

CONSTRUCTURE: Benchmarking CONCEPT STRUCTURE REASONING for Multimodal Large Language Models

Anonymous ACL submission

Abstract

Multimodal Large Language Models (MLLMs) have shown promising results in various tasks, but their ability to perceive the visual world with deep, hierarchical understanding similar to humans remains uncertain. To address this gap, we introduce CONSTRUCTURE, a novel concept-level benchmark to assess MLLMs' hierarchical concept understanding and reasoning abilities. Our goal is to evaluate MLLMs across four key aspects: 1) Understanding atomic concepts at different levels of abstraction; 2) Performing upward abstraction reasoning across concepts; 3) Achieving downward concretization reasoning across concepts; and 4) Conducting multi-hop reasoning between sibling or common ancestor concepts. Our findings indicate that even state-of-the-art multimodal models struggle with concept structure reasoning (e.g., GPT-4o averages a score of 62.1%). We summarize key findings of MLLMs in concept structure reasoning evaluation. Moreover, we provide key insights from experiments using CoT prompting and fine-tuning to enhance their abilities.

1 Introduction

The basic level is the level in a taxonomy at which things are normally named, in the absence of reasons to the contrary. 'Dog' is a basic level category, 'boxer' a subordinate category, 'quadruped' a superordinate category.

— John R. Taylor

According to the prototype theory (Taylor, 2019) in cognitive science, humans perceive the visual world hierarchically, with basic, subordinate, and superordinate categories. People interpret the world differently based on these conceptual levels in diverse environments. As shown in Figure 1, humans can not only understand basic concepts,

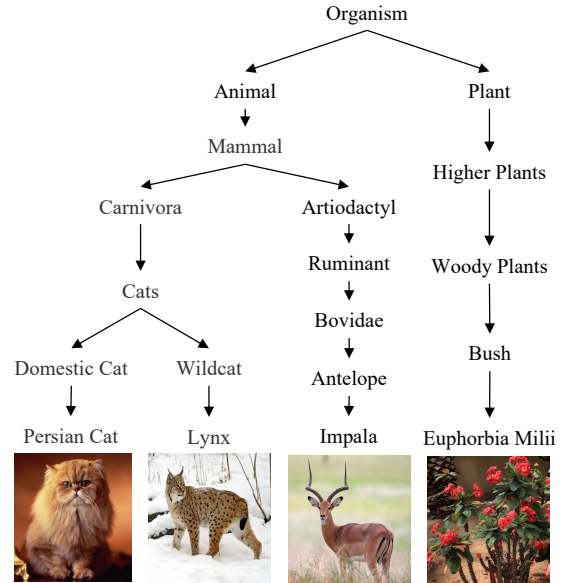


Figure 1: Demo of Concept Structure.

but also have a deep concept structure cognition in minds (Medin, 1989; Suresh et al., 2023). This raises the question: do multimodal AI systems, like Multimodal Large Language Models (e.g. GPT-4V (OpenAI, 2023)), exhibit similar concept structure cognition ability?

Recently, pretrained multimodal large language models (MLLMs) have transcended the confines of text-only modalities, gaining a deeper conceptual understanding of the world and demonstrating remarkable performance across a spectrum of downstream tasks. As a result, there is a growing importance and challenge in conducting comprehensive evaluations of these models to better understand their capabilities and pinpoint areas for enhancement. Inspired by the cognitive processes underlying human conceptual understanding, we posit that similar deep and structured visual conceptual cognition capabilities are pivotal for the profound comprehension of MLLMs. This compelling question demands exploration: Do MLLMs truly grasp

Benchmark	Assessment of cognitive related abilities	Data Source	Answer Type	Evaluator	Size
LVLm-eHub (Xu et al., 2023)	Visual Reasoning, Visual Commonsense	Datasets	MC/OE	Metrics/LLMs/User	332k
MME (Fu et al.)	Commonsense Reasoning, Numerical Calculating, Text Translation, Code Reasoning	Datasets	MC	Accuracy	2,194
MMBench (Liu et al., 2023)	Reasoning(e.g. Physical Relation Reasoning, Social Relation Reasoning)	Datasets/Handcraft/LLMs	MC	Accuracy	2,974
SEED-Bench (Li et al., 2023a)	Visual Reasoning, Spatial Relation	Handcraft/LLMs	MC	Accuracy	19k
MM-Vet (Yu et al., 2023)	Spatial awareness, Knowledge, Math	Datasets/Handcraft	OE	LLMs	218
EgoThink (Cheng et al., 2023)	Scene Reasoning, Planning	Handcraft	OE	LLMs	700
CONSTRUCTURE(Ours)	Concept Structure Reasoning	Datasets/Handcraft	MC	Accuracy	2,064

Table 1: Comparison of recent comprehensive evaluation benchmarks of MLLMs and our proposed benchmark M²C²-Bench.

and internalize concepts, or do they merely learn the superficial concept alignment through pretraining?

However, addressing this question requires a deeper investigation into the underlying mechanisms and limitations of MLLMs. As shown in Table 1, the deep visual concept structure cognition remains largely unexplored in existing benchmarks. On one hand, these benchmarks seldom consider visual cognitive capabilities. On the other hand, existing benchmarks for visual cognition typically focus solely on assessing visual reasoning ability related to the whole image content, overlooking the evaluation of deep and structural visual concept cognition. For example, SEED-bench ?? evaluates spatial relationships or visual reasoning tasks based on the content of images, while MME ?? assesses common-sense reasoning, numerical computation, code inference, and text translation. LVLm-eHub ?? evaluates visual reasoning and visual common sense. However, all these benchmarks require answering questions based on the entire content of an image. Therefore, there is a pressing need to construct an benchmark specifically designed to assess the deep visual concept structure cognition in MLLMs.

In this work, we propose the novel CONSTRUCTURE benchmark focusing on deep visual concept structure cognition. To fully uncover the deep visual cognitive ability of MLLMs across the overall concept structure beyond the Figure 1, we consider the following four key capabilities, ranging from *atomic concept understanding*, *concept abstraction reasoning*, *concept concretization reasoning*, to *common ancestor reasoning*. The first capability is to evaluate the atomic concept understanding at different levels. The last three capabilities are to evaluate reasoning ability on concept structure, including upward abstraction reasoning from child concepts to parent concepts, downward concretization reasoning from parent concepts to child concepts and multi-hop reasoning between sibling concepts or common ancestor concepts, respectively.

Based on the our proposed CONSTRUCTURE benchmark, we conduct comprehensive experiments to evaluate concept structure cognition capabilities of fourteen popular MLLMs (including 6 api-based MLLMs and 8 open-sourced MLLMs). We conclude the main findings as follows:

1) Current MLLMs possess a certain level of conceptual understanding, but their performance in concept structure reasoning is poor. The best model, GPT-4o, only achieved a score of 0.621, indicating significant room for improvement.

2) In concept structure reasoning tasks, MLLMs perform the worst in common ancestor reasoning. The main reasons for errors are inconsistencies in the reasoning process and constraint violations. This demonstrates that adhering to multiple constraints and maintaining consistency in the reasoning process are key challenges to enhancing MLLMs’ concept structure cognitive abilities.

3) MLLMs still need to improve their ability to reason about hierarchical relationships in concept structures.

4) MLLMs’ performance deteriorates as the concept hierarchy deepens and granularity increases. Therefore, improving fine-grained concept recognition and the ability to reason about related fine-grained concept structures is crucial for enhancement.

Our evaluation results reveal the limitations of MLLMs in concept structure cognition, providing a comprehensive and clear analysis that directs further improvements for MLLMs. Furthermore, we improved MLLMs’ concept structure reasoning abilities through few-shot CoT prompting and fine-tuning methods, and elucidated three key insights in the discussion.

2 Related Work

2.1 Multimodal Large Language Models

Building upon the monumental achievements of large language models (LLMs) (Du et al., 2022; OpenAI, 2023; Zhu et al., 2023), recent advance-

ments in multimodal LLMs (MLLMs) have firmly established LLMs as their foundational backbone. Initially, MLLMs commence pre-training using large-scale image-text paired datasets (Yang et al., 2023; Li et al., 2023b) or by integrating random visual and textual data (Ye et al., 2023; Bai et al., 2023). This foundational phase is further enriched by leveraging extensive image-text instruction datasets (Dai et al., 2024). Recent studies (Liu et al., 2024) have increasingly employed fine-tuning strategies, significantly enhancing MLLMs’ capacity to deliver superior performance in various downstream tasks and benchmarks.

2.2 Evaluations of MLLMs

Current benchmarks aim to comprehensively assess MLLMs’ capabilities in multimodal cognition tasks like reasoning. For example, Lvlm-ehub (Xu et al., 2023) evaluates visual reasoning and common-sense cognition across 47 benchmarks. MME ?? covers inference, numerical computation, translation, and code reasoning. SEED-bench ?? focuses on visual and spatial reasoning, while MM-VET ?? evaluates spatial relationships, knowledge, and math abilities. EgoThink ?? assesses scene reasoning and planning. Our CONSTRUCTURE benchmark specifically targets hierarchical concept structure reasoning in MLLMs’ understanding of visual concepts.

3 CONSTRUCTURE Benchmark

In this section, we first elaborate on the capability and question sets used to assess concept structure cognition abilities. Following that, we outline the process of constructing the test dataset.

3.1 Evaluation Capability

As shown in Figure 2, we evaluate the following four key capabilities to uncover the concept structure cognition ability of MLLMs., ranging from *atomic concept understanding*, *concept abstraction reasoning*, *concept concretization reasoning*, to *common ancestor reasoning*. We explain why each capability is needed with a question inspired by human concept cognition and introduce how to evaluate the capability with examples.

Atomic Concept Understanding. *How can MLLMs understand atomic concepts at different levels of abstraction?* Human cognition of concepts has different levels of abstraction, and we can not only understand concrete concepts like “Persian

Cat”, but also abstract concepts like “Mammal”. How is MLLMs capable of understanding concepts of various levels of abstraction? As shown in upper left part of Figure 2, to evaluate atomic concept understanding ability, we query MLLMs with a simple discriminant question (*i.e.* “*Is the concept depicted in the image a {concept_name}?*”) with true or false options related to specific concepts at various abstraction levels. To answer this question, MLLMs need to have a multi-level understanding of visual concepts.

Concept Abstraction Reasoning. *Can MLLMs perform upward abstraction reasoning across concepts at different levels of abstraction?* Human beings can categorize concrete concepts into upper level abstract concepts, *e.g.* categorize “penguins” and “sharks” to “birds” and “fish”, respectively. How well do MLLMs perform this kind of abstraction reasoning process? As shown in upper right part of Figure 2, to evaluate concept abstraction reasoning ability, we query MLLMs with multiple-choice questions to select the most abstract and general concept from candidate options. The image is aligned correctly with several concepts in candidate options, MLLMs need to recognize them and figure out one has the most abstract level. To answer this question, MLLMs need firstly recognize correct options and then reason out the most abstract concept aligned with the image.

Concept Concretization Reasoning. *Can MLLMs achieve downward concretization reasoning across concepts at various levels of abstraction?* Human beings can refine abstract concepts to lower level concrete concepts, *e.g.* recognize from the animal categories (“Cats”) to fine-grained breed (“Persian Cat”). How well do MLLMs perform this kind of concretization reasoning process? As shown in lower left part of Figure 2, to evaluate concept concretization reasoning ability, we query MLLMs with multiple-choice questions to select the most specific and accurate concept from candidate options. The image is aligned correctly with several concepts in candidate options, MLLMs need to recognize them and figure out the most concrete one. To answer this question, MLLMs need firstly recognize correct options and then reason out the most concrete concept aligned with the image.

Common Ancestor Reasoning. *Can MLLMs perform multi-hop reasoning between sibling concepts or common ancestor concepts?* Human perception of visual concepts follows a hierarchical struc-

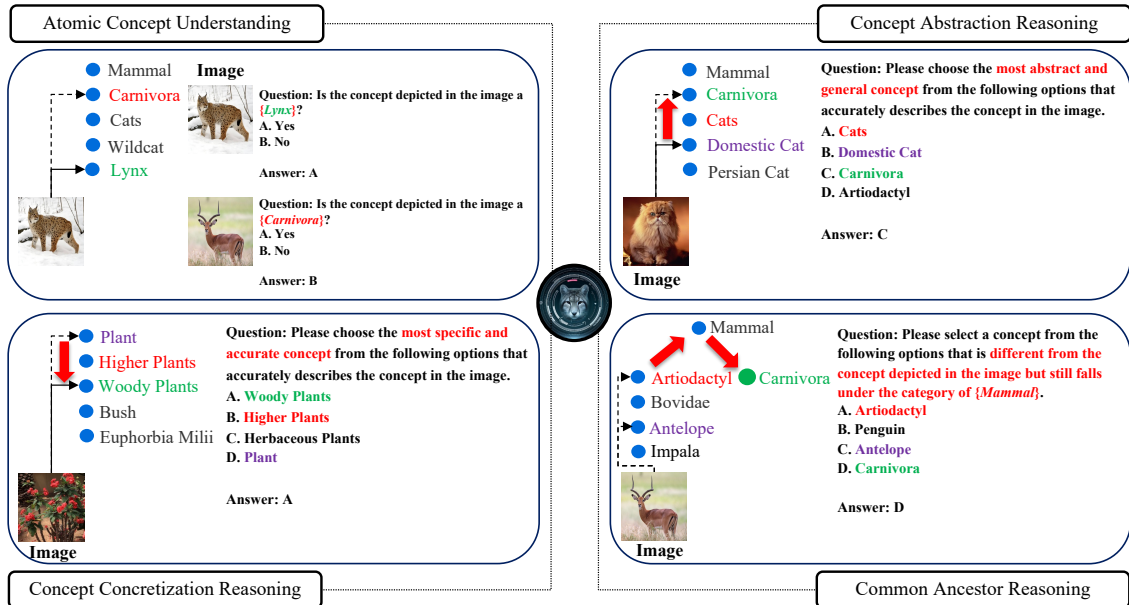


Figure 2: CONSTRUCTURE Benchmark.

ture. For instance, we recognize that “dogs” and “cats” have a higher-level common ancestor concept “mammal”, and within the category of “dogs”, there are subcategories like “pet dog” and “hunting dog”. Can MLLMs fully grasp the various hierarchical levels of a concept? Can they reason that two concepts have a common ancestor concept or are they in a sibling relationship? As shown in lower right part of Figure 2, to evaluate common ancestor reasoning ability, we query MLLMs with multiple-choice questions to select a concept which is different from as well as share the same parent or common ancestor concept with the concept depicted in the image. To answer this question, MLLMs need have the ability to reason across sibling concepts or common ancestor concepts, which indicates that MLLMs need to have the structure cognition ability of the related concepts.

3.2 Data collection

In this section, we introduce details of data collection to construct our CONSTRUCTURE benchmark.

Collecting Concept Taxonomy and Construct Concept Chains. We use isA relations in Chinese Open WordNet (Wang and Bond, 2013) as our data source of concept structure, which encompasses 19.7K isA relations. We remove the identity isA relation in the raw data and construct an isA concept tree (with a root node). During construction, we drop potential isA relations to avoid the circular

dependency. After that, we recursively search for the isA concept chains with a length of 5 as our candidates for the next stage of sample generation.

Collecting Visual Concept Image with Manual Check. Since concept in collected candidate chains may not be visual concepts, we leverage M²ConceptBase (Zha et al., 2023) (a multimodal knowledge base has rich concept-image alignments) to filter out visual concepts. Then we carefully check whether the lowest concept in the concept chain corresponds to the image in the knowledge base correctly, and search for correct images from the Internet for lower quality or wrong images. We also put in considerable manual efforts to check the correctness of the concept chains. We remove all wrong candidate concept chains or wrong isA relation part (rooted in raw data from Chinese Open Wordnet) in candidate chains, and finally get totally 646 chains with length ranging from 3 to 5, each chains has correctly aligned image with the lowest level concept.

Sample Construction. For each type of question, we generate different options using concept chains and taxonomy, including hard negative options. For atomic concept understanding, we randomly assign 50% of concepts as “Yes” paired with correct images, and 50% with incorrect images from unrelated concepts. In concept abstraction reasoning, for each chain of length N , we generate $N - 1$ questions with options set to current-level concept, lower-level concept, positive upper-level concept,

Subset	Train	Valid	Test	Total
Atomic Concept Understanding	2,168	313	637	3,118
Concept Abstraction Reasoning	1,738	247	490	2,475
Concept Concretization Reasoning	1,717	246	489	3,210
Common Ancestor Reasoning	1,611	225	448	2,284
Total	7,234	1,031	2,064	10,329

Table 2: Statistics of CONSTRUCTURE.

and negative upper-level concept (in random order). Similarly, in concept concretization reasoning, options include current-level concept, upper-level concept, positive lower-level concept, and negative lower-level concept. For common ancestor reasoning in chains of length N , options cover current-level concept, random-level concept, and brother or brother-son concept (as the answer option), arranged randomly. We conduct rigorous programmatic and manual checks to eliminate unreasonable options that might lead to multiple correct answers.

Statistics. As shown in Table 2, our CONSTRUCTURE benchmark has totally 10,329 samples. We split the benchmark into training, validation, and test sets in a ratio of 7:1:2, totaling 7234, 1031, and 2064 samples, respectively. The test set is exclusively used for evaluating MLLMs, while the training and validation sets are employed for fine-tuning MLLMs to enhance their conceptual structure reasoning capabilities. The breakdown of question types is detailed in Table 2. We further elaborate on the details of the CONSTRUCTURE benchmark in Table 4. Our benchmark comprises 1315 visual concepts, 646 concept chains, and 644 images, with an average concept chain length per sample of 4.93. The breakdown of subsets is detailed in Table 4. As shown in Figure 5, our CONSTRUCTURE benchmark covers a wide range of topics, including "animal", "plant", "device", "clothing", "chemicals", and more.

4 Experiments

4.1 Experimental setups

Multimodal Large Language Models. We evaluate six api-based Multimodal Large Language Models (MLLMs) and eight open-source MLLMs. All open-source MLLMs are \sim 7B models. We conduct zero-shot setups in overall evaluation for all MLLMs across our CONSTRUCTURE benchmark. Since our benchmark includes both Chinese and English versions, we evaluate MLLMs with Chinese capabilities in a Chinese context and those without in an English context.

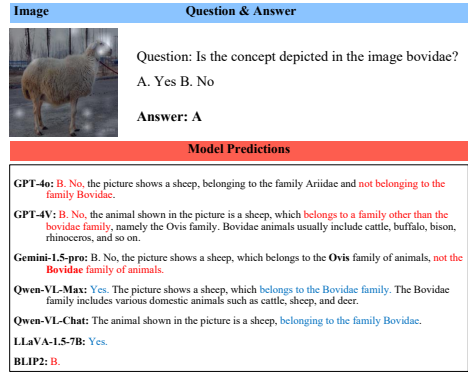


Figure 3: Case study of Atomic Concept Understanding Task.

Evaluation Methods. In the CONSTRUCTURE benchmark, test samples are organized in a multiple-choice format. The questions are divided into different sets: atomic concept understanding questions, which include "true" and "false" as the two options, and three other reasoning question sets, each offering four challenging options. Only one of these options is the correct answer. We input the image and the question to query MLLMs to get a response. Following (Huang et al., 2024), we then employ regular expression matching to parse the model outputs and calculate the accuracy by comparing it with the ground truth.

4.2 Main Results

Overall Results. We first present the overall results of the evaluated models on our CONSTRUCTURE benchmark as shown in Table 3. Current MLLMs, whether open-source or proprietary, demonstrate some level of concept understanding, but their ability to comprehend conceptual structures remains limited. The top-performing model, GPT-4o-0513, achieved an average score of only 0.621, with even lower average scores of 0.498 across the three conceptual structure reasoning tasks. There is significant variability in performance across different reasoning tasks, and all MLLMs struggle particularly with Common Ancestor Reasoning, indicating its ongoing challenges. Overall, GPT-4o-0513 and GPT-4-vision-preview stand out among API-based models. Although API-based models generally outperform open-source models, open-source models like BLIP2 and Qwen-VL-Chat achieve competitive results, surpassing models such as Gemini. This indicates that open-source approaches are capable of producing high-performing models in certain areas.

Model	Concept Understanding	Abstraction Reasoning	Concretization Reasoning	Common Ancestor	Avg. Score
API-based Models					
gpt-4-vision-preview	0.692	0.748	0.601	0.241	0.586
gpt-4o-0513	0.896	0.657	0.663	0.145	0.621
gemini-pro-vision	0.733	0.584	0.486	0.040	0.489
gemini-1.5-pro	0.746	0.692	0.407	0.100	0.513
claude3_sonnet	0.666	0.571	0.501	0.397	0.546
Qwen-VL-Max	0.763	0.543	0.631	0.143	0.545
Open-source Models					
BLIP2	0.794	0.484	0.442	0.326	0.535
InstructBLIP	0.0	0.061	0.143	0.069	0.063
MiniGPT-4	0.455	0.257	0.186	0.172	0.283
mPLUG_Owl	0.480	0.192	0.225	0.194	0.289
VisualGLM	0.281	0.027	0.045	0.016	0.107
Chinese_LLaVA	0.532	0.531	0.303	0.201	0.406
LLaVA-1.5	0.670	0.347	0.337	0.246	0.422
Qwen-VL-Chat	0.794	0.316	0.543	0.245	0.502

Table 3: Evaluation Results on CONSTRUCTURE Benchmark.

Results on Atomic Concept Understanding.

MLLMs generally perform well on the atomic concept understanding task. Among the API-based models, GPT-4o achieves nearly 90% accuracy, indicating a robust understanding of visual concepts across different levels. In the open-source models, both BLIP2 and Qwen-VL-Chat achieve a score of 0.794, surpassing most API-based models and only slightly behind GPT-4o. We observe that models performing well in the atomic concept understanding task also tend to excel in the other three reasoning tasks, demonstrating a positive correlation between atomic concept understanding and concept structure reasoning performance. InstructBLIP, however, performs poorly across most tasks, frequently outputting incorrect answers or gibberish. VisualGLM also struggles, often disregarding task instructions and merely generating descriptions of the images. Figure 3 below illustrates the performance of different MLLMs on an atomic concept understanding task. In the example, GPT-4o, GPT-4V, and Gemini-1.5-pro correctly identified the image as a sheep but incorrectly answered that a sheep is not a bovine animal (when, in fact, it is). However, Qwen-VL-Max and Qwen-VL-Chat correctly identified that a sheep belongs to the bovine family. This indicates that not all MLLMs possess comprehensive knowledge of concept structures.

Results on Concept Abstraction Reasoning. In the concept abstraction reasoning task, GPT-4V achieved the highest score of 0.748, followed by Gemini-1.5-pro with 0.657. Among open-source models, Chinese_LLaVA performed the best, scoring 0.531. Figure 6 presents the responses of different MLLMs to a specific question in this task. In this example, GPT-4V, Qwen-VL-Chat, LLaVA-

1.5, and BLIP2 all provided correct answers. However, both GPT-4o and Gemini-1.5-pro answered incorrectly, misunderstanding the concept abstraction and hierarchical relationship (*i.e.* waterbirds include both swimming birds and wading birds).

Results on Concept Concretization Reasoning.

In the concept concretization reasoning task, GPT-4o once again achieved the highest score of 0.663, followed by Qwen-VL-Max with a score of 0.631. Among the open-source models, Qwen-VL-Chat performed the best with a score of 0.543, with BLIP2 coming in second at 0.442. Figure 7 illustrates the performance of different MLLMs on a specific question in this task. In this example, Gemini-1.5-Pro and BLIP2 answered correctly. Gemini-1.5-Pro provided the correct reasoning process, while the other models answered incorrectly: GPT-4o gave an incorrect response without engaging in reasoning, GPT-4v misidentified the concept in the image (it’s a Mahi Mahi or dolphinfish, not a saury pike), and Qwen-VL-Max selected an overly broad option due to conservative answering.

Results on Common Ancestor Reasoning. In the common ancestor reasoning task, all MLLMs struggled, indicating it as the most challenging task for MLLMs in concept structure reasoning. The best performer was claude3_sonnet, achieving a score of 0.397. GPT-4V and GPT-4o scored 0.241 and 0.145, respectively. Among the open-source models, BLIP2 scored 0.326, surpassing both GPT-4V and GPT-4o but falling short of claude3_sonnet. Figure 8 illustrates the performance of different MLLMs on a specific question in this task. In this example, the image represents the structural formula of a chemical compound, acetamide. The

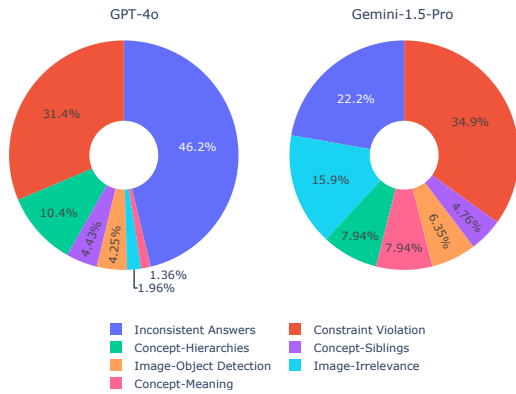


Figure 4: Distribution of Error Types for GPT-4o and Gemini-1.5-Pro.

question requires the model to adhere to two constraints: 1) it must differ from the concept shown in the image, and 2) it must also be a “compound”. Only Qwen-VL-Chat provided the correct answer in the example. GPT-4o and Gemini-1.5-pro correctly identified that option B’s “solution” is a mixture rather than a “compound” (a pure substance), but they provided an incorrect answer in the final step of reasoning, revealing inconsistency in their reasoning process. GPT-4v gave a hallucinative answer, and its reasoning process was also incorrect. Qwen-VL-Max’s answer did not meet the first requirement. Therefore, we observe that MLLMs perform poorly in tasks like Common Ancestor Reasoning due to various reasons, such as errors in understanding images or concepts, inconsistency in generation or reasoning processes, and others. Next, we will further analyze the types of model errors.

4.3 Analysis

Error Type Analysis. To further analyze the error type distribution of MLLMs, we examined two models: GPT-4o and Gemini-1.5-Pro. As shown in Figure 4, we finally categorize errors into seven types, focusing on the capabilities of MLLMs regarding images, concepts, and language. The error types are as follows: 1)Image-Irrelevance: This error occurs when models provide answers that do not take the content of the provided images into account. 2)Image-Object Detection: This refers to models failing to recognize key objects or misidentifying them. For example, a cow might be incorrectly detected as an antelope. 3)Concept-Hierarchies: models have difficulties in understanding hierarchical relationships between concepts and identifying which one is in a higher level.

4) Concept-Meaning: models fail to comprehend the meanings of options because they don’t know which concepts include the objects in the images. 5) Concept-Siblings: models have difficulties in recognizing sibling relationships between concepts. For example, the models doesn’t recognize that "antelopes" and "yaks", which are hyponyms of "bovine animals", are sibling nodes, distinct from "deer". 6) Constraint Violation: models have difficulties in understanding the constraints of a given task. Most of errors in common ancestor reasoning fall into this category. 7) Inconsistent Answers: models produce conflicting answers or multiple answers inside the analysis of answers, unable to reach a definitive conclusion.

Analyzing the Impact of Concept Abstraction Levels. We analyze the impact of conceptual abstraction levels on model performance in understanding conceptual structures. We find that as the level of concept abstraction increases, the model performance declines. See detailed analyses in appendix A.3.

4.4 CoT Reasoning and Finetuning

As we know, Chain-of-Thoughts (CoT) prompting and fine-tuning are two methods capable of enhancing the reasoning abilities of LLMs. Therefore, in this section, we conducted experiments aimed at enhancing the conceptual structure reasoning of MLLMs using CoT prompting and fine-tuning. Specifically, we focus on the last three reasoning tasks in our CONSTRUCTURE benchmark. We will start by presenting the overall results, followed by a detailed discussion of the results for each experimental setting.

Overall Results. As shown in Table 5, across the three concept structure reasoning tasks, GPT-4V emerged as the top performer among all baselines, achieving an average score of 0.537, followed by GPT-4o at 0.498. We conducted separate experiments with Zero-shot CoT prompting and Few-shot CoT prompting based on GPT-4o. The results revealed that Few-shot CoT significantly enhanced GPT-4o’s concept structure reasoning capabilities, increasing its average score from 0.498 to 0.699—an improvement of over 20 points. This improvement surpassed GPT-4v by more than 16 points. Notably, in the Common Ancestor Reasoning task, the score rose from 0.145 to 0.529, marking an impressive increase of nearly 38 percentage points. Conversely, Zero-shot CoT did not yield significant

540	improvements.		
541	Furthermore, through fine-tuning on the con-		
542	structed training and validation sets, we achieved		
543	the most substantial performance enhancement		
544	with the open-source MLLM Qwen-VL-Chat,		
545	reaching a score of 0.74. This performance sur-		
546	passed even the best baseline, GPT-4V, which		
547	scored 0.537.		
548	Zero-shot CoT. We evaluate and report on two em-		
549	pirically derived Zero-shot CoT prompts. Our ex-		
550	periments reveal that these prompts do not enhance		
551	model performance on conceptual structure reason-		
552	ing tasks. Specifically, we randomly sample 200		
553	error cases from the GPT-4o conceptual structure		
554	reasoning tasks and apply five empirically designed		
555	Zero-shot CoT prompts.(see detailed prompt in		
556	Appendix B). Based on the error correction rate,		
557	we select the two best-performing Zero-shot CoT		
558	prompts for evaluation on the full test set and report		
559	the results. As shown in Table 5, although these		
560	Zero-shot CoT prompts achieve up to a 30% er-		
561	ror correction rate in the sampled error cases, they		
562	lead to a performance decline when applied to the		
563	full test set, with scores dropping from 0.498 to		
564	0.479 and 0.451, respectively. This indicates that		
565	while these prompts slightly improve performance		
566	on error cases, particularly in the Common An-		
567	cestor reasoning task (from 0.145 to 0.165), they		
568	negatively impact Concretization Reasoning (from		
569	0.662 to 0.575 and 0.495), leading to an overall		
570	decline in performance. This suggests that prompts		
571	like “Let’s think step by step.” do not significantly		
572	help with concept structure reasoning problems.		
573	Few-shot CoT. We use Few-shot CoT prompts		
574	written by human experts for each question type		
575	(detailed prompts can be found in the Appendix B).		
576	The human-designed Few-shot CoT samples thor-		
577	oughly address issues identified in previous error		
578	analyses, such as GPT-4o not following question		
579	instructions and having inconsistent reasoning pro-		
580	cesses. The CoT prompts instruct the model to		
581	answer questions according to the given instruc-		
582	tions and to analyze each option one by one. In		
583	the samples, the prompts first break down the mul-		
584	tiiple constraints of the question and then analyze		
585	whether each option meets these constraints. By		
586	guiding GPT-4o through this reasoning paradigm,		
587	its ability to solve conceptual structure reasoning		
588	problems is significantly enhanced. This demon-		
589	strates that following a proper reasoning process		
590	and possessing strong reasoning capabilities are		
	crucial for improving the model’s concept structure		591
	reasoning.		592
	Fine-tuning. Additionally, we conduct fine-tuning		593
	experiments on the Qwen-VL-Chat model using		594
	our training and validation sets. The fine-tuned		595
	model, based on the specialized concept structure		596
	reasoning dataset, achieves a top performance score		597
	of 0.740. To further understand the impact of		598
	atomic concept understanding data, we perform an		599
	ablation study by removing the atomic concept un-		600
	derstanding subset from the training data and train-		601
	ing for the same number of epochs. This results in		602
	a decrease in model performance, demonstrating		603
	that the inclusion of atomic concept understanding		604
	data is beneficial for enhancing the model’s concept		605
	structure reasoning abilities.		606
	Discussion. We summarize three key insights from		607
	experiments on enhancing MLLMs’ concept struc-		608
	ture reasoning abilities: 1) Zero-CoT prompting		609
	does not significantly improve MLLMs’ perfor-		610
	mance in concept structure reasoning. This may		611
	be because step-by-step reasoning prompts can in-		612
	crease the likelihood of the model focusing on er-		613
	roneous interference items, leading to reasoning		614
	errors. 2) Few-shot CoT prompting that integrates		615
	human expert reasoning process priors can signif-		616
	icantly enhance MLLMs’ concept structure rea-		617
	soning abilities. 3) Fine-tuning based on concept		618
	structure reasoning data can significantly improve		619
	MLLMs’ performance, demonstrating the signifi-		620
	cant value of our proposed concept structure rea-		621
	soning dataset.		622
	5 Conclusion		623
	We introduce the CONSTRUCTURE benchmark,		624
	which evaluates MLLMs’ cognitive and reasoning		625
	abilities in tasks like atomic concept understanding,		626
	concept abstraction reasoning, concept concretiza-		627
	tion reasoning, and common ancestor reasoning.		628
	Our findings highlight significant challenges in		629
	concept structure reasoning for MLLMs. The top-		630
	performing model, GPT-4o, achieved an average		631
	score of 0.621, indicating room for improvement.		632
	We summarize current evaluations of MLLMs in		633
	concept structure reasoning, analyze reasons for		634
	their underperformance, and provide key insights		635
	from experiments using CoT prompting and fine-		636
	tuning to enhance their abilities. Our discoveries		637
	offer crucial guidance for advancing MLLMs’ cog-		638
	nitve capabilities in concept structure reasoning.		639

640 Limitation

641 Since our concept chains are based on a Chinese
642 taxonomy, there may be some language bias during
643 translation into English. Despite extensive manual
644 checks, our data annotations might still contain a
645 few inaccuracies due to errors in the raw data, in-
646 fluenced by the annotators’ understanding of the
647 correct answers. Additionally, we primarily eval-
648 uate the capabilities of MLLMs within a Chinese
649 context. Given the varying proficiency of different
650 models in Chinese and English, the results may
651 exhibit some variations.

652 References

653 Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang,
654 Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou,
655 and Jingren Zhou. 2023. Qwen-vl: A frontier large
656 vision-language model with versatile abilities. *arXiv
657 preprint arXiv:2308.12966*.

658 Sijie Cheng, Zhicheng Guo, Jingwen Wu, Kechen Fang,
659 Peng Li, Huaping Liu, and Yang Liu. 2023. Can
660 vision-language models think from a first-person per-
661 spective? *arXiv preprint arXiv:2311.15596*.

662 Wenliang Dai, Junnan Li, Dongxu Li, Anthony
663 Meng Huat Tiong, Junqi Zhao, Weisheng Wang,
664 Boyang Li, Pascale N Fung, and Steven Hoi.
665 2024. Instructblip: Towards general-purpose vision-
666 language models with instruction tuning. *Advances
667 in Neural Information Processing Systems*, 36.

668 Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding,
669 Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. Glm:
670 General language model pretraining with autoregres-
671 sive blank infilling. In *Proc. of ACL*.

672 Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin,
673 Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jin-
674 rui Yang, Xiawu Zheng, Ke Li, Xing Sun, Rongrong
675 Ji, and TencentYoutu Lab. Mme: A comprehensive
676 evaluation benchmark for multimodal large language
677 models.

678 Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei
679 Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu,
680 Chuan Cheng Lv, Yikai Zhang, Yao Fu, et al. 2024.
681 C-eval: A multi-level multi-discipline chinese evalua-
682 tion suite for foundation models. *Advances in Neural
683 Information Processing Systems*, 36.

684 Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yix-
685 iao Ge, and Ying Shan. 2023a. Seed-bench: Bench-
686 marking multimodal llms with generative compre-
687 hension. *arXiv preprint arXiv:2307.16125*.

688 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi.
689 2023b. Blip-2: Bootstrapping language-image pre-
690 training with frozen image encoders and large lan-
691 guage models. In *International conference on ma-
692 chine learning*, pages 19730–19742. PMLR.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae
Lee. 2024. Visual instruction tuning. *Advances in
neural information processing systems*, 36. 693
694
695

Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li,
Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi
Wang, Conghui He, Ziwei Liu, et al. 2023. Mm-
bench: Is your multi-modal model an all-around
player? *arXiv preprint arXiv:2307.06281*. 696
697
698
699
700

Douglas L Medin. 1989. Concepts and conceptual struc-
ture. *American psychologist*, 44(12):1469. 701
702

OpenAI. 2023. Gpt-4 technical report. *ArXiv*. 703

Siddharth Suresh, Kushin Mukherjee, Xizheng Yu, Wei-
Chun Huang, Lisa Padua, and Timothy Rogers. 2023.
Conceptual structure coheres in human cognition but
not in large language models. In *Proceedings of the
2023 Conference on Empirical Methods in Natural
Language Processing*, pages 722–738. 704
705
706
707
708
709

John R. Taylor. 2019. 2. [prototype theory](#). In Claudia
Maienborn, Klaus Heusinger, and Paul Portner, edi-
tors, *Semantics - Theories*, pages 29–56. De Gruyter
Mouton, Berlin, Boston. 710
711
712
713

Shan Wang and Francis Bond. 2013. Building the chi-
nese open wordnet (cow): Starting from core synsets.
In *Proceedings of the 11th Workshop on Asian Lan-
guage Resources*, pages 10–18. 714
715
716
717

Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao,
Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang,
Yu Qiao, and Ping Luo. 2023. Lvlm-ehub: A com-
prehensive evaluation benchmark for large vision-
language models. *arXiv preprint arXiv:2306.09265*. 718
719
720
721
722

Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng
Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan
Wang. 2023. The dawn of llms: Preliminary
explorations with gpt-4v (ision). *arXiv preprint
arXiv:2309.17421*, 9(1):1. 723
724
725
726
727

Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye,
Ming Yan, Yiyang Zhou, Junyang Wang, An-
wen Hu, Pengcheng Shi, Yaya Shi, et al. 2023.
mplug-owl: Modularization empowers large lan-
guage models with multimodality. *arXiv preprint
arXiv:2304.14178*. 728
729
730
731
732
733

Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang,
Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan
Wang. 2023. Mm-vet: Evaluating large multimodal
models for integrated capabilities. *arXiv preprint
arXiv:2308.02490*. 734
735
736
737
738

Zhiwei Zha, Jiaan Wang, Zhixu Li, Xiangru Zhu, Wei
Song, and Yanghua Xiao. 2023. M2conceptbase: A
fine-grained aligned multi-modal conceptual knowl-
edge base. *arXiv preprint arXiv:2312.10417*. 739
740
741
742

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and
Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing
vision-language understanding with advanced large
language models. *arXiv preprint arXiv:2304.10592*. 743
744
745
746




Figure 9: Level-wise Accuracy of Api-based MLLMs.

and general concept. A fox belongs to the Canidae family, so Canid is the more abstract concept., the correct answer is D. Canid."

Few CoT Prompt for Concept Concretization Reasoning.

"This is an image of a humpback whale. Please select from the following options the one that correctly describes the concept in the image and is the most specific and accurate concept.. Sirenians . Cetaceans . Aquatic mammals . Baleen whales: Let's analyze each option according to the prompt requirements. First, the prompt requires 1) selecting an option that correctly describes the concept in the image, and 2) is the most specific and accurate concept. Next, we analyze each option:.. Sirenians - The concept in the image is a humpback whale, not a sirenian, so this option does not meet the first requirement.. Cetaceans - The concept in the image is a humpback whale, which is a cetacean, so this option meets the first requirement.. Aquatic mammals - The concept in the image is a humpback whale, which is an aquatic mammal, so this option meets the first requirement.. Baleen whales - The concept in the image is a humpback whale, which is a type of baleen whale, so this option meets the first requirement. Now, between options B. Cetaceans, C. Aquatic mammals, and D. Baleen whales, the most specific concept is Baleen whales., the correct

answer is D. Baleen whales."

Few CoT Prompt for Common Ancestor Reasoning.

"This is an image of a butterflyfish. Please select from the following options the concept that is different from the image but belongs to the same 'Perciformes' group.. Sturgeon . Surgeonfish . Crocodile . Perciformes: Let's analyze each option according to the prompt requirements. First, the prompt requires selecting a concept that 1) is different from the image, and 2) belongs to the 'Perciformes' group. Next, we analyze each option:.. Sturgeon - The concept in the image is a butterflyfish, which is different from a sturgeon, so this option meets the first requirement, but sturgeon belongs to the order Acipenseriformes, not Perciformes, so this option does not meet the second requirement.. Surgeonfish - The concept in the image is a butterflyfish, not a surgeonfish, so this option meets the first requirement. Surgeonfish belong to the order Acanthuriformes, which is not Perciformes, so this option does not meet the second requirement.. Crocodile - Crocodiles are reptiles and do not belong to Perciformes, so this option does not meet the second requirement.. Perciformes - The concept in the image is a butterflyfish, which belongs to the Perciformes order, so this option meets the second requirement. Now, between options A. Sturgeon and D. Perciformes,




Figure 10: Level-wise Accuracy of Open-source MLLMs.

890 *we select the concept that is different but belongs*
 891 *to the same group, which is D. Perciformes., the*
 892 *correct answer is B. Surgeonfish."*