

MANBench: Is Your Multimodal Model Smarter than Human?

Han Zhou¹, Qitong Xu², Yiheng Dong¹, Xin Yang^{1*}

¹School of Electronic Information and Communications, Huazhong University of Science and Technology

²School of Basic Medicine, Huazhong University of Science and Technology
{micdz, esoox, yiheng, xinyang2014}@hust.edu.cn

Abstract

The rapid advancement of Multimodal Large Language Models (MLLMs) has ignited discussions regarding their potential to surpass human performance in multimodal tasks. In response, we introduce **MANBench** (Multimodal Ability Norms Benchmark), a bilingual benchmark (English and Chinese) comprising 1,314 questions across nine tasks, spanning knowledge-based and non-knowledge-based domains. MANBench emphasizes intuitive reasoning, seamless cross-modal integration, and real-world complexity, providing a rigorous evaluation framework.

Through extensive human experiments involving diverse participants, we compared human performance against state-of-the-art MLLMs. The results indicate that while MLLMs excel in tasks like Knowledge and Text-Image Understanding, they struggle with deeper cross-modal reasoning tasks such as Transmorphic Understanding, Image Consistency, and Multi-image Understanding. Moreover, both humans and MLLMs face challenges in highly complex tasks like Puzzles and Spatial Imagination.

MANBench highlights the strengths and limitations of MLLMs, revealing that even advanced models fall short of achieving human-level performance across many domains. We hope MANBench will inspire efforts to bridge the gap between MLLMs and human multimodal capabilities. The code and dataset are available at <https://github.com/micdz/MANBench>.

1 Introduction

Recent advancements in Large Language Models (LLMs) have demonstrated exceptional capabilities in text comprehension, reasoning, and generation (Liu et al., 2024a; Dubey et al., 2024; Jaech et al., 2024). Building on this foundation, Multimodal Large Language Models (MLLMs) have

emerged, significantly enhancing the ability to process and generate multimodal content (OpenAI, 2024b; Team, 2025; Anthropic, 2024; Team et al., 2023). These models have achieved substantial success in tasks such as image captioning, visual question answering, and visual dialog, showcasing their potential for multimodal understanding and generation.

As artificial intelligence continues to evolve rapidly, a critical question arises: Can the general multimodal capabilities of MLLMs eventually surpass those of humans? While humans possess sophisticated multimodal abilities, existing benchmarks for evaluating these capabilities exhibit several significant limitations. First, many benchmarks fail to comprehensively assess the full spectrum of multimodal reasoning, often focusing on narrow tasks or requiring domain-specific knowledge (Lu et al., 2024; Chen et al., 2024b; Zhang et al., 2024b; Wang et al., 2024c). Second, human evaluations in these benchmarks are frequently inadequate, relying on small sample sizes and limited diversity. These shortcomings hinder robust conclusions about human performance (Zhang et al., 2024b; Wang et al., 2024c; Yue et al., 2024; Russakovsky et al., 2014; Radford et al., 2021).

To address these issues, we introduce **MANBench** (Multimodal Ability Norms Benchmark), a comprehensive benchmark designed to evaluate the multimodal capabilities of both humans and MLLMs. MANBench comprises nine tasks, encompassing 1,314 questions and 2,231 images. It aims to provide a fair and rigorous assessment framework, ensuring equitable comparisons between human and MLLM performance. Specifically, our benchmark addresses the following key limitations of existing benchmarks: (1) it separates questions requiring prior knowledge from those that do not, (2) it ensures that all questions demand reasoning rather than simple retrieval, (3) it mandates the integration of textual and visual informa-

*Corresponding author.

tion, and (4) it includes a wide range of reasoning difficulties to challenge both humans and MLLMs.

We conducted a large-scale human evaluation using the MANBench dataset, systematically comparing the performance of humans and MLLMs across a variety of multimodal tasks. Twelve MLLMs, including state-of-the-art models such as GPT-4o (OpenAI, 2024b), Qwen2.5-VL-72B-Instruct (Team, 2025), and Gemini-1.5 Pro (Team et al., 2023), were evaluated on MANBench. Additionally, we recruited 575 participants from diverse backgrounds to complete the tasks. Our results reveal that while MLLMs excel in tasks such as Knowledge and Text-Image Understanding, they struggle with tasks requiring deeper cross-modal reasoning, including Transmorphic Understanding, Image Consistency, and Multi-image Understanding. Notably, both humans and MLLMs encounter challenges in highly complex tasks such as Puzzles and Spatial Imagination.

In summary, our contributions are as follows: (1) We introduce MANBench, a bilingual benchmark with 1,314 questions across nine tasks. (2) We conduct extensive human experiments with 575 diverse participants. (3) We evaluate 12 state-of-the-art MLLMs. (4) We provide insights into the strengths and limitations of MLLMs, highlighting the challenges in bridging the gap between MLLMs and human multimodal abilities.

2 Related Work

2.1 Multimodal Large Language Models

As LLMs advance, integrating visual content has become a critical step in the development of interactive intelligent assistants. MLLMs focus on aligning vision and language modalities. CLIP (Radford et al., 2021) is foundational in this field, using contrastive learning on image-text pairs. Approaches like MiniGPT-4 (Zhu et al., 2023) and LLaVA (Liu et al., 2024b,c), connect vision encoders and LLMs using techniques like Q-Former and MLP, leveraging frozen language models with a limited number of trainable parameters, demonstrating promising results. Furthermore, a range of MLLMs trained on vast datasets are now accessible to the public, pushing the boundaries of visual comprehension and generalization, including InternVL1.5, InternVL2.5 series (Chen et al., 2024c) and Qwen-VL (Wang et al., 2024e).

2.2 Evaluations of MLLMs

Reasoning ability is a fundamental indicator of intelligence. Numerous benchmarks have been developed to evaluate reasoning in text-only modalities, including MathQA (Amini et al., 2019), GSM1k (Zhang et al., 2024a), and LiveBench (White et al., 2024). These benchmarks present questions that require advanced reasoning and knowledge, making them difficult for most people to solve.

To comprehensively assess the multimodal reasoning capabilities of both humans and MLLMs, we propose the development of a novel text-image benchmark that addresses certain limitations in existing evaluation frameworks. As illustrated in Figure 1, current benchmarks exhibit four primary areas for improvement.

Limitation 1: Overreliance on Prior Knowledge in Certain Samples Existing benchmarks frequently emphasize prior knowledge over reasoning capability (Lu et al., 2024; Chen et al., 2024b; Zhang et al., 2024b; Wang et al., 2024c), leading to variability in evaluating human performance. To address this issue, MANBench introduces a clear distinction between Knowledge-based and non-knowledge-based tasks, facilitating a more balanced evaluation of cognitive processes.

Limitation 2: Potential for Knowledge Retrieval Without Reasoning Current benchmarks sometimes enable solutions through direct knowledge retrieval (Chen et al., 2024a), which limits their ability to effectively evaluate reasoning skills. Our proposed dataset mitigates this with carefully constructed questions that explicitly require reasoning processes, while clearly distinguishing knowledge-based tasks.

Limitation 3: Insufficient Integration of Multimodal Information Some existing benchmarks fail to effectively utilize the integration potential of textual and visual data (Chen et al., 2024a). Our framework addresses this by introducing tasks that necessitate the concurrent analysis of both modalities, providing a more comprehensive assessment of multimodal reasoning capabilities.

Limitation 4: Narrow Difficulty Spectrum Current benchmarks often exhibit a limited range of difficulty levels, potentially underestimating the abilities of both humans and models (Fu et al., 2024; Meng et al., 2024). To mitigate this, our dataset introduces a meticulously balanced range of reasoning complexities, enabling a more thor-

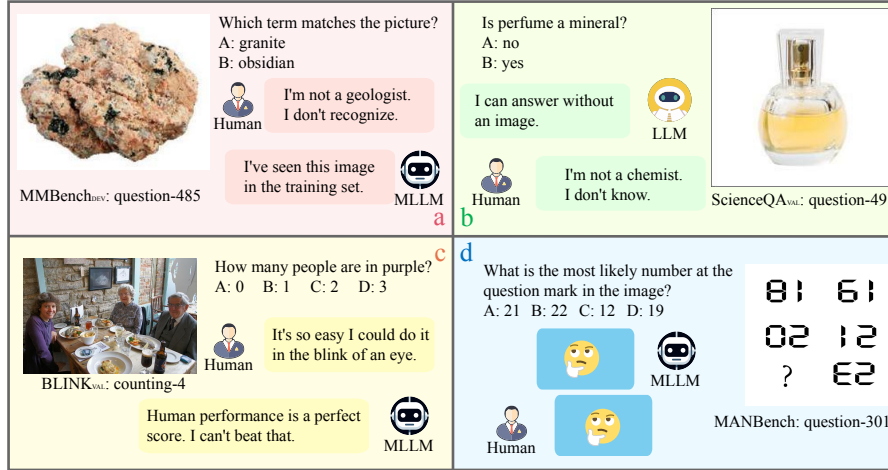


Figure 1: Limitations of existing text-image benchmarks. (a) Some samples necessitate prior knowledge rather than reasoning capabilities (Liu et al., 2024d). (b) Certain samples can be answered accurately without utilizing the image content (Lu et al., 2022). (c) Other samples lack sufficient reasoning complexity (Fu et al., 2024). (d) In contrast, our proposed questions require comprehensive reasoning, effective integration of textual and visual information, and a diverse range of reasoning difficulties.

ough evaluation across varying levels of difficulty.

2.3 Evaluations of Humans

Several researchers have conducted experiments to evaluate human multimodal capabilities. In the early stages of research, studies primarily focused on human image processing abilities, such as image classification (Russakovsky et al., 2014) and few-shot classification (Radford et al., 2021), rather than the integrated capabilities between images and text.

Recently, there has been a surge in studies assessing human multimodal capabilities within the domain of Visual Question Answering (VQA). These studies explore specialized areas such as MathVerse (Zhang et al., 2024b), MATH-Vision (Wang et al., 2024c), and MMMU (Yue et al., 2024). However, due to their specialized nature, it is difficult to generalize these tests to the broader population when evaluating overall human multimodal capabilities. The reliance on domain-specific knowledge implies that most questions can only be effectively answered by experts in the respective fields.

A critical limitation of both early and recent studies is the relatively small sample sizes, which may restrict the representativeness of their findings. To ensure the reliability and robustness of conclusions regarding human capabilities, studies typically require larger and more diverse participant pools to account for variability and ensure findings are broadly representative.

3 The Multimodal Ability Norms Benchmark

Our goal is to develop a benchmark that accurately evaluates the multimodal capabilities of both humans and MLLMs.

3.1 Overview of MANBench

As illustrated in Figure 3, MANBench comprises nine tasks, each containing more than 110 questions, resulting in a total of 1,314 questions and 2,231 images. All questions in the dataset are bilingual, with text provided in both English and Chinese. The translations have been meticulously reviewed by the coauthors to ensure semantic equivalence across languages.

The questions are categorized into two types: knowledge-based and non-knowledge-based. Knowledge-based questions require participants to possess prior knowledge in subjects such as Biology, Engineering, Geography, Physics, and Chemistry. In contrast, non-knowledge-based questions do not demand advanced domain-specific knowledge.

3.2 Key Features of MANBench

Intuitive Reasoning without Domain Bias: The benchmark is designed to eliminate reliance on prior knowledge by emphasizing universal perception and logical inference, accessible through immediate visual-textual analysis. The tasks are care-

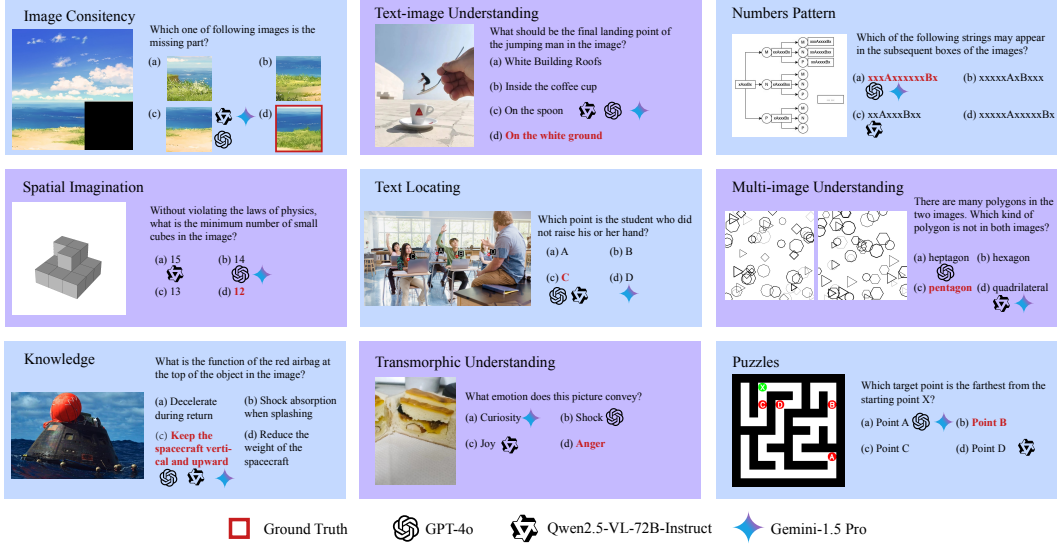


Figure 2: Qualitative results on MANBench. For each task, we display the selections made by GPT-4o, Qwen2.5-VL-72B-Instruct, and Gemini-1.5 Pro, with the red-highlighted choice indicating the ground truth.

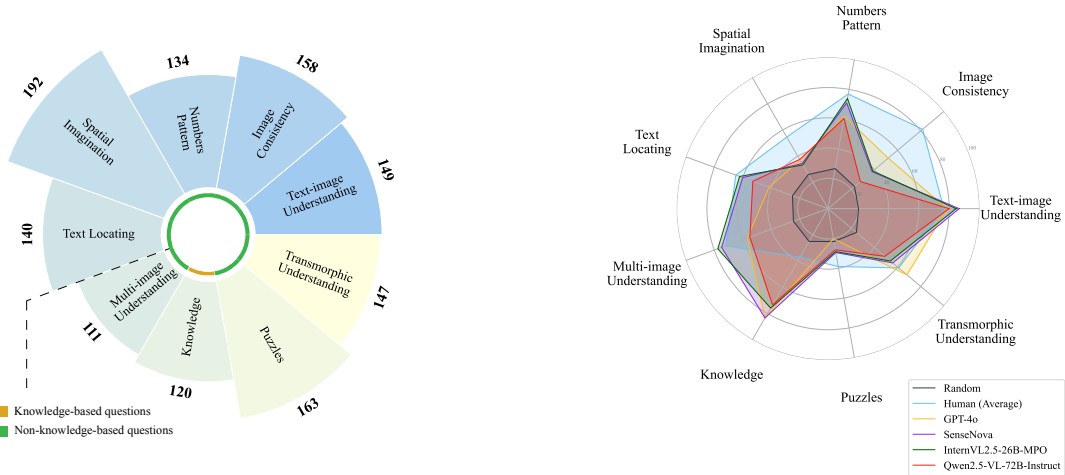


Figure 3: Statistics of MANBench. The benchmark comprises 9 tasks, categorized into knowledge-based and non-knowledge-based questions.

Figure 4: Performance comparison on the MANBench English subset among human average accuracy and some selected MLLMs. Please refer to Table 1 for more results.

fully constructed to avoid domain-specific knowledge, ensuring ease of use and fairness in benchmarking.

Seamless Cross-Modal Integration: Questions necessitate joint interpretation of images and text, demanding causal or functional relationships between modalities rather than superficial keyword-to-image alignment. This design compels subjects to synthesize contextual meaning from both inputs, such as inferring intent from visual details paired with textual hints.

Clarity-Focused Task Design: The tasks prioritize clear and unambiguous question stems with minimal distractors, ensuring that the challenges

arise from reasoning difficulty rather than linguistic ambiguity or misleading cues.

Robustness to Real-World Complexity: By simulating high-density, noisy environments and requiring multi-step reasoning, the benchmark discourages shortcut solutions. This ensures that models demonstrate practical applicability in solving complex, real-world scenarios.

3.3 Dataset Collection Process

MANBench evaluates multimodal capabilities across nine tasks, all of which utilize a multiple-choice question format. The dataset is constructed

from diverse sources, including generated images, real-world images, social media images, and images from existing benchmark datasets.

Spatial Imagination: This task systematically evaluates the spatial cognition capabilities of MLLMs and humans through a multidimensional assessment framework. By utilizing generated controllable scene images, real-world images, and 3D IQ test images, we design a variety of question types, including 2D-to-3D spatial transformations, the disassembly and reassembly of three-dimensional structures, spatial distance measurements, and object quantity estimations. Through the simulation of real-world scenarios and abstract geometric problems, the task comprehensively examines subjects' mastery of core competencies, such as spatial topological relationships and geometric transformation principles.

Numbers Pattern: This task is designed to comprehensively evaluate the numerical reasoning capabilities of MLLMs and humans. By integrating generated original images with real-world photographic materials, we design question types that span precise object counting, dynamic quantity estimation, and complex numerical pattern deduction. These multidimensional questions systematically assess the ability of models to process numerical information, ranging from foundational concepts to advanced reasoning.

Puzzles: This task aims to evaluate the inductive and reasoning abilities of MLLMs and humans in abstract graphic patterns. The data sources for this task are diverse and unique: On the one hand, it utilizes visual classic IQ test (Fu et al., 2024), presenting problems through basic shape examples and four candidate images to assess subjects' logical reasoning capabilities based on shape attributes, positional relationships, and sequential patterns. On the other hand, we develop maze schematics to evaluate subjects' ability to estimate path distances in planar mazes. As a type of problem rarely addressed in existing datasets, this task provides a novel means of assessing the ability to extract and reason about information from complex graphics.

Multi-image Understanding: This task aims to assess the ability of MLLMs and humans to synthesize and interpret information from multiple images. We collect a diverse set of images from the Internet and MUIRBench (Wang et al., 2024a), as well as generate additional images, carefully constructing a series of questions that require models to accurately identify commonalities among mul-

tle images. These questions require the model not only to comprehend the content of each image individually but also to compare and reason across multiple images to discern shared features or patterns.

Text-image Understanding: This task evaluates the ability of MLLMs to comprehensively integrate and understand both image and text information. Unlike traditional visual quizzes that rely on superficial descriptions, this task emphasizes three specific types of questions: phenomenon attribution, functional reasoning about objects, and topic generalization. The dataset for this task is constructed by integrating resources from MMBench (Liu et al., 2024d), MMStar (Chen et al., 2024a), and social media.

Image Consistency: The task is designed to assess the ability of MLLMs and humans to understand the consistency of an image locally and holistically. By integrating existing benchmarks, social media, and generated visuals, we implement a non-uniform grid partitioning strategy. Specifically, fixed-size cropping is applied to the bottom-right corner of source images, while the remaining area is retained as a reference. Subjects are required to identify the visually consistent completion unit from a set of four candidate images. Compared to the jigsaw task of BLINK (Fu et al., 2024), our design improves evaluation precision for cross-scale visual-semantic alignment and detail-oriented reasoning by expanding candidate options and incorporating more complex visual patterns.

Transmorphic Understanding: The aim of this task is to evaluate the generalization ability of MLLMs and humans to trigger cross-modal abstract associations from visual inputs and overcome the limitations of figurative semantic matching. For this purpose, we collect images from social media and require subjects to infer underlying human emotions based on individual images or extrapolate from the visual characteristics of objects to implicitly associated scenes.

Text Locating: This task aims to evaluate the ability of subjects to accurately locate targets in images based on textual descriptions, which is crucial for multimodal interaction and image-text fusion applications. To construct the task, we collect images from social media and annotate them with multiple markers, accompanied by a textual description. These markers follow the design principles of BLINK (Fu et al., 2024). Then, subjects are required to select the target point based on the de-

scription. Unlike traditional tests that focus primarily on basic image-text matching, this task emphasizes a deeper comprehension of textual information and efficient retrieval capabilities in complex scenarios, thereby offering more comprehensive and precise evaluation results.

Knowledge: This task is designed to assess the knowledge base, conceptual associations, and cross-modal knowledge transfer in humans and MLLMs. With data from Wikipedia, Oxford-Pets (Parkhi et al., 2012), ScienceQA (Lu et al., 2022) and MMStar (Chen et al., 2024a), we generate a series of questions that require subjects to recognise images, reason about them according to their knowledge system, and select the matching textual options.

4 Experiments

We conducted extensive experiments involving 575 human participants and 12 MLLMs to evaluate their multimodal abilities.

4.1 Experimental Setup

4.1.1 Human Experiments

The human experiments were conducted in two stages.

In the first stage, a preliminary pilot study was conducted to evaluate the quality of the questions. Five native Chinese-speaking university students were recruited to assess the questions. They completed all tasks in the dataset and provided feedback on clarity, ambiguity, and difficulty. Their response times and accuracy were analyzed to determine the difficulty level of the questions.

In the second stage of our study, we expanded our participant pool to include a larger and more diverse group. Given the impracticality of requiring participants to complete the entire dataset of 1,314 questions in a large-scale study—which would take approximately 10 hours—we divided the dataset into ten subsets, each containing around 131 questions. This division was guided by a difficulty gradient established during the pilot study. However, the limited number of participants in the pilot study may have introduced some imperfections in the gradient. Despite this limitation, the division facilitated a more manageable and scalable evaluation process for our broader study.

We then recruit 575 Chinese participants from diverse backgrounds to engage in the experiments. Additionally, we gather demographic data from

the participants, including age, gender, and education level, to enhance the analysis. As shown in Figure 5, the participants are distributed across a wide range of ages. For more participants statistics, please refer to Appendix A.2.




Figure 5: Age and gender distribution of participants.

We meticulously designed the user interface for the human experiments to ensure clarity and consistency, allowing participants to clearly comprehend the questions, provide responses, and receive the same information as the MLLMs. Participants were permitted to omit perceived as overly challenging or ambiguous. For further details regarding the experimental setup, methodology, and additional resources, please refer to Appendix A.

4.2 MLLM Experiments

We adopt standard experimental setups following the VLMEvalKit (Duan et al., 2024), where the temperature is fixed at 0 and the retry count is set to 10. Please refer to Appendix B for more details on our computing infrastructure.

The MANBench benchmark is evaluated on 14 recent MLLMs. Specifically, we select 6 closed-source models: GPT-4o (OpenAI, 2024b), GPT-o1 (OpenAI, 2024a), SenseNova (SenseTime, 2024), Step-1o (StepFun, 2025), Gemini-1.5-Pro (Team et al., 2023), Claude-3.5-Sonnet (Anthropic, 2024). Additionally, we evaluate 8 open-source models: Deepseek-VL2 (Wu et al., 2024), Qwen2-VL-72B-Instruct (Wang et al., 2024d), Qwen2.5-VL-72B (Team, 2025), QVQ-72B-Preview (Team, 2024), InternVL2-26B, InternVL2-8B (Chen et al., 2024c), InternVL2.5-26B-MPO, InternVL2.5-78B-MPO (Wang et al., 2024f).

	Overall	Text-image Understanding (149)	Numbers Pattern (134)	Image Consistency (158)	Spatial Imagination (192)	Text Locating (140)	Multi-image Understanding (111)	Knowledge (120)	Puzzles (163)	Transmorphic Understanding (147)
Human (Average)	62.26	76.46	76.88	81.55	54.56	64.82	71.69	37.05	38.83	60.90
Human (Best)	90.87	89.00	91.79	100.00	94.57	94.39	92.86	74.75	86.42	90.54
Random	24.05	20.13	26.87	22.78	26.04	25.00	24.32	25.00	22.09	24.49
Open-source MLLMs										
Deepseek-VL2	45.43	<u>76.51</u>	57.46	31.01	31.77	44.29	36.04	<u>70.83</u>	26.38	44.90
InternVL2-26B	41.86	61.74	54.48	24.68	29.69	33.57	43.24	<u>67.50</u>	27.61	46.26
InternVL2-8B	36.15	59.06	36.57	27.22	27.08	31.43	31.53	<u>50.83</u>	28.22	38.78
InternVL2.5-26B-MPO	56.32	84.56	73.88	38.61	33.33	62.14	77.48	<u>75.83</u>	28.83	53.74
InternVL2.5-78B-MPO	59.82	84.56	70.90	45.57	38.02	62.86	<u>75.68</u>	<u>80.00</u>	38.65	60.54
QVQ-72B-Preview	50.00	<u>76.51</u>	58.21	31.01	31.77	42.86	63.96	82.50	34.36	46.94
Qwen2-VL-72B-Instruct	46.19	<u>77.85</u>	58.96	28.48	32.29	38.57	48.65	<u>79.17</u>	23.31	43.54
Qwen2.5-VL-72B-Instruct	49.92	<u>79.87</u>	60.45	27.85	37.50	52.86	54.95	<u>73.33</u>	27.61	48.98
Closed-source MLLMs										
Claude-3.5-Sonnet	54.87	<u>81.88</u>	57.46	46.20	36.98	50.00	69.37	<u>78.33</u>	27.61	<u>62.59</u>
GPT-4o	53.81	<u>78.52</u>	61.94	51.90	36.98	41.43	56.76	<u>83.33</u>	20.25	<u>68.03</u>
GPT-o1	59.97	88.59	73.88	53.16	34.90	54.29	64.86	85.83	33.74	<u>68.03</u>
Gemini-1.5-Pro	55.10	<u>86.58</u>	66.42	33.54	41.67	39.29	69.37	<u>83.33</u>	34.36	57.82
SenseNova	56.85	<u>86.58</u>	70.90	37.97	34.38	60.00	74.77	<u>83.33</u>	29.45	55.78
Step-1o	54.79	<u>83.89</u>	68.66	39.87	30.73	44.29	60.36	<u>78.33</u>	33.74	70.07

Table 1: Results of different models on the MANBench English subset. The first row shows task names and number of test data. The best performance in each task is in bold. Values exceeding human average performance are underlined. For detailed results of Chinese subset, please refer to Table 5. For the information of the best human performance, please refer to Appendix A.5. The overall score is calculated by weighting the task scores according to the number of questions in each task.

4.3 Main Results

4.3.1 MLLMs Performance

The overall performance of MLLMs on MANBench is shown in Table 1. The best MLLMs performance is achieved by the closed-source model GPT-o1 with the average score of 59.97, followed by the open-source model InternVL2.5-78B-MPO with a score of 59.82.

Especially on the Puzzles task, the accuracy of MLLMs is approximately 30%, which is comparable to random guessing. This result highlights the significant difficulty of the Puzzles task for MLLMs and suggests that there remains substantial room for improvement in this type of image reasoning task.

4.3.2 Human Performance

We evaluate the performance of MLLMs by comparing their results to human performance benchmarks. As detailed in Table 1, two key metrics are employed: average human performance and best human performance. The average human performance serves as the baseline for assessing MLLM capabilities. To compute this baseline, the accuracy for each task C_i is determined as follows: For every question j within a task, the accuracy p_{ij} is calculated. The task accuracy P_i is then derived by averaging the accuracies of all questions within the

task:

$$P_i = \frac{1}{n_i} \sum_{j=1}^{n_i} p_{ij}$$

where n_i represents the total number of questions in task C_i .

In Figure 6, we observe that children under the age of fourteen perform comparably to adults on MANBench, indicating that our dataset is not highly sensitive to age differences. Nevertheless, our findings also reveal a decline in human performance as age increases.




Figure 6: Performance comparison across different age groups.




Figure 7: Performance of GPT-4o with and without images on MANBench English subset.

4.4 Analysis

Analysis of Table 1 shows humans generally outperform MLLMs across most tasks, except in Knowledge and Text-Image Understanding. This suggests that MLLMs excel in tasks requiring prior knowledge and basic text-image alignment but struggle in tasks demanding deeper cross-modal reasoning. To further investigate the factors influencing task performance, we conducted additional analyses to address three key questions: the role of prior knowledge, the necessity of visual content, and the sufficiency of reasoning difficulty.

Is prior knowledge crucial? As shown in Figure 6, children under the age of fourteen perform comparably to adults on MANBench, despite likely possessing less prior knowledge. This indicates that prior knowledge is not crucial for task performance in MANBench, emphasizing the dataset’s focus on reasoning rather than domain-specific expertise.

Is the visual contents necessary? To investigate the necessity of visual content in the tasks, we conducted an ablation study on GPT-4o by removing the images from the questions. For more experiment details, please refer to Appendix C.2. As shown in Figure 7, GPT-4o demonstrated performance divergence across 9 task categories. In the Image Consistency task, due to the lack of images, GPT-4o refused to answer most questions, which confirmed strong visual dependency. And in the Transmorphic Understanding task, the model’s accuracy slightly exceeded random baseline. A further manual review revealed that this higher accuracy resulted from the model’s preference for positive emotions, such as “joy”, which appear more frequently as correct answers in the task. For other task types, the model’s performance showed no statistically significant deviation from random chance.




Figure 8: Average human response time for all questions in MANBench. Different colors indicate distinct categories within MANBench. Each point represents the average response time and accuracy for a specific question.

Is the reasoning difficulty sufficient? To assess the reasoning difficulty of the tasks, we measured the response time of participants for each question. The results, presented in Figure 8, reveal a wide range of reasoning response times among humans on MANBench, which correlates with varying levels of accuracy. This indicates that MANBench encompasses a diverse spectrum of question difficulties, effectively evaluating the reasoning capabilities of both humans and MLLMs.

5 Conclusion

The advancement of artificial intelligence reflects humanity’s enduring pursuit to replicate and understand the complexities of human cognition. MANBench, as a meticulously designed benchmark, contributes to this endeavor by evaluating the multimodal capabilities of both humans and MLLMs. Our experiments reveal that while MLLMs excel in tasks such as text generation and knowledge retrieval, they encounter significant challenges in tasks requiring nuanced cross-modal reasoning. Even state-of-the-art models, including GPT-4o, Qwen2.5-VL-72B-Instruct, and Gemini-1.5 Pro, achieve less than 60% accuracy on MANBench, falling short of average human performance and remaining far from the upper bounds of human capability.

MANBench provides a structured framework for assessing the strengths and limitations of MLLMs, offering valuable insights into areas requiring further development. It highlights the importance of moving beyond superficial alignment toward deeper, integrative multimodal reasoning. While MANBench is a step forward, it is not without limitations, and its findings should be interpreted as part of a broader effort to advance AI. As the field

progresses, benchmarks like MANBench can serve as tools to guide research and inspire the development of systems that better approximate human understanding. Ultimately, the goal is not merely to surpass benchmarks but to foster meaningful progress in bridging the gap between artificial and human intelligence.

6 Limitations

Low population coverage: The experimental subjects in our study were predominantly drawn from a specific group within the Chinese linguistic and cultural context, which may restrict the cross-linguistic and cross-cultural adaptability of our findings. Cultural differences, variations in knowledge bases, and distinct cognitive patterns could significantly impact the generalizability of the results. Additionally, the educational backgrounds and age distributions of our participants do not fully represent the broader human population, potentially limiting the generalizability of the human benchmarks established in this study. Please refer to Appendix A.2 for more details on participant demographics.

Image Attribution and Copyright Compliance: The images utilized in our dataset were obtained from publicly available online sources. We have diligently ensured that all images included in this paper comply with applicable copyright laws and are accompanied by proper attribution. However, if you are the copyright holder of any image featured in our work and believe its usage violates your licensing terms, please contact us immediately. We are committed to resolving any legitimate copyright concerns in a timely and respectful manner. Please refer to Appendix D for more details on data licensing.

7 Ethical Considerations

Informed Consent: Prior to submitting personal information through the web interface, all participants were fully informed about the purpose of data collection and its exclusive use for scientific research.

Protection of Minors: We ensured that all minor participants completed the survey under the supervision of a guardian.

Privacy Protection: Participant privacy was strictly protected, with no personally identifiable information being analyzed or disclosed during the data analysis process. Appendix A.5 was written

only after obtaining the explicit consent of the participant involved.

Risk of Amplifying Educational Bias: The overrepresentation of highly educated participants (e.g., 40.5% with bachelor's degrees) in the dataset may inadvertently reinforce existing societal biases. If misused, such data could deepen disparities in education and employment by framing high academic achievement as a universal competency standard.

References

- Aida Amini, Saadia Gabriel, Peter Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. Mathqa: Towards interpretable math word problem solving with operation-based formalisms. *arXiv preprint arXiv:1905.13319*.
- Anthropic. 2024. Claude 3.5 sonnet. <https://www.anthropic.com/news/claude-3-5-sonnet>.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. 2024a. Are we on the right way for evaluating large vision-language models? *arXiv preprint arXiv:2403.20330*.
- Pengcheng Chen, Jin Ye, Guoan Wang, Yanjun Li, Zhongying Deng, Wei Li, Tianbin Li, Haodong Duan, Ziyang Huang, Yanzhou Su, et al. 2024b. Gmaimmbench: A comprehensive multimodal evaluation benchmark towards general medical ai. *arXiv preprint arXiv:2408.03361*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024c. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. 2024. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 11198–11201.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A Smith, Wei-Chiu Ma, and Ranjay Krishna. 2024. Blink: Multimodal large language models can see but not perceive. *arXiv preprint arXiv:2404.12390*.

- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024a. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024c. Llava-next: Improved reasoning, ocr, and world knowledge.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2024d. *Mmbench: Is your multi-modal model an all-around player?* *Preprint*, arXiv:2307.06281.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2024. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In *International Conference on Learning Representations (ICLR)*.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*.
- Fanqing Meng, Jin Wang, Chuanhao Li, Quanfeng Lu, Hao Tian, Jiaqi Liao, Xizhou Zhu, Jifeng Dai, Yu Qiao, Ping Luo, et al. 2024. Mmiu: Multimodal multi-image understanding for evaluating large vision-language models. *arXiv preprint arXiv:2408.02718*.
- OpenAI. 2024a. Gpt-4o. <https://openai.com/index/hello-gpt-4o/>.
- OpenAI. 2024b. Gpt-o1. <https://openai.com/index/openai-o1-system-card/>.
- Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. 2012. *Cats and dogs*. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3498–3505.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2014. *Imagenet large scale visual recognition challenge*.
- SenseTime. 2024. Sensenova. <https://platform.sensenova.cn/home>.
- StepFun. 2025. Step-1o. <https://platform.stepfun.com/>.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Qwen Team. 2024. Qvq: To see the world with wisdom.
- Qwen Team. 2025. Qwen2.5-vl.
- Fei Wang, Xingyu Fu, James Y. Huang, Zekun Li, Qin Liu, Xiaogeng Liu, Mingyu Derek Ma, Nan Xu, Wenxuan Zhou, Kai Zhang, Tianyi Lorena Yan, Wenjie Jacky Mo, Hsiang-Hui Liu, Pan Lu, Chunyuan Li, Chaowei Xiao, Kai-Wei Chang, Dan Roth, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2024a. *Muirbench: A comprehensive benchmark for robust multi-image understanding*. *Preprint*, arXiv:2406.09411.
- Fei Wang, Xingyu Fu, James Y Huang, Zekun Li, Qin Liu, Xiaogeng Liu, Mingyu Derek Ma, Nan Xu, Wenxuan Zhou, Kai Zhang, et al. 2024b. Muirbench: A comprehensive benchmark for robust multi-image understanding. *arXiv preprint arXiv:2406.09411*.
- Ke Wang, Juntong Pan, Weikang Shi, Zimu Lu, Mingjie Zhan, and Hongsheng Li. 2024c. *Measuring multimodal mathematical reasoning with math-vision dataset*. *Preprint*, arXiv:2402.14804.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024d. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024e. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.
- Weiyun Wang, Zhe Chen, Wenhao Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, and Jifeng Dai. 2024f. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. *arXiv preprint arXiv:2411.10442*.

Colin White, Samuel Dooley, Manley Roberts, Arka Pal, Ben Feuer, Siddhartha Jain, Ravid Shwartz-Ziv, Neel Jain, Khalid Saifullah, Siddhartha Naidu, Chinmay Hegde, Yann LeCun, Tom Goldstein, Willie Neiswanger, and Micah Goldblum. 2024. [Livebench: A challenging, contamination-free llm benchmark](#).

Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, Zhenda Xie, Yu Wu, Kai Hu, Jiawei Wang, Yaofeng Sun, Yukun Li, Yishi Piao, Kang Guan, Aixin Liu, Xin Xie, Yuxiang You, Kai Dong, Xingkai Yu, Haowei Zhang, Liang Zhao, Yisong Wang, and Chong Ruan. 2024. [Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding](#). *Preprint*, arXiv:2412.10302.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhui Chen. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of CVPR*.

Hugh Zhang, Jeff Da, Dean Lee, Vaughn Robinson, Catherine Wu, Will Song, Tiffany Zhao, Pranav Raja, Dylan Slack, Qin Lyu, et al. 2024a. A careful examination of large language model performance on grade school arithmetic. *arXiv preprint arXiv:2405.00332*.

Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Yu Qiao, et al. 2024b. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? In *European Conference on Computer Vision*, pages 169–186. Springer.

Deyao Zhu, Jun Chen, Xiaqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.

A Human Experiments

A.1 Participant Recruitment

Participants were primarily recruited by widely sharing the survey link on social media platforms, targeting a diverse audience. Additionally, personal networks were used to broaden the participant pool. Furthermore, offline testing was conducted, with participants completing surveys under researcher supervision. Notably, in all tests, we ensured that minors participated under the supervision of guardians. Each participant received a compensation of 2 USD, which is deemed fair and reasonable considering the estimated completion time of 20 minutes per survey.

Education	Number	Percentage
Bachelor	233	40.5%
Associate	106	18.4%
High School	74	12.9%
Underage	61	10.6%
Master or above	58	10.1%
Junior High School	35	6.1%
Primary School	8	1.4%
Total	575	

Table 2: Education background of the participants.

A.2 Participants Statistics

The education level distribution of the participants is shown in Table 2. The age and gender distribution of participants is shown in Figure 5.

A.3 Experiments Interface

Human Multi-Modal Ability Evaluation

Please fill in the following information

Name

Age (*)

Gender (*)

Language English

Invite Code (*)

START

If you have already registered,

CONTINUE

Figure 9: Screenshot of the User Personal Information Form: Capture essential details like name, age and gender.

Users can access the human experiments through a web interface. The interface is divided into two sections: the User Personal Information Form (Figure 9) and the User Question Interface (Figure 10). The User Personal Information Form captures essential details like name, age, gender, and education level. The User Question Interface allows users to engage with and respond to a series of questions in a structured format.

A.4 Experiment Procedure

Before beginning the experiments, users are required to complete the User Personal Information Form. Upon submitting the form, users are pro-

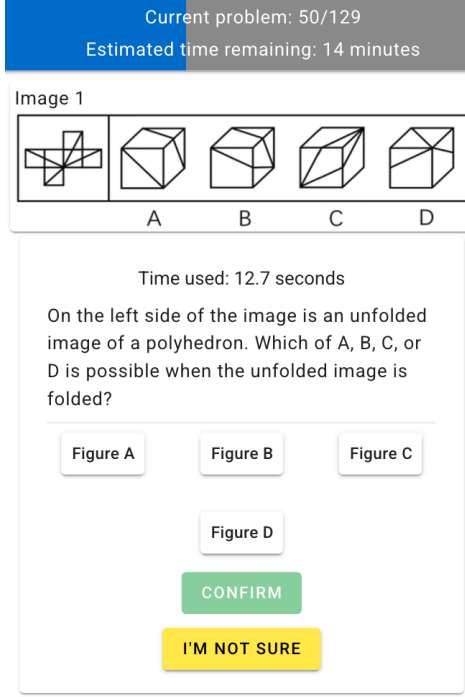


Figure 10: Screenshot of the User Question Interface: Engage with and respond to a series of questions in an organized format.

vided with the following three instructions: 1) Do not search for any information during the test. 2) Avoid distractions during the test. If necessary, you may temporarily exit and re-enter the test interface later. 3) If the image is unclear, click to enlarge it.

The questions are displayed in a random order, and each question is accompanied by a timer. While there is no time limit for completing the test, users may choose to skip questions. However, once a question is answered or skipped, it cannot be revisited.

A.5 Information of the Best Human Performance

The best human performance sample was selected from the first stage of the human experiments. The individual is a 21-year-old Chinese student, with a strong background in mathematics and computer science. With prior experience in algorithm competitions, they demonstrates exceptional reasoning abilities and the capacity to solve complex problems both quickly and accurately. Additionally, they achieved the highest GPA in their class. It is important to note that this sample does not represent the highest level of expertise globally but reflects the best human performance observed within the scope of our experiments.

Category	Correct/Total	Score
Text-image Understanding	89/100	89.00
Image Consistency	158/158	100.00
Numbers Pattern	123/134	91.79
Spatial Imagination	174/184	94.57
Text Locating	101/107	94.39
Multi-image Understanding	78/84	92.86
Knowledge	74/99	74.75
Puzzles	70/81	86.42
Transmorphic Understanding	67/74	90.54
Average		90.48

Table 3: Best human performance on each category.

Model Name	Version / Test Time
Claude-3.5-Sonnet	claude-3-5-sonnet-20241022
Gemini-1.5-Pro	gemini-1.5-pro-002
SenseNova	2025-02-08
Step-1o	step-1o-vision-32k
GPT-4o	gpt-4o-2024-11-20
GPT-o1	gpt-o1-2024-12-17

Table 4: Closed-source model names and their corresponding version or test time.

The best human performance in each category is shown in Table 3. They was tested on the Chinese subset of the MANBench dataset. They achieved a perfect score in the Image Consistent category. They answered a total of 1021 questions, which is 77.7% of the total questions in MANBench. Their total time spent on the test is 14.3 hours. Their performance is significantly higher than the average human performance across all tasks.

B MLLM Experiments

All evaluations of open-source models were conducted using four NVIDIA A800 GPUs. The execution time for each model varied significantly depending on its architectural complexity and number of parameters. In total, the comprehensive evaluation of all open-source MLLMs required approximately 200 A800-hours.

Closed-source models were evaluated using publicly available APIs. The version of each closed-source model is listed in Table 4.

C Analysis

C.1 Chinese Subset Results

The results of MLLMs on the Chinese subset are shown in Table 5 and Figure 11.

	Overall	Text-image Understanding (149)	Image Consistency (158)	Numbers Pattern (134)	Spatial Imagination (192)	Text Locating (140)	Multi-image Understanding (111)	Knowledge (120)	Puzzles (163)	Transmorphic Understanding (147)
Human (Average)	62.26	76.46	81.55	76.88	54.56	64.82	71.69	37.05	38.83	60.90
Human (Best)	90.87	89.00	100.00	91.79	94.57	94.39	92.86	74.75	86.42	90.54
Random	24.05	20.13	22.78	26.87	26.04	25.00	24.32	25.00	22.09	24.49
Open-source MLLMs										
Deepseek-VL2	43.15	73.83	26.58	57.46	29.69	42.86	38.74	<u>66.67</u>	22.09	42.18
InternVL2-26B	40.18	61.07	29.11	55.22	32.29	31.43	34.23	<u>57.50</u>	25.77	42.18
InternVL2-8B	34.47	50.34	24.68	41.79	26.04	26.43	45.05	<u>49.17</u>	22.09	34.69
InternVL2.5-26B-MPO	55.25	<u>83.89</u>	32.91	70.15	29.17	62.14	<u>76.58</u>	<u>74.17</u>	30.67	59.86
InternVL2.5-78B-MPO	60.12	88.59	37.97	75.37	37.50	65.71	78.38	83.33	33.74	61.90
QVQ-72B-Preview	49.85	73.83	25.32	61.19	32.29	48.57	64.86	<u>74.17</u>	36.81	48.98
Qwen2-VL-72B-Instruct	48.33	<u>78.52</u>	25.95	63.43	31.25	45.00	57.66	<u>73.33</u>	25.15	51.70
Qwen2.5-VL-72B-Instruct	52.05	<u>77.85</u>	38.61	67.16	33.85	60.71	<u>72.07</u>	<u>69.17</u>	26.99	40.82
Closed-source MLLMs										
Claude-3.5-Sonnet	54.79	<u>83.22</u>	44.30	64.18	36.98	43.57	69.37	<u>75.00</u>	26.99	<u>65.99</u>
GPT-4o	54.49	<u>81.88</u>	45.57	52.99	34.90	49.29	71.17	80.83	25.77	<u>65.99</u>
GPT-o1	60.12	88.59	56.96	72.39	33.33	53.57	70.27	80.83	30.67	72.79
Gemini-1.5-Pro	52.66	<u>81.88</u>	39.24	55.22	29.69	50.71	67.57	<u>79.17</u>	31.29	57.82
SenseNova	57.69	<u>84.56</u>	36.71	73.13	35.42	62.14	80.18	80.83	30.06	58.50
Step-1o	54.79	<u>84.56</u>	33.54	69.40	33.85	50.00	68.47	<u>76.67</u>	28.83	<u>66.67</u>

Table 5: Results of different models on the MANBench Chinese subset. The first row shows task names and number of test data. The best performance in each task is in bold. Values exceeding human average performance are underlined. For the information of the best human performance, please refer to Appendix A.5. The overall score is calculated by weighting the task scores according to the number of questions in each task.

C.2 Ablation Study

Since the MLLMs tend to refuse to answer the questions when the images are removed, we use a prompt to compel the MLLMs to select an answer. The prompt is “Even if the images are not provided or you are not sure about the answer, you are forced to choose one of the options”. The results are shown in Figure 7.

D Artifact Licensing and Distribution

Data Licensing: MANBench utilizes images from MUIRBench (Wang et al., 2024b), ScienceQA (Lu et al., 2022), MMStar (Chen et al., 2024a), OxfordPets (Parkhi et al., 2012), MMBench (Liu et al., 2024d), and BLINK (Fu et al., 2024). MUIRBench, ScienceQA, MMStar, and OxfordPets are licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0). MMBench and BLINK are licensed under the Apache License 2.0. MANBench is released under the CC BY-NC-SA 4.0 license, ensuring full compliance with the copyright requirements of the source datasets. All images included in this paper have been carefully reviewed to ensure compliance with applicable copyright laws and are accompanied by proper attribution. If you are the copyright holder of any image featured in this work and believe its usage violates your licensing terms, please contact us promptly. We are

committed to addressing any legitimate copyright concerns in a timely and respectful manner.

Code Licensing: The evaluation code of MANBench is developed based on VLMEvalKit (Duan et al., 2024), licensed under the Apache License 2.0. MANBench’s code is also released under the Apache License 2.0, ensuring open-source accessibility for the community. The code and dataset are available at <https://anonymous.4open.science/r/MANBench-4742/>.

E Data Collection and Privacy Protection Measures

Our dataset comprises images and questions that were meticulously curated and rigorously evaluated to ensure the exclusion of sensitive or inappropriate content. Prior to finalizing the dataset, we conducted a pilot study to assess the quality of the questions and images. This pilot involved a small cohort of participants who provided feedback on the clarity, relevance, and potential offensiveness of the materials. Their feedback was instrumental in refining the dataset to align with the highest standards of quality and ethical integrity.

Demographic information, including optional fields such as name, age, gender, and education level, was collected from participants. This information were securely stored in an encrypted




Figure 11: Performance comparison on the MANBench Chinese subset among human average accuracy, and some selected MLLMs. Please refer to Table 5 for more results.

database with access restricted to authorized personnel, ensuring full compliance with data protection regulations.

To uphold the ethical integrity of the research and protect participant privacy, rigorous anonymization protocols were implemented before the release of the experimental dataset. Specifically, all personally identifiable information (PII) was systematically removed, and data points were carefully processed to prevent any potential re-identification or traceability to individual participants. Furthermore, we strictly adhered to ethical research practices by obtaining written informed consent from all participants prior to data collection.

F Datasheet for MANBench

Motivation for Dataset Creation

Why was the dataset created? (e.g., were there specific tasks in mind, or a specific gap that needed to be filled?)

The dataset, **MANBench**, was created to address significant limitations in existing benchmarks for evaluating multimodal capabilities of both humans and Multimodal Large Language Models (MLLMs). Existing benchmarks often focus on narrow tasks, require domain-specific knowledge, and rely on insufficient human evaluations characterized by small sample sizes and limited diversity. MANBench aims to provide a comprehensive and fair assessment framework to rigorously compare human and MLLM performance across a wide range of multimodal reasoning tasks.

What (other) tasks could the dataset be used for? Are there obvious tasks for which it should *not* be used?

MANBench could be used for tasks such as:

- Multimodal reasoning: Evaluating the ability to integrate textual and visual information.
- Cross-modal understanding: Assessing tasks such as Transmorphic Understanding, Image Consistency, and Multi-image Understanding.
- Complex problem-solving: Testing capabilities in highly complex tasks, including Puzzles and Spatial Imagination.

MANBench is not suitable for tasks that rely extensively on domain-specific knowledge or for simple retrieval tasks, as it is specifically designed to emphasize reasoning and multimodal fusion.

Has the dataset been used for any tasks already? If so, where are the results so others can compare (e.g., links to published papers)?

Yes, the MANBench dataset has already been utilized for tasks involving both Multimodal Large Language Models (MLLMs) and human participants. Specifically:

MLLM Evaluation: The dataset was employed to evaluate 12 state-of-the-art MLLMs, including GPT-4o, Qwen2.5-VL-72B-Instruct, and Gemini-1.5 Pro. The results revealed that while MLLMs excel in tasks such as Knowledge and Text-Image Understanding, they struggle in tasks requiring deeper cross-modal reasoning, such as Transmorphic Understanding, Image Consistency, and Multi-image Understanding.

Human Testing: The dataset was also used to conduct large-scale human evaluations. A total of 575 participants from diverse backgrounds were recruited to complete the tasks in MANBench. These human evaluations established a performance baseline for comparing MLLMs and highlighted areas where both humans and MLLMs encounter similar difficulties, such as Puzzles and Spatial Imagination.

Who funded the creation of the dataset? If there is an associated grant, provide the grant number.

No external funding was used in the creation of the MANBench dataset.

Dataset Composition

What are the instances? (that is, examples; e.g., documents, images, people, countries) Are there multiple types of instances? (e.g., movies, users, ratings; people, interactions between them; nodes, edges)

The instances in MANBench consist of 1,314 bilingual(English and Chinese) questions and 2,231 images, organized across nine tasks which are in visual question answer(VQA) format.

Are relationships between instances made explicit in the data (e.g., social network links, user/movie ratings, etc.)?

The relationships between instances are implicit rather than explicit. Each question and its associated image(s) are designed to evaluate specific multimodal reasoning abilities. However, there are no explicit relationships, such as social network links or user/movie ratings, between the instances.

How many instances of each type are there?

The MANBench benchmark comprises a total of 1,314 questions distributed across nine distinct multimodal tasks, each designed to rigorously evaluate the capabilities of both humans and Multimodal Large Language Models (MLLMs). The distribution of tasks is as follows: Text-image Understanding (149 instances), Image Consistency (158 instances), Numbers Pattern (134 instances), Spatial Imagination (192 instances), Text Locating (140 instances), Multi-image Understanding (111 instances), Knowledge (120 instances), Puzzles (163 instances), and Transmorphic Understanding (147 instances).

What data does each instance consist of? "Raw" data (e.g., unprocessed text or images)? Features/attributes? Is there a label/target associated with instances? If the instances are related to people, are subpopulations identified (e.g., by age, gender, etc.) and what is their distribution?

Each instance consists of textual questions (including the prompt, options, answers, and the task to which it belongs) and images.

Is everything included or does the data rely on external resources? (e.g., websites, tweets, datasets) If external resources, a) are there guarantees that they will exist, and remain constant, over time; b) is there an official archival version. Are there licenses, fees or rights associated with *any* of the data?

Everything included in the MANBench dataset is self-contained and does not rely on external resources.

Are there recommended data splits or evaluation measures? (e.g., training, development, testing; accuracy/AUC)

We follow standard setups as in the VLMEvalKit (Duan et al., 2024), where the temperature is set to 0 and retry is set to 10.

What experiments were initially run on this dataset? Have a summary of those results and, if available, provide the link to a paper with more information here.

The initial experiments involved evaluating 12 state-of-the-art MLLMs on MANBench and comparing their performance to that of 575 human participants across 9 tasks. The results showed that MLLMs excel in tasks like Knowledge and Text-Image Understanding but struggle in tasks requiring deeper cross-modal reasoning.

Data Collection Process

How was the data collected? (e.g., hardware apparatus/sensor, manual human curation, software program, software interface/API; how were these constructs/measures/methods validated?)

The MANBench dataset was collected by authors manually.

Who was involved in the data collection process? (e.g., students, crowdworkers) How were they compensated? (e.g., how much were crowdworkers paid?)

Only the authors were involved in the data collection process. No crowdworkers or external contributors were used.

Over what time-frame was the data collected? Does the collection time-frame match the creation time-frame?

The MANBench dataset was collected over a period of 3 months. The collection time-frame matches the creation time-frame.

How was the data associated with each instance acquired? Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part of speech tags; model-based guesses for age or language)? If the latter two, were they validated/verified and if so how?

All the data associated with each instance was directly observable and manually curated by the authors.

Does the dataset contain all possible instances? Or is it, for instance, a sample (not necessarily random) from a larger set of instances?

The MANBench dataset comprises a comprehensive set of instances designed to evaluate a wide range of multimodal reasoning tasks. However, it is not exhaustive and may not cover all possible multimodal reasoning scenarios.

If the dataset is a sample, then what is the population? What was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)? Is the sample representative of the larger set (e.g., geographic coverage)? If not, why not (e.g., to cover a more diverse range of instances)? How does this affect possible uses?

The MANBench dataset is not a sample but a comprehensive collection of instances designed to evaluate a wide range of multimodal reasoning tasks. The instances are carefully selected to cover a diverse range of multimodal reasoning scenarios and are representative of the larger set of multimodal reasoning tasks.

Is there information missing from the dataset and why? (this does not include intentionally dropped instances; it might include, e.g., redacted text, withheld documents) Is this data missing because it was unavailable?

No, there is no information missing from the dataset. All instances are complete and self-contained.

Are there any known errors, sources of noise, or redundancies in the data?

No, the MANBench dataset has been carefully curated to ensure accuracy and consistency. There are no known errors, sources of noise, or redundancies in the data. If any errors are identified, they will be corrected in future updates.

Data Preprocessing

What preprocessing/cleaning was done? (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values, etc.)

We conducted manual preprocessing to ensure the quality of the data.

Was the “raw” data saved in addition to the preprocessed/cleaned data? (e.g., to support unanticipated future uses)

Yes, the raw data was saved by the authors in addition to the preprocessed/cleaned data.

Is the preprocessing software available?

No, the preprocessing software is not available.

Does this dataset collection/processing procedure achieve the motivation for creating the dataset stated in the first section of this datasheet?

Yes, the dataset collection and processing procedures fulfill the motivation for creating the dataset by providing a comprehensive and fair assessment framework to rigorously compare human and Multimodal Large Language Model (MLLM) performance across a wide range of multimodal reasoning tasks.

Dataset Distribution

How is the dataset distributed? (e.g., website, API, etc.; does the data have a DOI; is it archived redundantly?)

The MANBench dataset is distributed through the official [Hugging Face repository](#).

When will the dataset be released/first distributed? (Is there a canonical paper/reference for this dataset?)

The dataset will be released on Feb 10th, 2025.

What license (if any) is it distributed under? Are there any copyrights on the data?

The MANBench dataset is released under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0).

Are there any fees or access/export restrictions?

No. The MANBench dataset is freely available for non-commercial use, sharing, and adaptation.

Dataset Maintenance

Who is supporting/hosting/maintaining the dataset? How does one contact the owner/curator/manager of the dataset (e.g. email address, or other contact info)?

The first author of the paper “MANBench: Is Your Multimodal Model Smarter than Human?” is responsible for supporting, hosting, and maintaining the MANBench dataset. However, during the review process, due to the double-blind policy, we are unable to provide contact information. Once the paper is accepted, the contact information will be updated.

Will the dataset be updated? How often and by whom? How will updates/revisions be documented and communicated (e.g., mailing list, GitHub)? Is there an erratum?

Yes. The dataset may be updated periodically to address errors, improve the quality of the data, or add new instances. Updates will be documented

and communicated through the official [Hugging Face repository](#).

If the dataset becomes obsolete how will this be communicated?

We will communicate the obsolescence of the dataset through the official [GitHub repository](#).

Is there a repository to link to any/all papers/systems that use this dataset?

Yes, once the paper is accepted, we will provide a link to the official [GitHub repository](#) where all papers and systems utilizing the MANBench dataset will be linked.

If others want to extend/augment/build on this dataset, is there a mechanism for them to do so? If so, is there a process for tracking/assessing the quality of those contributions. What is the process for communicating/distributing these contributions to users?

Yes, we encourage researchers to extend, augment, and build on the MANBench dataset. MANBench is released under the CC BY-NC-SA 4.0 license, which permits non-commercial use, sharing, and adaptation. Researchers can contribute to the dataset by submitting pull requests to the official [Hugging Face repository](#). All contributions will be reviewed by the dataset maintainers to ensure quality and relevance before being incorporated into the dataset.

Legal & Ethical Considerations

If the dataset relates to people (e.g., their attributes) or was generated by people, were they informed about the data collection? (e.g., datasets that collect writing, photos, interactions, transactions, etc.)

The MANBench dataset does not generate any data that relates to people, but we collected data related to people from existing datasets.

If it relates to other ethically protected subjects, have appropriate obligations been met? (e.g., medical data might include information collected from animals)

No, the dataset does not relate to other ethically protected subjects.

If it relates to people, were there any ethical review applications/reviews/approvals? (e.g. Institutional Review Board applications)

No, the dataset only contains data related to people from existing datasets, which have already been reviewed and approved by the respective institutions.

If it relates to people, could this dataset expose people to harm or legal action? (e.g., financial social or otherwise) What was done to mitigate or reduce the potential for harm?

No, the dataset does not contain any information that could expose people to harm or legal action.

If it relates to people, does it unfairly advantage or disadvantage a particular social group? In what ways? How was this mitigated?

No, the dataset does not unfairly advantage or disadvantage any particular social group.

If it relates to people, were they provided with privacy guarantees? If so, what guarantees and how are these ensured?

Yes, all personally identifiable information (PII) was systematically removed, and data points were carefully processed to prevent any potential re-identification or traceability to individual participants. The dataset is designed to be fully anonymized and does not contain any sensitive information.

Does the dataset comply with the EU General Data Protection Regulation (GDPR)? Does it comply with any other standards, such as the US Equal Employment Opportunity Act?

Yes, the dataset complies with the EU General Data Protection Regulation (GDPR) and other relevant data protection standards.


Does the dataset contain information that might be considered sensitive or confidential? (e.g., personally identifying information)

No, the dataset does not contain any sensitive or confidential information.

Does the dataset contain information that might be considered inappropriate or offensive?

No, the dataset does not contain any inappropriate or offensive information.


Image Consistency



Prompt: Disclaimer: This is not to make unfair assumptions about the people in the image and you just need to give your assessment on this question. You don't need to identify the real people. You just need to analyze based on the information I gave you. Given the first image with the lower right corner missing, can you tell which one of following images is the missing part? Imagine which image would be more appropriate to place in the missing spot. You can also carefully observe and compare the edges of the images. \nPlease select the correct answer from the following options:\n(A) the second image\n(B) the third image\n(C) the fourth image\n(D) the fifth image

Ground Truth: (B)

GPT-4o: To determine which image best fits the missing lower right corner of the first image, we need to match the colors and alignments of the lines from the partial image with those in the options.\n\nUpon careful observation, option (B) the third image is the correct choice. The green, purple, and light green lines perfectly match and align with the curvature of the incomplete pattern in the first image.



Prompt: Disclaimer: This is not to make unfair assumptions about the people in the image and you just need to give your assessment on this question. You don't need to identify the real people. You just need to analyze based on the information I gave you. Given the first image with the lower right corner missing, can you tell which one of following images is the missing part? Imagine which image would be more appropriate to place in the missing spot. You can also carefully observe and compare the edges of the images. \nPlease select the correct answer from the following options:\n(A) the second image\n(B) the third image\n(C) the fourth image\n(D) the fifth image

Ground Truth: (C)

GPT-4o: To determine which image fits the missing lower right corner of the first image, we need to look for continuity in both the natural features (such as the path and vegetation) and any lines or patterns that seem to extend from the main image.\n\nObserving the main image, the missing portion should continue the path that is partially visible. The image should show a similar rock and plant environment, especially matching the lower part of the staircase.\n\nComparing the options:\n- (A) the second image shows rocks and some greenery but does not align with the staircase.\n- (B) the third image shows the continuation of the path and stairs with similar vegetation and fence posts, making it a suitable match.\n- (C) the fourth image is a close-up of stairs, but these steps don't match the end of those in the main image.\n- (D) the fifth image mostly portrays greenery without any evident continuation of the path or stairs.\n\nTherefore, the option that completes the missing part of the image with a natural continuation of the path and structure is (B) the third image.

Figure 12: Examples of data in MANBench with GPT-4o predictions.(1/9)

Text-image Understanding



Image 1

Prompt: What is the most likely scenario for the image?
Please select the correct answer from the following options:
(A) At a family birthday party
(B) At the end of a sports competition
(C) During the concert
(D) At office meetings
Ground Truth: (C)

GPT-4o: The most likely scenario for the image is:
(C) During the concert
This is a depiction of a performer crowd-surfing during a concert, which is a popular activity in live music settings.



Image 1

Prompt: Which of the following options explains the strip of light at the axial location of the image?
Please select the correct answer from the following options:
(A) The river that connects to the lake in the lower left corner reflects the sunlight.
(B) Highways that reach into the city reflect the sun's rays
(C) Streetlights on city streets have not been turned off
(D) Searchlights on top of the tallest buildings light up the city
Ground Truth: (B)

GPT-4o: The correct answer is:
(A) The river that connects to the lake in the lower left corner reflects the sunlight.
Explanation: The strip of light seen in the image corresponds to sunlight reflecting off the surface of the water in the river, which runs through the city. The water's reflective property causes it to appear illuminated, creating the bright strip extending toward the lake.

Figure 13: Examples of data in MANBench with GPT-4o predictions.(2/9)

Multi-image Understanding




Image 1




Image 2

Prompt: Which number is missing in both images?\nPlease select the correct answer from the following options:\n(A) 8\n(B) 7\n(C) 3\n(D) 2
Ground Truth: (A)

GPT-4o: The missing number in both images is (C) 3.



Image 1



Image 2

Prompt: What changes from the first image to the second image?\nPlease select the correct answer from the following options:\n(A) The person in the image stops pouring water\n(B) The person in the image starts pouring water\n(C) The person in the image puts the kettle on the table\n(D) The person in the image picks up the kettle
Ground Truth: (A)

GPT-4o: The correct answer is (B) The person in the image starts pouring water.

Figure 14: Examples of data in MANBench with GPT-4o predictions.(3/9)

Knowledge



Image 1

Prompt: Which of the following species does the cat in the picture belong to?
Please select the correct answer from the following options:
(A) Maine Coon
(B) Ragdoll
(C) Egyptian Mau
(D) Persian
Ground Truth: (C)

GPT-4o: The cat in the picture appears to be an **Egyptian Mau**. This breed is known for its spotted coat, slender build, and graceful appearance. So the correct answer is: **(C) Egyptian Mau**




Image 1

Prompt: What is the function of the eight black objects arranged longitudinally in the picture?
Please select the correct answer from the following options:
(A) Provide power to the space station
(B) Cool the space station
(C) Decompose waste gas for the space station
(D) Heat the space station
Ground Truth: (A)

GPT-4o: The eight black objects arranged longitudinally in the picture are **solar panels**. Their primary function is to: **(A) Provide power to the space station.**
Solar panels convert sunlight into electrical energy, which is used to power the systems and equipment of the International Space Station (ISS).

Figure 15: Examples of data in MANBench with GPT-4o predictions.(4/9)

Spatial Imagination




Image 1

Prompt: On the left side of the image is an unfolded image of a polyhedron. Which of A, B, C, or D is possible when the unfolded image is folded?
Please select the correct answer from the following options:
(A) Figure A
(B) Figure B
(C) Figure C
(D) Figure D

Ground Truth: (A)

GPT-4o: To solve this problem, let's analyze how the given unfolded polyhedron (net) on the left side folds into a 3D shape. The unfolded image shows four triangular faces arranged in a specific layout.
Step 1: Identify the shape of the polyhedron
The net corresponds to a tetrahedron because it consists of four triangular faces. When folded, the triangles must meet at a single vertex, forming a closed 3D triangular pyramid.
Step 2: Analyze the patterns
The patterns (solid black, checks, stripes, etc.) on the triangular faces need to align when the net is folded. By mentally folding the net into a pyramid, you can visualize how these patterns will appear on the finalized shape.
Step 3: Match the folded appearance with the options
Carefully match the patterns on the faces of the net with the possible 3D shapes shown in Figures A, B, C, and D. Folding the net as a tetrahedron:
- The solid black triangle will always touch the striped and checked triangles.
- The checked triangle will border the striped and the dotted triangles.
- The dotted triangle sits opposite the solid black triangle.
Only **Figure A** correctly represents this arrangement when the net is folded.
Final Answer: **(A) Figure A**




Image 1




Image 2




Image 3




Image 4




Image 5

Prompt: The first image is a structure built by cubes. The first layer has 25 cubes, the second layer has 9 cubes, and the third layer has 1 cube. Now the structure is split into two parts, each subsequent image contains two parts of the structure. Which image is the correct one?
Please select the correct answer from the following options:
(A) the second image
(B) the third image
(C) the fourth image
(D) the fifth image

Ground Truth: (B)

GPT-4o: To select the correct image demonstrating the structure split into two parts, let's analyze the first image and its subsequent layers.
1. **Structure Description:**
- The first layer contains a base of **25 cubes** (a 5x5 square).
- The second layer sits on top and has **9 cubes** (a 3x3 square centered on the base).
- The third layer is the topmost, with a single cube (a 1x1 square in the center).
2. **Splitting the Structure:**
- When splitting this structure into two parts, there should be a clear division of the total cube configuration so that both parts represent logical portions of the original structure while matching the logical layout from the original image.
3. **Analyzing the Options:**
- **Second Image (A):** The division is incorrect because the cube arrangement differs from the original layering and would not reassemble properly into the original structure.
- **Third Image (B):** Again, the division is inconsistent with the logical layers, so this cannot be correct.
- **Fourth Image (C):** This image portrays a split where one part captures the entire top and middle layers, leaving the bottom layer separate, preserving the correct cube alignment for reassembly. **This is the correct split.**
- **Fifth Image (D):** While closer, this split still does not align logically with the original cube configuration.
Correct Answer: **(C) The fourth image**

Figure 16: Examples of data in MANBench with GPT-4o predictions.(5/9)

Transmorphic Understanding



Image 1

Prompt: What emotion does this picture convey?\nPlease select the correct answer from the following options:\n(A) Boredom\n(B) Satisfaction\n(C) Joy\n(D) Anger

Ground Truth: (D)

GPT-4o: The image depicts a bag that unintentionally resembles a face with a frown and a wide-open \"mouth.\" Based on this expression, the correct emotional interpretation would be:\n\n**(D) Anger.**




Image 1

Prompt: Based on the shape of the object in the picture, what might be the first thing that comes to mind?Please select the correct answer from the following options:\n(A) A big tree with lots of leaves\n(B) A vast expanse of grassland\n(C) Elegant Ballet Dancer\n(D) A muscular man

Ground Truth: (C)

GPT-4o: The correct answer is: \n\n**(C) Elegant Ballet Dancer**\n\nThe shape of the tree resembles the pose of a graceful ballet dancer, likely due to the positioning of the trunk and branches.

Figure 17: Examples of data in MANBench with GPT-4o predictions.(6/9)

Puzzles




Image 1

Prompt: The image shows a maze, where white cells are accessible and others are walls. The green circle is the start point X, and the red circles are the target points A, B, C, D. Which target point is the closest to the start point X? Please select the correct answer from the following options:\n(A) A\n(B) B\n(C) C\n(D) D

Ground Truth: (C)

GPT-4o: To determine which target point is the closest to the start point X (green circle), we'll visually analyze the maze and trace the path to points A, B, C, and D.\n\nHere's the analysis:\n\n1. **Point A**: Point A (red circle) is below point X, requiring several turns to reach it. The path includes a curve downward, moving right, and coming back left.\n\n2. **Point B**: Point B is in the bottom-left corner, quite distant from X. The path to B seems significantly longer compared to the others.\n\n3. **Point C**: Point C is located very close to X and is accessible via a very short path. It appears to be the closest target by far.\n\n4. **Point D**: Point D is farther to the lower-right of X, requiring a longer path compared to reaching C or A.\n\nThus, **Point C** is the closest target to the starting point X.\n\n**Correct Answer:** (C) C




Image 1

Prompt: During the IQ test, you will be presented with existing picture example, and four picture options. Your task is to identify the one picture that follows the same pattern or rule established by the previous pictures. Here are some strategies to help you determine the right choice: 1. Look for patterns related to quantity or numbers that progress from one picture to the next. 2. Check if there is a consistent way the images are rotated or flipped in sequence. 3. Identify a common feature that each successive picture shares with the previous one, while also paying attention to any variations that might indicate a regular progression or change. Select from the following another choices. Please select the correct answer from the following options:\n(A) Picture A\n(B) Picture B\n(C) Picture C\n(D) Picture D

Ground Truth: (B)

GPT-4o: To identify the correct answer, let's analyze the sequence and identify the pattern.\n\nIn the top row, the black shapes maintain a specific structure but appear to rotate **90 degrees clockwise** as you move from one image to the next.\n\nThe sequence is as follows:\n\n1. First image: The shape in its original orientation.\n\n2. Second image: Rotated 90 degrees clockwise.\n\n3. Third image: Rotated another 90 degrees clockwise (total 180 degrees from the first image).\n\nTo predict the next image in the sequence, the shape needs to be rotated another **90 degrees clockwise** (total 270 degrees from the original image).\n\nNow, compare the answer choices:\n\n- (A): Does not match the expected 270-degree clockwise rotation.\n\n- (B): Does not match the expected orientation.\n\n- (C): Matches the shape rotated 270 degrees clockwise.\n\n- (D): Does not match the expected orientation.\n\nThe correct answer is **(C)**.

Figure 18: Examples of data in MANBench with GPT-4o predictions.(7/9)

Text Locating



Image 1

Prompt: Four points are circled in the image, labeled A, B, C, and D. Which best fits the following statement: which point is the traffic light in front of the black vehicle?
Please select the correct answer from the following options:
(A) A
(B) B
(C) C
(D) D
Ground Truth: (A)

GPT-4o: The traffic light in front of the black vehicle is labeled **A**.




Image 1

Prompt: Four points are circled in the image, labeled A, B, C, and D. Which best fits the following statement: which point is the black box located on the table?
Please select the correct answer from the following options:
(A) A
(B) B
(C) C
(D) D
Ground Truth: (A)

GPT-4o: Based on the image provided, the black box located on the table is highlighted near **Point B**.
Correct answer: **(B) B**

Figure 19: Examples of data in MANBench with GPT-4o predictions.(8/9)

Numbers Pattern




Image 1

Prompt: What is the most likely number at the question mark in the image?
Please select the correct answer from the following options:
(A) 3
(B) 4
(C) 8
(D) 9
Ground Truth: (D)

GPT-4o: To solve this, let's analyze the image:
1. The figure associated with the number 3 is a triangle (3 sides).
2. The figure associated with the number 4 is a star with 4 points.
3. The figure associated with the number 5 is a star with 5 points.
Following this pattern, the number corresponds to the number of points or vertices of the shape. The figure with the "?" has 8 points.
Correct answer: (C) 8



Image 1

Prompt: How many horses are in the image?
Please select the correct answer from the following options:
(A) 6
(B) 7
(C) 8
(D) 9
Ground Truth: (A)

GPT-4o: The correct answer is **(A) 6**. There are 6 horses visible in the image.

Figure 20: Examples of data in MANBench with GPT-4o predictions.(9/9)