-6B	53.2	79.0	68.0	33.0	49.0	26.0	55.5	39.0	65.0
LLaVA-Next-13B	53.0	86.5	85.5	52.0	53.0	21.5	70.5	48.0	85.5
TransCore-M	52.7	83.5	75.0	54.5	53.0	17.0	62.0	50.5	50.2
QWen-VL-Chat	52.5	91.5	75.5	48.5	43.5	36.0	66.0	34.0	47.5
Claude3V-Haiku	52.2	65.5	65.0	56.5	53.0	0.0	66.0	0.0	51.5
XComposer	52.1	86.0	71.0	50.5	45.5	24.0	59.5	25.5	49.4
mPLUG-Owl2	52.0	82.5	41.5	45.0	51.0	18.5	67.5	44.0	46.3
RBDash-v1-13B	51.8	85.5	78.5	60.5	50.0	22.5	62.5	48.5	47.8
LLaVA-v1.5-13B	51.7	87.5	70.5	56.0	51.5	25.5	62.0	47.0	46.9
CogVLM-Chat	51.6	89.5	88.0	42.0	46.5	32.5	68.0	30.5	47.2
ShareGPT4V-7B	51.5	80.5	61.0	51.0	54.5	28.0	64.0	37.0	47.2
LLaVA-Next-7B	51.1	79.0	73.5	46.0	55.0	18.5	66.5	49.0	78.5
LLaVA-v1.5-13B-XTuner	51.1	82.0	64.5	47.0	53.0	20.5	59.5	43.0	45.9
LlaVA-InternLM2-7B	50.8	84.0	67.0	52.0	53.0	12.5	64.0	48.0	48.0
LLaVA-v1.5-7B-Xtuner	50.2	83.5	51.0	51.5	52.5	11.0	62.0	46.0	45.3
SharedCaptioner	49.9	79.0	65.5	47.0	54.0	18.5	60.5	46.0	44.6
LLaVA-InternLM-7b	49.7	81.0	61.5	40.5	51.5	9.5	58.0	46.5	45.7
LLaVA-v1.5-7B	49.5	83.5	66.5	48.0	54.0	15.5	64.5	45.5	45.9
LLaMA-Adapter-v2-7B	40.4	33.0	28.5	28.5	16.0	26.5	43.0	27.5	32.6
VisualGLM.6b	38.6	39.5	60.0	30.0	28.0	25.0	25.0	20.0	32.3
Frequency	31.7	29.5	29.0	26.0	27.0	28.0	31.0	30.0	32.5
Random	28.5	25.5	24.0	25.5	25.0	27.0	29.0	23.0	29.8