

ECBench: Can Multi-modal Foundation Models Understand the Egocentric World? A Holistic Embodied Cognition Benchmark

Ronghao Dang^{1*}, Yuqian Yuan^{2*}, Wenqi Zhang^{2*}, Yifei Xin¹, Boqiang Zhang¹,
Long Li², Liuyi Wang³, Qinyang Zeng³, Xin Li^{1†}, Lidong Bing⁴

¹Alibaba DAMO Academy, ²Zhejiang University, ³Tongji University, ⁴Shanda AI Research Institute

dangronghao.drh@alibaba-inc.com

Abstract

The enhancement of generalization in robots by large vision-language models (LVLMs) is increasingly evident. Therefore, the embodied cognitive abilities of LVLMs based on egocentric videos are of great interest. However, current datasets for embodied video question answering lack comprehensive and systematic evaluation frameworks. Critical embodied cognitive issues, such as robotic self-cognition, dynamic scene perception, and hallucination, are rarely addressed. To tackle these challenges, we propose ECBench, a high-quality benchmark designed to systematically evaluate the embodied cognitive abilities of LVLMs. ECBench features a diverse range of scene video sources, open and varied question formats, and 30 dimensions of embodied cognition. To ensure quality, balance, and high visual dependence, ECBench uses class-independent meticulous human annotation and multi-round question screening strategies. Additionally, we introduce ECEval, a comprehensive evaluation system that ensures the fairness and rationality of the indicators. Utilizing ECBench, we conduct extensive evaluations of proprietary, open-source, and task-specific LVLMs. ECBench is pivotal in advancing the embodied cognitive capabilities of LVLMs, laying a solid foundation for developing reliable core models for embodied agents. All data and code is available at <https://github.com/Rh-Dang/ECBench>.

1. Introduction

Spurred by the significant advancements in large language models (LLMs), AI researchers are now more optimistic than ever about achieving artificial general intelligence. Many pioneering studies aim to cognize the physical world by developing models capable of understanding multimodal inputs, such as MiniGPT-4 [42], Video-LLaMA [40], etc.

However, these models are typically trained and evaluated on non-egocentric inputs. Therefore, it is crucial to examine the barriers to their application in embodied environments.

A comprehensive and reliable benchmark is essential for evaluating the multimodal understanding capabilities of models in egocentric embodied contexts. Such benchmarks are very scarce due to the challenges involved in egocentric video collection and embodied questions annotation. ScanQA [1] and SQA3D [26] are representative close-vocabulary benchmarks for 3D scene understanding. These task-specific datasets are insufficient to comprehensively assess task-generalized LVLMs. Therefore, OpenEQA [27] introduces the first open-vocabulary benchmark to assess foundational models' question-answering capabilities in indoor embodied videos. Despite these efforts, the current evaluation of LVLMs in embodied scenarios has the following limitations:

- **Not Systematic:** Current benchmarks focus on independent embodied abilities, such as object recognition and counting. They lack a comprehensive top-down analysis of embodied cognition requirements, leading to limitations in evaluation hierarchy and dimensions.
- **Lack of Robot-Centric:** As illustrated by the robot-centric example in Fig. 1, robots often need to address questions related to their own embodiment, such as the distance to a target, or their historical trajectory. However, benchmarks like OpenEQA focus solely on third-person scenario questions, significantly overlooking the evaluation of robots' self-awareness.
- **Lack of Dynamics:** In the real physical world, scene dynamics are perpetually ongoing, as exemplified by video content changes on the screen in the dynamic scene example of Fig. 1. For complex tasks like "Revert the screen content to before you faced the whiteboard," a robot must recognize these dynamics and accurately recall their timing and process. However, current embodied question answering benchmarks typically overlook these dynamic aspects, defaulting to static context assumptions.
- **Hallucination issue:** Although the hallucination phe-

*Equal contribution.

†Corresponding author.

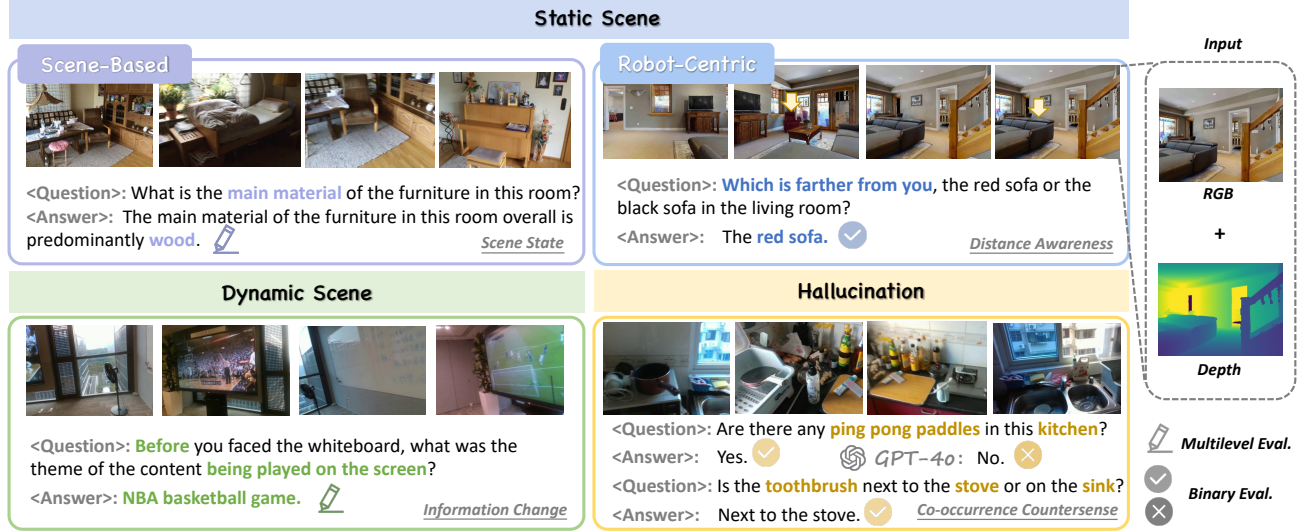


Figure 1. **Illustration of question answering (QA) format and representative cognitive dimensions from ECBench.** There are 386 RGB-D videos, 4,324 QA pairs, and 30 distinct embodied cognitive abilities, spanning across various aspects such as perception, reasoning, self-awareness, dynamic capturing, and hallucination. ECEval employs distinct evaluation methods for different types of answers.

nomenon has been extensively analyzed within LVLMs, embodied-based question answering presents unique hallucination challenges. For instance, as depicted in the hallucination example of Fig. 1, LVLMs like GPT-4o often rely too much on common sense in the scene when answering questions in counterintuitive scenes, resulting in incorrect answers. These embodied hallucination issues remain unexplored in the academic literature.

To address the aforementioned challenges, we propose ECBench, a comprehensive video question answering benchmark designed to evaluate the embodied cognitive abilities of LVLMs. ECBench covers three sets: static scenes, dynamic scenes, and hallucinations. Within the static scene category, we systematically categorize questions into scene-based and robot-centric types, encompassing 19 cognitive abilities such as spatial reasoning and trajectory review. Notably, we innovatively introduce robot-centric cognitive questions, aiming for the models to develop self-awareness and understand the relationship between their own entity and the environment. Robot-centric cognition is crucial when robots frequently move and perform complex tasks, such as searching and transporting.

Additionally, we propose the first open-world question answering task for dynamic scenes, encompassing dynamics both within and beyond immediate visibility. To highlight the model’s capability in perceiving scene dynamics, all questions are related to changes within the scene. We focus on four categories: spatial dynamics, information dynamics, quantity dynamics, and state dynamics.

We further evaluate hallucination issues specific to embodied scenes from seven detailed perspectives, grouped into two aspects: over-confidence in common sense and

over-confidence in user input. Over-confidence in common sense impairs LVLMs’ ability to handle counterintuitive scenarios; for instance, they might struggle to accept the visual fact of a toothbrush being next to a stove (Fig. 1). Over-confidence in user input hinders LVLMs’ ability to recognize ambiguous, incorrect, or missing references in user input, substantially reducing the robot’s interactivity during task execution.

During dataset collection, to fulfill the evaluation requirements for different aspects of embodied cognition, we permit annotators to flexibly employ various question formats, including open-ended and multiple-choice questions. For evaluation, we propose ECEval, a novel assessment framework for question-answering tasks in embodied contexts. ECEval amalgamates binary scoring and multi-level scoring, enabling accurate evaluation of answer correctness for both open-ended and closed-ended questions (Fig. 1). Uniquely, we execute a multi-round screening strategy supported by GPT-4o blind testing to minimize the proportion of questions that can be answered solely based on common sense. Overall, the benchmark encompasses 4,324 unique question-answering (QA) pairs provided by volunteers, covering a total of 30 fine-grained evaluation angles.

Based on ECBench, we comprehensively evaluate the embodied cognitive abilities of various LVLMs. Our assessment encompasses open-source and proprietary general LVLMs [13, 25, 38], as well as embodied / egocentric LVLMs [3, 35]. Notably, all mainstream LVLMs exhibit poor performance in dynamic scenes and embodied hallucination issues. Furthermore, in static scenes, robot-centric questions (first-person) present greater challenges compared to scene-based questions (third-person). These

Benchmark	Visual Input			Question Answering Pairs			Evaluation Domain			
	Real Scenes	RGB-D	Egocentric	Open Vocab.	Generation	QA Form	Number of Capabilities	Robot Centric	Halluc.	Dynamic Scene
ScanQA [1]	✓	✓	✓		Automatic/Human	Open	—			
SQA3D [26]	✓	✓	✓		Human	Open	—	✓		
Env-QA [10]		✓	✓		Template	Open	5			✓
RoboVQA [34]	✓		✓	✓	Automatic	Open	8			✓
OpenEQA [27]	✓	✓	✓	✓	Human	Open	7			
MMBench-Video [9]	✓			✓	Human	Open	26		✓	✓
MVBench [20]	✓			✓	Template/Automatic	Close	20			✓
ECBench(Ours)	✓	✓	✓	✓	Human	Open/Close	30	✓	✓	✓

Table 1. **Comparing ECBench and widely adopted Embodied / General VideoQA benchmarks.** ECBench has significant advantages in terms of quality, diversity and evaluation dimensions.

results indicate that contemporary LVLMs only possess third-person cognition in static scenes, while they struggle to achieve first-person understanding in dynamic scenes. We aim for ECBench to drive the development of LVLMs toward more practical embodied scenarios.

2. Related Work

We compare ECBench with popular general and embodied video question answering (VideoQA) benchmarks in Tab. 1.

2.1. General VideoQA Benchmarks

VideoQA technology promotes the model understanding of the world. Initially, popular evaluation benchmarks concentrate on specific domains such as movies [36], TV series [11, 18], and video games [30]. However, with the advent of LVLMs, the scope of model applicability has broadened. To evaluate these LVLMs in all aspects, comprehensive benchmarks [9, 19, 20, 28] have emerged. MMBench-video [9] uses a carefully constructed capability classification. MVBench [20] introduces a novel framework of spatial-temporal task construction. Nonetheless, general VideoQA benchmarks focus on YouTube videos and seldom incorporate embodied-specific QA formats. Our ECBench maintains the comprehensive and systematic nature of general benchmarks while far exceeding them in embodied cognition (e.g. 3D spatial understanding, robot self-awareness).

2.2. Embodied VideoQA Benchmarks

In embodied scenarios, VideoQA-based evaluations can effectively gauge a model’s understanding of its environment and tasks. ScanQA [1], SQA3D [26], and Env-QA [10] are typical datasets for traditional scene question answering with a closed vocabulary. These datasets have a strong text bias and relatively limited question forms. RoboVQA [34], EgoPlan-Bench [4], and PCA-Bench [2] focus on testing task-planning abilities of LVLMs but lack direct evaluation of cognitive capabilities. OpenEQA [27] represents the first open vocabulary benchmark in the realm of embodied

video understanding. However, it lacks robot-centric questions and dynamic scenes. ECBench is the first work to systematically analyze the embodied cognition of LVLMs, allowing researchers to explore the frontiers of LVLMs in robots.

3. ECBench

In this section, we introduce the construction and detailed statistics of ECBench.

3.1. Benchmark Construction

Video Collection Traditional embodied question answering videos typically involve the comprehensive scanning of 3D scenes. Alternatively, some datasets, like OpenEQA and Env-QA, create more diverse videos by manually simulating motion trajectories in virtual scenes [8, 15]. In contrast, ECBench enhances video diversity by integrating virtual scenes with real scenes and combining static scenes with dynamic scenes. Although current LVLMs can’t process depth information, ECBench still includes depth maps that are crucial for embodied tasks.

Firstly, we employ an open-world object navigation agent [16] and an active embodied question-answering agent [33] to capture robotic authentic video streams in the HM3D environment [32]. These videos exhibit increased instances of hesitation and pauses compared to those manually gathered by humans, thereby offering a more realistic representation of robotic perception during task execution. Secondly, we carefully select 191 real scan videos from ScanNet [7] and MultiScan [29]. Lastly, utilizing the Intel RealSense depth camera, we collect videos of counterintuitive scenes and dynamic scenes from the real world, providing a data basis for evaluating hallucination problems and dynamic perception.

Capability Taxonomy Motivated by the popular general VideoQA benchmarks [9, 20, 28, 41], we implement an enhanced systematic approach to categorize various embod-

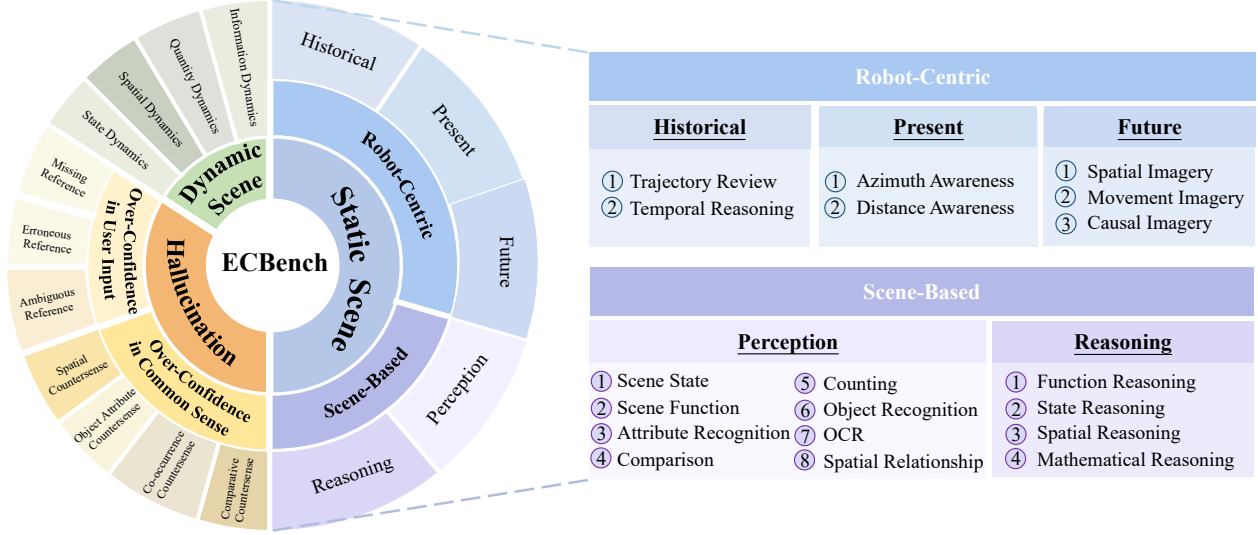


Figure 2. **Overview of embodied cognition dimensions in ECBench.** ECBench includes three subsets: static scenes, dynamic scenes, and hallucination, evaluating a total of 30 embodied cognitive abilities.

ied cognitive abilities (Fig. 2). ECBench is categorized into three primary sets: static scene QA, dynamic scene QA, and hallucination evaluation.

Static scene QA employs a hierarchical classification system divided into three levels of cognitive abilities. At the L1 level, there are two higher-order cognitive functions: scene-based cognition and robot-centric cognition. Scene-based cognition is concerned with understanding scene information and can be further divided into perception and reasoning. In contrast, robot-centric cognition necessitates a level of self-awareness in the model, enabling it to understand its position and impact within the physical environment. We divide robot-centric cognition into three parts: historical cognition, present cognition, and future cognition. Overall, static scene QA encompasses 19 distinct cognitive abilities, comprehensively covering the cognitive demands of robots in static environments. The detailed classification is presented in Fig. 2.

Dynamic scene QA can be divided into four categories according to their dynamic characteristics: spatial, state, information, and quantity dynamics. We design videos and questions specifically tailored to highlight these dynamic features. Emphasis is also placed on the temporal and spatial dynamics within robot movement in our questions.

In embodied scenarios, hallucinations are not merely a bias under 2D vision but are also closely related to the model’s over-confidence in common sense and user input in 3D scenarios. We propose four categories of counter-intuitive questions (Fig. 2) to evaluate LVLMS’ commonsense hallucinations. Furthermore, over-confidence in user input can inhibit the model’s ability to identify ambiguous references, missing references, and erroneous references in user instructions, significantly impairing the effectiveness

of human-robot interaction.

Examples of each category are presented in Appendix 6.

Construction of Question-Answer Pairs Current benchmarks for embodied VideoQA utilize two forms of QA pair construction. One is to use template sentences and object annotations within 3D environments to automatically create standard QA pairs, which tend to have a limited textual form [10]. Another is to use LVLMS like GPT-4 to generate QA pairs, supplemented by automatic or manual filtering mechanisms [34]. However, these generated QA pairs only rely on single-frame information, lacking the ability to evaluate the model’s cognitive ability to entire 3D scenes and dynamic events. Therefore, we employ manual annotation methods to construct ECBench.

ECBench contains 30 cognitive abilities, which will result in a significant category imbalance if traditional video-by-video question-annotation methods are used. Therefore, we implement a category-independent annotation approach by assigning volunteers a predetermined number of tasks related to each cognitive ability. This method ensures a balanced distribution between the quantity of QA pairs associated with both rare and common abilities.

After completing all data labeling, we conduct a filtering process for commonsense questions. It is widely recognized that LVLMS possess extensive prior knowledge in common sense but often struggle to align with visual reality [14, 24, 39]. Our objective is to reduce the number of questions that rely solely on common sense, allowing ECBench to concentrate on visual cognition. Specifically, we employ GPT-4o to answer all the questions six times without visual inputs, followed by manual revisions of questions consistently answered correctly to enhance their reliance on vi-




Figure 3. **Data analysis of ECBench** reflects a rich diversity of scenario categories, video sources, and evaluation dimensions.

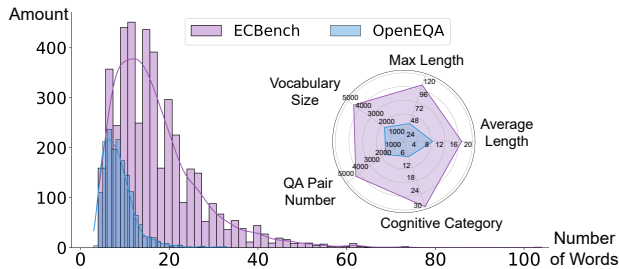


Figure 4. **Comparison with OpenEQA [27] on textual data**, including the distribution of question lengths, average question length, maximum question length, vocabulary size, number of questions, and number of capabilities.

sual input. This process is iterated thrice. Ultimately, we perform stringent cross-validation on all questions to guarantee the accuracy of QA pairs. The final ECBench includes static scene QA, dynamic scene QA, and hallucination QA, systematically evaluating LVLMs from various embodied cognitive abilities.

Evaluation Framework In ECBench, most answers are close-ended and explicit, while some open-ended answers involve an element of ambiguity. To address this issue, we propose ECEval, a combined evaluation system integrating both multi-level and binary judgments. We first partition ECBench into two parts: open-ended and close-ended. For the close-ended part, we prompt GPT-4o to assign a straightforward binary score of either 0 or 1. For the open-ended part, we manually annotate a *0.5-point answer* for each open-ended question. In the evaluation, open-ended answers are scored by GPT-4o on a scale from 0 to 1 in increments of 0.2, based on *0.5-point answers*. This evaluation framework facilitates more precise scoring for questions with definitive answers and provides a more nuanced scoring methodology for open-ended inquiries. Further details can be found in Appendix 7.

3.2. Dataset Statistics

Video Source As illustrated in Fig. 3b, ECBench comprises 386 RGBD videos exploring indoor scenes, sourced

from ScanNet [7], MultiScan [29], HM3D [32], and self-collections. We select 140 and 51 real-world videos from the ScanNet and MultiScan datasets, respectively, chosen based on the richness of scene information. Additionally, ECBench includes 111 first-person perspective videos of agents performing tasks in the HM3D virtual environment, with 49 videos originating from the Embodied Question Answering Agent [33] and 62 videos from the Object Navigation Agent [16]. We manually collect scenes that are less common in public datasets, featuring 44 counterintuitive scenes and 40 dynamic scenes. As depicted in Fig. 3a, each video’s scene category has been manually annotated. Overall, ECBench comprises 12 distinct indoor scenes. Beyond common scenes like bedrooms and kitchens, ECBench also includes rare scenes such as stairways and recreation rooms, providing a comprehensive evaluation of the model’s adaptability across diverse environments.

QA Pairs As shown in Fig. 3c, ECBench comprises 4,324 QA pairs, systematically evaluating LVLMs from 30 cognitive perspectives. In previous benchmarks, such as OpenEQA [27], the QA pairs only include a limited range of ability categories. ECBench not only provides a more systematic classification and analysis of embodied cognitive abilities but also introduces 534 hallucination QA pairs and 248 dynamic scene QA pairs. As illustrated in Fig. 4, the average question length in ECBench is 16.88, which is twice that of OpenEQA. ECBench includes more complex and diverse questions and contains a vocabulary of 4,513 words, which is 2.6 times larger than OpenEQA. Further dataset analysis is presented in Appendix 8.

4. Experiment

Based on ECBench, we conduct extensive evaluations of native Video-LVLMs, open-source and proprietary Image-LVLMs, as well as embodied and egocentric LVLMs. For the open-source models, we utilize their default hyperparameters for inference. Further experimental details and prompts are presented in Appendix 9.

Model	Overall Mean	Static Scene			Dynamic Scene					Hallucination		
		SB	RC	Mean	ID	QD	SPD	STD	Mean	UI	CS	Mean
Blind LLMs												
GPT-4o	24.09	25.32	33.06	28.26	4.76	0.00	10.00	13.81	8.55	1.83	5.36	3.56
Native Video-LVLMs												
Video-LLaMA2-7B-[16f] [6]	33.87	41.23	33.52	38.30	16.51	5.71	16.97	11.19	13.31	4.03	24.29	13.93
Video-LLaVA-7B-[8f] [23]	35.25	41.21	37.25	39.71	18.41	14.29	19.20	16.67	17.66	1.10	26.97	13.75
Kangaroo-8B-[64f] [25]	36.23	40.79	40.01	40.49	19.20	17.14	16.97	15.95	17.42	0.37	33.49	16.55
LongVA-7B-[384f] [21]	40.47	49.03	38.56	45.05	28.57	8.57	13.94	16.67	17.82	8.06	33.49	20.49
Open-Source Image-LVLMs												
Idefics3-8B-[8f] [17]	32.76	36.31	36.49	36.38	11.43	8.57	13.64	14.53	12.66	5.86	30.73	18.01
Idefics3-8B-[20f]	33.35	36.56	38.09	37.14	18.10	8.57	10.61	19.29	15.16	4.76	28.89	16.55
InternVL2-40B-[8f] [5]	35.41	40.34	38.16	39.51	30.16	2.86	15.15	16.90	17.82	2.93	30.19	16.25
InternVL2-40B-[20f]	35.94	41.27	38.09	40.06	33.33	0.00	12.12	21.66	19.03	3.66	29.58	16.33
Qwen2VL-7B-[8f] [38]	37.28	43.23	38.31	41.36	23.81	5.71	19.67	16.90	17.82	5.49	33.33	19.10
Qwen2VL-7B-[20f]	39.57	46.21	39.33	43.60	32.70	8.57	19.70	18.33	20.97	4.40	39.08	21.35
Qwen2VL-72B-[8f]	41.57	48.28	43.52	46.47	33.65	8.57	18.18	16.43	20.16	4.76	33.56	18.84
Qwen2VL-72B-[20f]	44.62	52.40	43.95	49.19	37.14	11.43	18.18	24.05	24.03	5.86	42.38	23.71
Proprietary Image-LVLMs												
GPT-4v-[8f] [31]	39.16	45.09	40.16	43.22	29.52	8.57	22.12	20.24	21.45	9.16	32.11	20.37
GPT-4o-mini-[8f] [13]	39.71	46.00	40.35	43.86	28.25	17.14	16.97	19.05	20.56	9.16	33.26	20.94
GPT-4o-mini-[32f]	43.69	51.20	43.12	48.13	26.35	20.00	15.15	18.10	19.68	9.89	41.38	25.28
GPT-4o-[8f] [13]	45.96	52.95	45.91	50.27	36.83	11.43	21.21	23.33	24.52	9.52	45.67	27.19
GPT-4o-[32f]	50.35	59.74	49.04	55.06	25.71	14.29	22.73	24.04	22.74	7.69	57.01	31.80
Embodied / Egocentric LVLMs												
GeLM-7B-[180f] [3]	21.54	25.72	23.70	24.95	5.08	5.71	8.18	3.51	5.48	0.37	12.49	6.29
AlanaVLM-7B-[64f] [35]	34.75	40.38	36.61	38.95	19.05	5.71	19.70	14.76	15.89	1.83	29.96	15.58
Human	94.96	97.45	93.21	95.86	91.77	97.21	94.32	96.25	94.55	82.74	96.91	89.21

Table 2. **Main evaluation results of various LVLMs on ECBench.** We evaluate in three major categories: Static Scenes, Dynamic Scenes, and Hallucinations. SB and RC represent Scene-Based and Robot-Centric, respectively. ID, QD, SPD, and STD stand for Information Dynamics, Quantity Dynamics, Spatial Dynamics, and State Dynamics, respectively. UI and CS represent Over-Confidence in User Input and Over-Confidence in Common Sense, respectively. -[Nf] indicates the method take N frames uniformly sampled from a video as input.

4.1. Main Results

Blind LLM In the first row of Tab. 2, we conduct a blind evaluation of ECBench using GPT-4o [13] without the video input. This score reflects the extent to which reliance solely on the common sense and guessing ability of LVLMs can achieve. ECBench includes 90.31% close-ended questions to ensure the accuracy of the evaluation. Despite this, the blind evaluation score for ECBench reaches only 24.09, which is a 28% reduction compared to the entirely open-ended OpenEQA (33.5). This indicates that ECBench significantly outperforms OpenEQA in terms of evaluation precision and visual dependency.

Native Video-LVLMs We assess popular general and long context native Video-LVLMs. The long video models are capable of incorporating more frames, which is critical for embodied videos with high dynamic views. Consequently, the SOTA model in the domain of long video understanding, LongVA-[384f] [21], achieves 40.47, making it the best among open-source Video-LVLMs. However, in dynamic scenes, models with more input frames do not

show a significant improvement. This suggests that, Video-LVLMs fundamentally lacks the ability to capture and understand dynamic scenes.

Open-Source Image-LVLMs Qwen2VL [38] exhibited impressive performance, achieving 44.16, surpassing GPT-4o-mini. Additionally, we observe that as model capabilities improved, the gains from increasing the number of input frames became more pronounced. When the input is expanded from 8 frames to 20 frames, the Qwen2VL-72B score increases by 3.05, whereas the performance of InternVL2-40B [5] remains virtually unchanged.

Proprietary Image-LVLMs The GPT series models are widely regarded as the most capable proprietary models for multimodal understanding. We evaluate several models from this series, including GPT-4v [31], GPT-4o-mini, and GPT-4o. GPT-4o-[32f] exhibits the strongest embodied cognitive abilities, achieving 50.35. This score is 12.8% higher than the best open-source Image-LVLMs and 24.4% higher than the best native Video-LVLMs.

Model	Overall Mean	Perception										Reasoning				
		SS	SF	AR	COM	COU	OR	OCR	SR	Mean		FR	STR	SPR	MR	Mean
AlanaVLM-7B-[64f]	40.38	52.60	59.37	46.21	39.08	28.75	52.21	7.84	34.68	40.44		55.46	49.46	39.54	21.49	40.28
LongVA-7B-[384f]	49.03	71.10	74.05	62.61	51.15	32.19	60.39	31.76	43.17	52.34		61.84	56.39	40.00	22.41	43.73
Qwen2VL-72B-[8f]	48.28	64.74	75.57	67.45	36.78	38.59	55.97	19.61	42.30	50.34		43.14	55.41	42.27	36.35	44.98
Qwen2VL-72B-[20f]	52.40	64.05	74.30	71.68	35.06	48.69	65.58	43.73	37.99	55.05		51.68	60.20	48.66	34.77	48.16
GPT-4o-[8f]	52.95	65.20	74.68	72.05	59.77	43.30	65.58	24.51	47.48	56.80		48.43	51.02	47.00	41.66	46.77
GPT-4o-[32f]	59.74	67.98	77.97	75.03	61.49	54.21	72.21	49.02	47.91	63.14		50.81	54.21	49.49	49.54	51.70

Table 3. **Evaluation of sub-abilities under the Scene-Based (SB) category in the static scene set.** SS: Scene State, SF: Scene Function, AR: Attribute Recognition, COM: Comparison, COU: Counting, OR: Object Recognition, SR: Spatial Relationship. FR: Function Reasoning, STR: State Reasoning, SPR: Spatial Reasoning, MR: Mathematical Reasoning.

Model	Overall Mean	Historical			Present			Future			
		TJR	TPR	Mean	AA	DA	Mean	SI	MI	CI	Mean
AlanaVLM-7B-[64f]	36.61	40.09	37.14	38.92	33.61	34.36	33.96	38.65	33.72	39.30	37.09
LongVA-7B-[384f]	38.56	40.47	36.43	38.86	37.81	41.03	39.32	35.27	36.52	42.57	37.85
Qwen2VL-72B-[8f]	43.52	43.96	50.00	46.36	34.70	43.08	38.65	43.67	43.67	49.12	45.26
Qwen2VL-72B-[20f]	43.95	44.34	48.57	46.02	39.00	41.54	40.19	43.86	45.12	47.49	45.37
GPT-4o-[8f]	45.91	47.26	54.29	50.06	37.44	46.67	41.79	42.03	46.09	51.81	46.32
GPT-4o-[32f]	49.04	52.17	53.57	52.73	38.17	51.28	44.35	47.92	47.63	55.91	50.15

Table 4. **Evaluation of sub-abilities under the Robot-Centric (RC) category in the static scene set.** TJR: Trajectory Review, TPR: Temporal Reasoning, AA: Azimuth Awareness, DA: Distance Awareness, SI: Spatial Imagery, MI: Movement Imagery, CI: Causal Imagery.

Embodied / Egocentric LVLMS Currently, there are several LVLMS designed to handle multimodal understanding in embodied and egocentric contexts. We evaluate two notable models: AlanaVLM [35] and GeLM [3]. Unfortunately, their performance in embodied cognition is underwhelming. GeLM, which is mainly trained using egocentric videos obtained from XR devices, faced a significant domain gap when applied to robotic videos, resulting in a low score of 21.54. AlanaVLM performs reasonably well on the OpenEQA benchmark, but it scores only 34.75 on ECBench. These findings suggest that ECBench offers more thorough evaluation dimensions, objectively reflecting the embodied cognitive capabilities of LVLMS.

4.2. Evaluation in Static Scene

Questions in static scenes are divided into two categories: Scene-Based and Robot-Centric. Scene-Based questions concentrate on assessing the LVLMS’ scene comprehension from a third-person viewpoint, whereas Robot-Centric questions focus on the relationship between the robot itself and the environment from a first-person viewpoint.

Scene-Based Questions Table 3 illustrates the ability of various models to answer Scene-Based questions. Among the eight categories of Perception questions, Counting, OCR, and Spatial Relationship are the areas where current LVLMS are least proficient. These deficiencies hinder robots from accurately recognizing the intricacies of their surrounding environments when completing tasks. For GPT-4o-[32f], the scores for Attribute Recognition and Object Recognition are 75.03 and 72.21, respectively. This

suggests that the advanced LVLMS can provide crucial cognition abilities for robots to navigate and manipulate individual objects. Among the four categories of Reasoning questions, Spatial Reasoning and Mathematical Reasoning provide greater challenges for the models. GPT-4o attains scores of 49.49 and 49.54 on these two difficult reasoning problems, clearly surpassing other models. LongVA and Qwen2VL excel in Function Reasoning and State Reasoning, achieving scores of 61.84 and 60.20, respectively. Overall, Reasoning problems impose higher demands on model capabilities compared to Perception questions.

Robot-Centric Questions Table 4 demonstrates the ability of various models to answer Robot-Centric questions, which are categorized as Historical, Present, and Future. As an embodied LVLMS, AlanaVLM achieves a score of only 36.61 in the Robot-Centric cognition, significantly lagging behind the SOTA generalist LVLMS. Current embodied LVLMS do not acknowledge the significance of self-awareness in robotics. LongVA-7B-[384f], as a model designed for long video understanding, performs exceptionally well in the Present category, with only a 0.87 gap from Qwen2VL-72B. The richer temporal visual information is crucial for robots to construct scene models and develop self-awareness (e.g. position and orientation). Notably, the number of theoretically answerable questions for robot-centric questions significantly increases with the enhancement of the input frame count. However, it is only in powerful proprietary models like GPT-4o that we can clearly observe the benefits of increasing input frame numbers. Thus,

Model	UI			CS			
	MRF	ERF	ARF	SC	OAC	COC	CC
AlanaVLM	1.69	3.41	0.00	28.33	24.09	28.96	43.24
LongVA	11.02	5.68	5.97	25.95	41.36	32.71	43.24
Qwen2VL	11.86	2.27	0.00	47.14	38.18	40.83	40.54
GPT-4o	14.41	4.55	0.00	45.95	55.45	66.46	59.46

Table 5. **Evaluation of sub-abilities in the Hallucination set.** MRF, ERF, and ARF represent Missing, Erroneous, and Ambiguous Reference. SC, OAC, COC, and CC represent Spatial, Object Attribute, Co-occurrence, and Comparative Countersense.

current LVLMs urgently need to improve their ability to capture egocentric temporal information.

4.3. Evaluation in Dynamic Scene

All Dynamic Scene questions involve changes occurring within the scene. Notably, numerous changes occur outside the robot’s immediate line of sight. LVLMs need to remember and compare scene information to effectively answer these questions. These questions are categorized into four types: Information Dynamics, Quantity Dynamics, Spatial Dynamics, and State Dynamics. As shown in Table 2, Quantity Dynamics presents the greatest challenge for LVLMs, with InternVL2-40B-[20f] even scoring zero on this task. Overall, current models show a complete inability to perceive the dynamic elements in the scene. However, for robots, the environment is constantly changing while they execute their tasks, which raises concerns about how to enable LVLMs to fully cognize these dynamics.

4.4. Evaluation of Hallucinatory Phenomena

The primary hallucination issues when completing embodied tasks are categorized into two types: Over-Confidence in User Input and Over-Confidence in Common Sense. We hope that LVLMs can recognize erroneous or ambiguous user inputs and correct them through subsequent dialogue. However, in the UI (User Input) category of Table 5, none of the LVLMs manages to handle this situation. In contrast, some questions in the CS (Common Sense) category are addressed with relatively good accuracy. Notably, GPT-4o achieves a score of 66.46 on co-occurrence countersense questions. This commendable performance is largely attributed to ongoing efforts in general multimodal hallucination research [12, 22, 37] regarding these problems. We hope that ECBench will draw the attention of more researchers to embodied hallucination issues and promote the development of more robust embodied cognitive LVLMs.

4.5. ECEval Evaluation Method

Currently, there are two mainstream evaluation methods for question-answering tasks in open-world settings. One is binary scoring, the other is multi-level scoring. In Fig. 5, we emphasize the advantages of our ECEval method relative

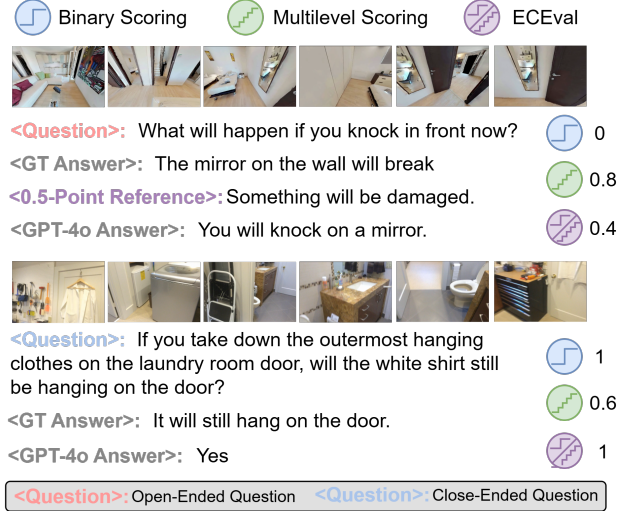


Figure 5. **Comparison of results between ECEval, Binary Scoring, and Multilevel Scoring**, for open-ended and closed-ended questions. Notably, only open-ended questions are annotated with *0.5-point answers*.

to these two typical evaluation methods. As shown in the open-ended example in Fig. 5, binary scoring fails to sensitively indicate changes in model abilities. Alternatively, the multi-level scoring without reference may result in inaccurate ratings due to GPT4o biases. ECEval has *0.5-point answers* as references for open-ended questions, allowing GPT4o to provide more accurate scoring. As shown in the close-ended example in Fig. 5, while the GPT4o answer differs in phrasing from the standard answer, the semantic content is entirely consistent. However, using multi-level scoring would prevent awarding full points. ECEval employs a binary evaluation mode for close-ended questions to avoid such issues. Thus, in a comprehensive benchmark like ECBench, relying solely on binary or multi-level evaluation modes is inherently flawed. Our proposed ECEval synthesizes the strengths of these evaluation methods, yielding more adaptable and precise evaluation outcomes.

5. Conclusion

This work proposes ECBench, a holistic open-world benchmark for embodied cognition. ECBench comprises three test sets: static scene, dynamic scene, and hallucination, enabling a detailed assessment of LVLMs’ capabilities across 30 embodied cognition tasks. In addition, we introduce ECEval, which ensures the accuracy and smoothness of the ECBench scoring. Extensive evaluations on ECBench reveal that various LVLMs struggle to perform well on common embodied situations, such as robot-centric, dynamic scenes and instruction ambiguity. We hope that ECBench will facilitate the advancement of LVLMs towards a more complex and diverse physical world.

Acknowledgments

This paper is supported by the National Natural Science Foundation of China under Grants (624B2105).

References

- [1] Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. Scanqa: 3d question answering for spatial scene understanding. In *CVPR*, pages 19129–19139, 2022. 1, 3, 10
- [2] Liang Chen, Yichi Zhang, Shuhuai Ren, Haozhe Zhao, Zefan Cai, Yuchi Wang, Peiyi Wang, Xiangdi Meng, Tianyu Liu, and Baobao Chang. Pca-bench: Evaluating multimodal large language models in perception-cognition-action chain. In *ACL*, pages 1086–1104. Association for Computational Linguistics, 2024. 3
- [3] Qirui Chen, Shangzhe Di, and Weidi Xie. Grounded multi-hop videoqa in long-form egocentric videos. *arXiv preprint arXiv:2408.14469*, 2024. 2, 6, 7
- [4] Yi Chen, Yuying Ge, Yixiao Ge, Mingyu Ding, Bohao Li, Rui Wang, Ruifeng Xu, Ying Shan, and Xihui Liu. Egoplan-bench: Benchmarking egocentric embodied planning with multimodal large language models. *arXiv preprint arXiv:2312.06722*, 2023. 3
- [5] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*, 2024. 6
- [6] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024. 6
- [7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas A. Funkhouser, and Matthias Nießner. Scan-net: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 2432–2443. IEEE Computer Society, 2017. 3, 5
- [8] Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Winson Han, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Procthor: Large-scale embodied AI using procedural generation. In *NeurIPS*, 2022. 3
- [9] Xinyu Fang, Kangrui Mao, Haodong Duan, Xiangyu Zhao, Yining Li, Dahua Lin, and Kai Chen. Mmbench-video: A long-form multi-shot benchmark for holistic video understanding. *arXiv preprint arXiv:2406.14515*, 2024. 3
- [10] Difei Gao, Ruiping Wang, Ziyi Bai, and Xilin Chen. Env-qa: A video question answering benchmark for comprehensive understanding of dynamic environments. In *ICCV*, pages 1655–1665. IEEE, 2021. 3, 4
- [11] Noa Garcia, Mayu Otani, Chenhui Chu, and Yuta Nakashima. Knowit VQA: answering knowledge-based questions about videos. In *AAAI*, pages 10826–10834. AAAI Press, 2020. 3
- [12] Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, et al. Hallusionbench: an advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *CVPR*, pages 14375–14385, 2024. 8
- [13] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024. 2, 6
- [14] Prannay Kaul, Zhizhong Li, Hao Yang, Yonatan Dukler, Ashwin Swaminathan, CJ Taylor, and Stefano Soatto. Throne: An object-based hallucination benchmark for the free-form generations of large vision-language models. In *CVPR*, pages 27228–27238, 2024. 4
- [15] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017. 3
- [16] Yuxuan Kuang, Hai Lin, and Meng Jiang. Openfmnav: Towards open-set zero-shot object navigation via vision-language foundation models. In *Findings of the Association for Computational Linguistics: NAACL 2024, Mexico City, Mexico, June 16-21, 2024*, pages 338–351. ACL, 2024. 3, 5
- [17] Hugo Laurençon, Andrés Marafioti, Victor Sanh, and Léo Tronchon. Building and better understanding vision-language models: insights and future directions. *arXiv preprint arXiv:2408.12637*, 2024. 6
- [18] Jie Lei, Licheng Yu, Mohit Bansal, and Tamara L. Berg. TVQA: localized, compositional video question answering. In *EMNLP*, pages 1369–1379. Association for Computational Linguistics, 2018. 3
- [19] Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. Seed-bench: Benchmarking multimodal large language models. In *CVPR*, pages 13299–13308. IEEE, 2024. 3
- [20] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Lou, Limin Wang, and Yu Qiao. Mvbench: A comprehensive multimodal video understanding benchmark. In *CVPR*, pages 22195–22206. IEEE, 2024. 3
- [21] Lei Li, Yuanxin Liu, Linli Yao, Peiyuan Zhang, Chenxin An, Lean Wang, Xu Sun, Lingpeng Kong, and Qi Liu. Temporal reasoning transfer from text to video. *arXiv preprint arXiv:2410.06166*, 2024. 6
- [22] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023. 8
- [23] Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection, 2024. 6
- [24] Jiazhen Liu, Yuhao Fu, Ruobing Xie, Runquan Xie, Xingwu Sun, Fengzong Lian, Zhanhui Kang, and Xirong Li. Phd: A prompted visual hallucination evaluation dataset. *arXiv preprint arXiv:2403.11116*, 2024. 4

- [25] Jiajun Liu, Yibing Wang, Hanghang Ma, Xiaoping Wu, Xiaoli Ma, Xiaoming Wei, Jianbin Jiao, Enhua Wu, and Jie Hu. Kangaroo: A powerful video-language model supporting long-context video input. *arXiv preprint arXiv:2408.15542*, 2024. 2, 6
- [26] Xiaojian Ma, Silong Yong, Zilong Zheng, Qing Li, Yitao Liang, Song-Chun Zhu, and Siyuan Huang. SQA3D: situated question answering in 3d scenes. In *ICLR*, 2023. 1, 3
- [27] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul McVay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent-Pierre Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Alexander Sax, and Aravind Rajeswaran. Openeqa: Embodied question answering in the era of foundation models. In *CVPR*, pages 16488–16498. IEEE, 2024. 1, 3, 5, 10
- [28] Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. Egoschema: A diagnostic benchmark for very long-form video language understanding. In *NeurIPS*, 2023. 3
- [29] Yongsan Mao, Yiming Zhang, Hanxiao Jiang, Angel X. Chang, and Manolis Savva. Multiscan: Scalable RGBD scanning for 3d environments with articulated objects. In *NeurIPS*, 2022. 3, 5
- [30] Jonghwan Mun, Paul Hongsuck Seo, Ilchae Jung, and Bohyung Han. Marioqa: Answering questions by watching gameplay videos. In *ICCV*, pages 2886–2894. IEEE Computer Society, 2017. 3
- [31] R OpenAI. Gpt-4 technical report. arxiv 2303.08774. *View in Article*, 2(5), 2023. 6
- [32] Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M. Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X. Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-matterport 3d dataset (HM3D): 1000 large-scale 3d environments for embodied AI. In *NeurIPS*, 2021. 3, 5
- [33] Allen Z Ren, Jaden Clark, Anushri Dixit, Masha Itkina, Anirudha Majumdar, and Dorsa Sadigh. Explore until confident: Efficient exploration for embodied question answering. *arXiv preprint arXiv:2403.15941*, 2024. 3, 5
- [34] Pierre Sermanet, Tianli Ding, Jeffrey Zhao, Fei Xia, Debidda Dwivedi, Keerthana Gopalakrishnan, Christine Chan, Gabriel Dulac-Arnold, Sharath Maddineni, Nikhil J. Joshi, Pete Florence, Wei Han, Robert Baruch, Yao Lu, Suvir Mirchandani, Peng Xu, Pannag Sanketi, Karol Hausman, Izhak Shafran, Brian Ichter, and Yuan Cao. Robovqa: Multimodal long-horizon reasoning for robotics. In *ICRA*, pages 645–652. IEEE, 2024. 3, 4
- [35] Alessandro Suglia, Claudio Greco, Katie Baker, Jose L Part, Ioannis Papaionnou, Arash Eshghi, Ioannis Konstas, and Oliver Lemon. Alanavlm: A multimodal embodied ai foundation model for egocentric video understanding. *arXiv preprint arXiv:2406.13807*, 2024. 2, 6, 7
- [36] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. Movieqa: Understanding stories in movies through question-answering. In *CVPR*, pages 4631–4640. IEEE Computer Society, 2016. 3
- [37] Lei Wang, Jiabang He, Shenshen Li, Ning Liu, and Ee-Peng Lim. Mitigating fine-grained hallucination by fine-tuning large vision-language models with caption rewrites. In *ACMMM*, pages 32–45. Springer, 2024. 8
- [38] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 2, 6
- [39] Mingrui Wu, Jiayi Ji, Oucheng Huang, Jiale Li, Yuhang Wu, Xiaoshuai Sun, and Rongrong Ji. Evaluating and analyzing relationship hallucinations in lvlms. *arXiv preprint arXiv:2406.16449*, 2024. 4
- [40] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. In *EMNLP*, pages 543–553. Association for Computational Linguistics, 2023. 1
- [41] Zicheng Zhang, Ziheng Jia, Haoning Wu, Chunyi Li, Zijian Chen, Yingjie Zhou, Wei Sun, Xiaohong Liu, Xiongkuo Min, Weisi Lin, et al. Q-bench-video: Benchmarking the video quality understanding of llms. *arXiv preprint arXiv:2409.20063*, 2024. 3
- [42] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigpt-4: Enhancing vision-language understanding with advanced large language models. In *ICLR*, 2024. 1

ECBench: Can Multi-modal Foundation Models Understand the Egocentric World? A Holistic Embodied Cognition Benchmark

Supplementary Material

6. Capability Taxonomy Details

In Sec. 3.1 of the main paper and in Fig. 2, we present the process of constructing the embodied cognition dimensions for ECBench. In this section, we will elaborate on the design principles and specific details underlying each capability dimension evaluated within ECBench. Furthermore, we will provide numerous examples for reference.

6.1. Static Scene

In static scene datasets, there are two types of questions: scene-based and robot-centric.

6.1.1. Scene-Based

Examples of all types of problems within the scene-based category can be found in Fig. 6 and Fig. 7. We categorize scene-based questions into two types: perception and reasoning. Perception-type questions assess the capability of large vision-language models (LVLMs) to directly capture information from the environment, whereas reasoning-type questions require LVLMs to address more complex reasoning problems based on the environment information. Consequently, perception serves as the foundation for LVLMs in solving reasoning problems.

Perception

- **Comparison (Fig. 6):** Comparing certain information within a scene, such as the size of objects and their distances.
- **Scene Function (Fig. 6):** The comprehension of the overall function of a scene, such as meetings, classes, or bathing.
- **Spatial Relationship (Fig. 6):** The understanding of spatial relationships between multiple objects, including concepts such as entangle, coverage, and verticality.
- **Attribute Recognition (Fig. 6):** The identification of attributes of individual objects, such as color, condition, and shape.
- **Object Recognition (Fig. 6):** The acknowledgment of the existence and position of individual objects.
- **Scene State (Fig. 6):** The identification of the overall state of a scene, including factors like tidiness, brightness, and clutter.
- **OCR (Fig. 7):** The recognition and analysis of textual content present within a scene.
- **Counting (Fig. 7):** To identify the quantity of certain objects in a scene or to comprehend reference resolution within a sequence.

Reasoning

- **Mathematical Reasoning (Fig. 7):** This involves possessing fundamental mathematical knowledge and utilizing it to address mathematical computation problems related to the physical world.
- **Spatial Reasoning (Fig. 7):** This requires an initial understanding of the spatial relationships between objects within a scene, followed by tackling more complex problems that arise from these spatial relationships.
- **State Reasoning (Fig. 7):** This necessitates a foundational comprehension of the states of objects in a scene, followed by resolving more complex issues stemming from these states.
- **Function Reasoning (Fig. 7):** This entails first understanding the functions of objects within a scene, and subsequently addressing the more intricate problems that arise as a consequence of these functions.

6.1.2. Robot-Centric

Examples of all types of questions within the robot-centric category can be found in Fig. 8 and Fig. 9. It is important to note that robot-centric questions inherently assume the following fact: the subject of the video footage is a mobile robot. This robot is capable of moving forward and backward, rotating left and right, and tilting its head to observe upward and downward. All questions categorized as robot-centric pertain to the robot itself, rather than being solely answerable based on isolated scene information. Based on the temporal dimensions of questions focus, we categorize robot-centric questions into three types: historical, present, and future.

Historical

- **Temporal Reasoning (Fig. 8):** Questions related to temporal elements such as time points, time intervals, and the sequencing of events based on the historical actions of the robot.
- **Trajectory Review (Fig. 9):** Retrospective examination of the robot’s actions and path, addressing issues pertinent to the robot’s spatial movement.

Present

- **Azimuth Awareness (Fig. 8):** The understanding of the positional relationship between the robot and its environment based on the robot’s current location and orientation (as indicated by the last frame).

- **Distance Awareness (Fig. 8, Fig. 9):** The comprehension of the distance relationship between the robot and its environment, derived from the robot’s current position and orientation (as indicated by the last frame).

Future

- **Movement Imagery (Fig. 8):** Envisioning the robot’s future movements and the resultant changes in its spatial position.
- **Spatial Imagery (Fig. 9):** Imagining a hypothetical scenario in which the robot occupies a certain location, thereby analyzing the spatial relationships between the robot and its surrounding environment.
- **Causal Imagery (Fig. 9):** Imagining the potential impacts that certain actions taken by the robot would have on the environment.

6.2. Dynamic Scene

- **Quantity Dynamic (Fig. 10):** The increases, decreases, and transfers of the quantity of objects within an environment.
- **Spatial Dynamic (Fig. 10):** The changes in the spatial positions of individual objects, such as changes in direction, alterations in speed, and movement trajectories. The changes in the spatial relationships among multiple objects, such as parallelism and perpendicularity, degrees of overlap, and arrangements mode.
- **State Dynamic (Fig. 10):** The state changes of objects, such as the on/off status of a light, the fullness or emptiness of a bottle, and the folding or unfolding of items. The state changes of scenes, such as brightness and darkness, as well as orderliness versus disorder.
- **Information Dynamic (Fig. 10):** The dynamics of textual information on a two-dimensional plane, such as the content, color, and typography of the text. The dynamics in the content displayed on screens, such as video playback and content transitions.

6.3. Hallucination

Over-Confidence in Common Sense To investigate whether the current LVLMs exhibit hallucinations due to an excessive reliance on scene knowledge when handling embodied tasks, we specifically collect a set of counterfactual scenarios and pertinent questions to assess the LVLMs.

- **Spatial Countersense (Fig. 11):** This refers to instances where spatial relationships do not conform to conventional expectations, such as a mouse pad placed atop a mouse, or a weighing scale situated on a dining table.
- **Object Attribute Countersense (Fig. 11):** This involves objects possessing attributes or characteristics that are relatively uncommon, for instance, a basketball smaller than a mobile phone, or a purple toilet.
- **Co-occurrence Countersense (Fig. 11):** Unrelated objects are put together, and objects typically associated with one another are deliberately separated. For example, a pot and a comb do not frequently co-occur, whereas a toothbrush and a toothbrush holder commonly do.
- **Comparative Countersense (Fig. 11):** Unexpected outcomes in comparative assessments, such as a computer being closer to a toilet than a roll of toilet paper is, or a garbage bin being positioned at a height greater than that of a sink.

Over-Confidence in User Input In online dialogues, the trust that LVLMs place in user input is understandable. However, in physical world interactions mediated by robots, questioning user instructions and seeking clarification is fundamental to ensuring smooth human-robot interaction.

- **Missing Reference (Fig. 12):** When certain key referents in the user’s instructions are absent from the scene, the model should identify these missing references.
- **Erroneous Reference (Fig. 12):** When the user’s instructions contain erroneous references, the model should autonomously correct these errors and provide feedback to the user.
- **Ambiguous Reference (Fig. 12):** When certain referents in the user’s instructions correspond to multiple entities in the scene, the model should enumerate these referents and inquire with the user for clarification on the intended reference.

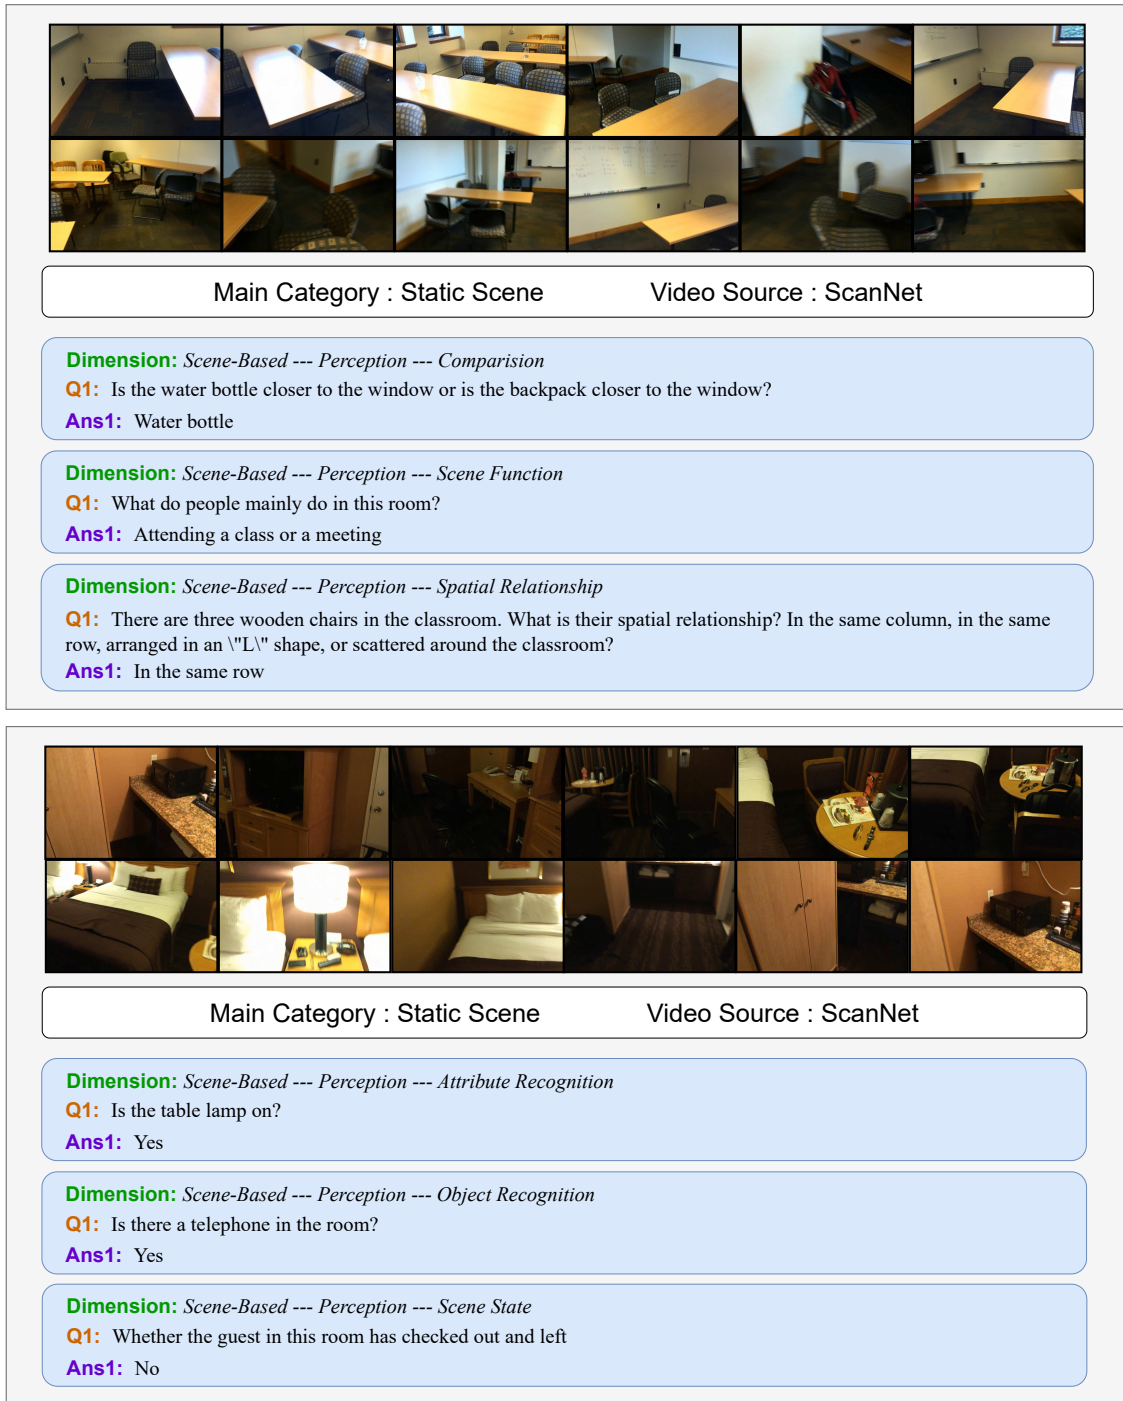


Figure 6. Visualization of question answering examples in the static scene test set. The main focus is on the scene-based category. Part 1 out of 2.

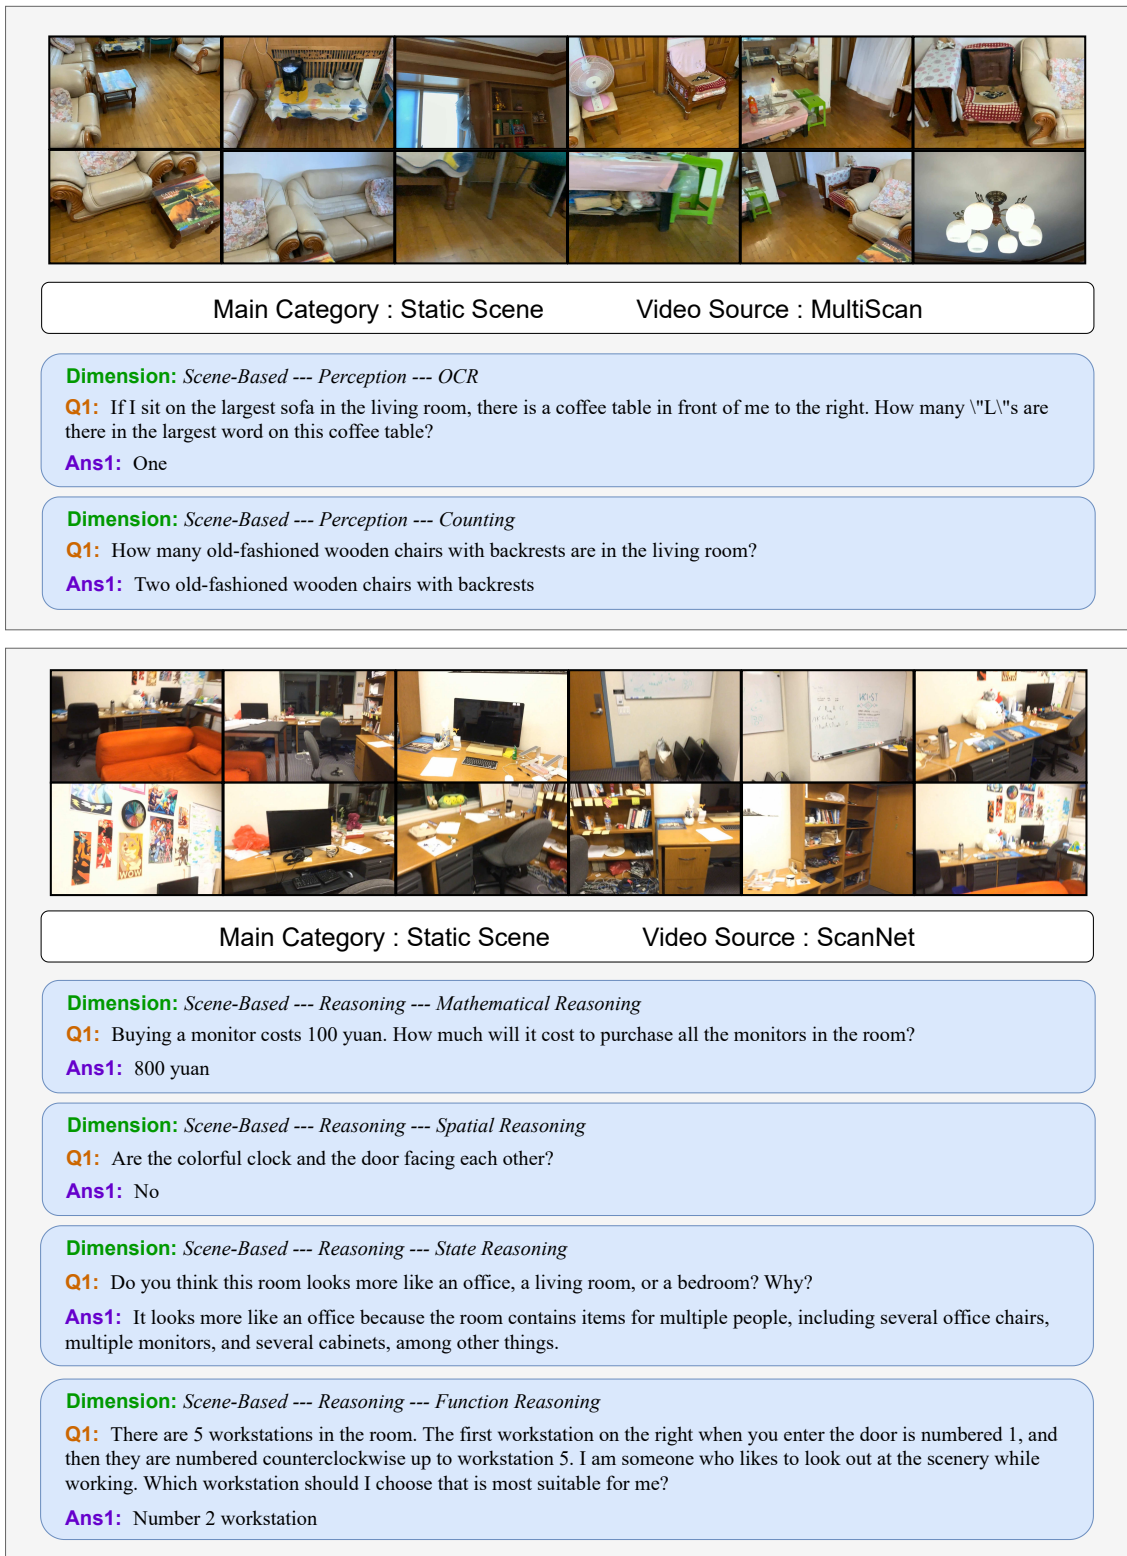


Figure 7. **Visualization of question answering examples in the static scene test set.** The main focus is on the scene-based category. Part 2 out of 2.

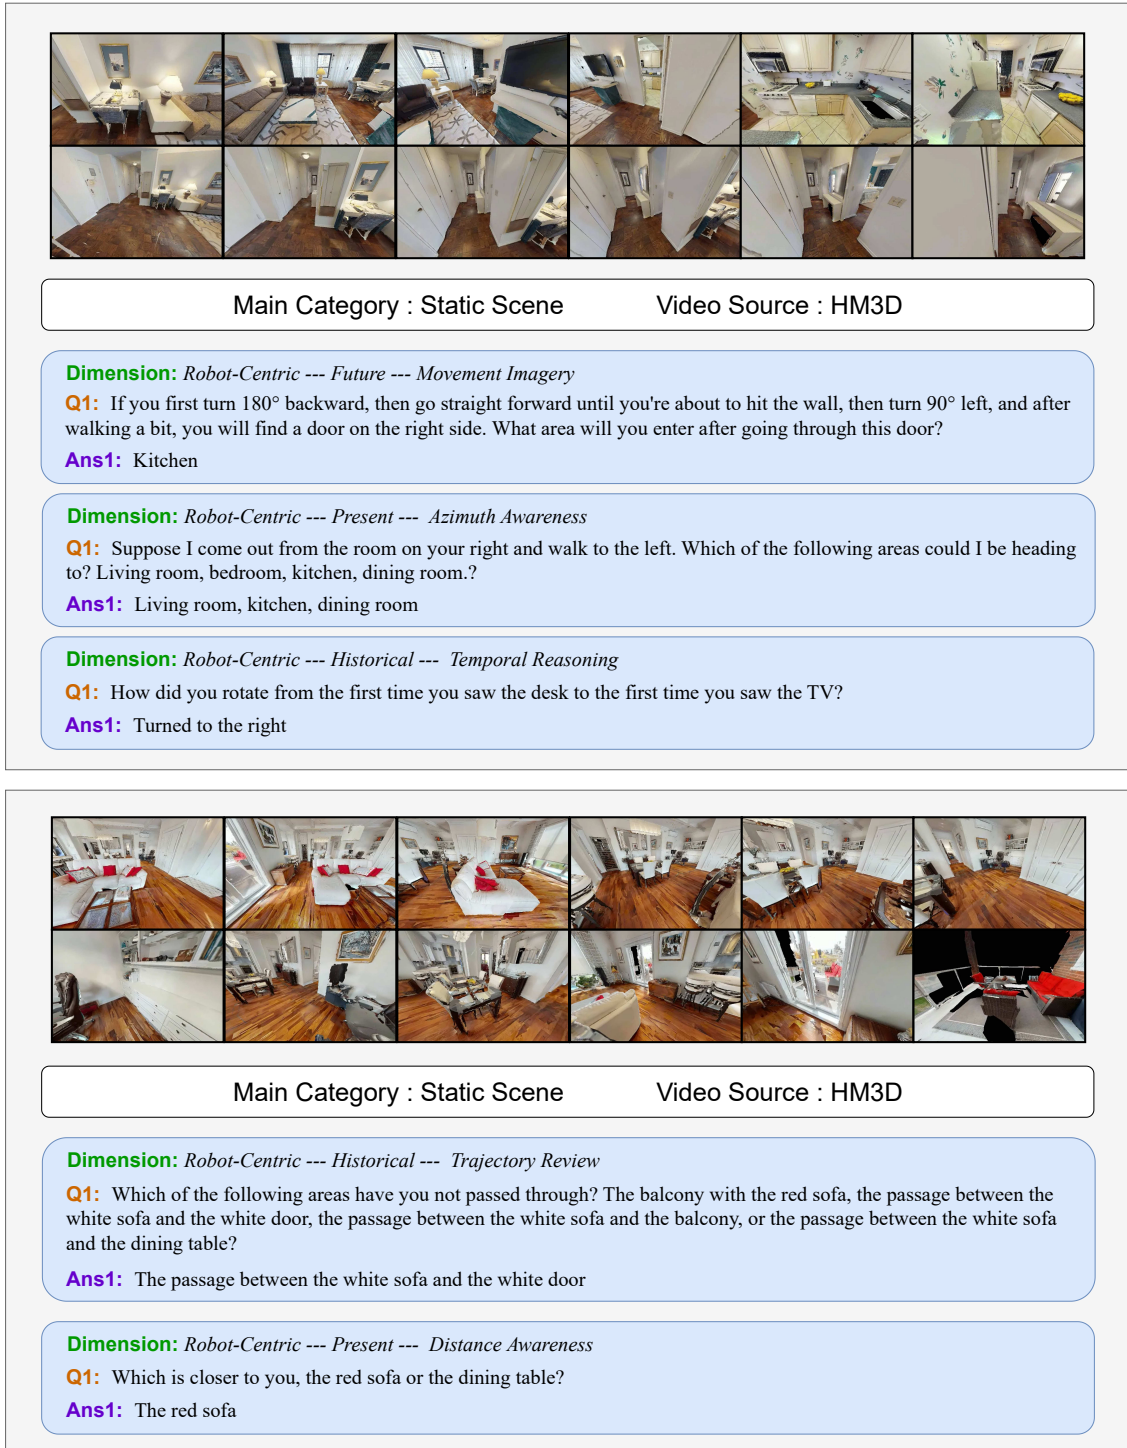


Figure 8. **Visualization of question answering examples in the static scene test set.** The main focus is on the robot-centric category. Part 1 out of 2.



Figure 9. **Visualization of question answering examples in the static scene test set.** The main focus is on the robot-centric category. Part 2 out of 2.

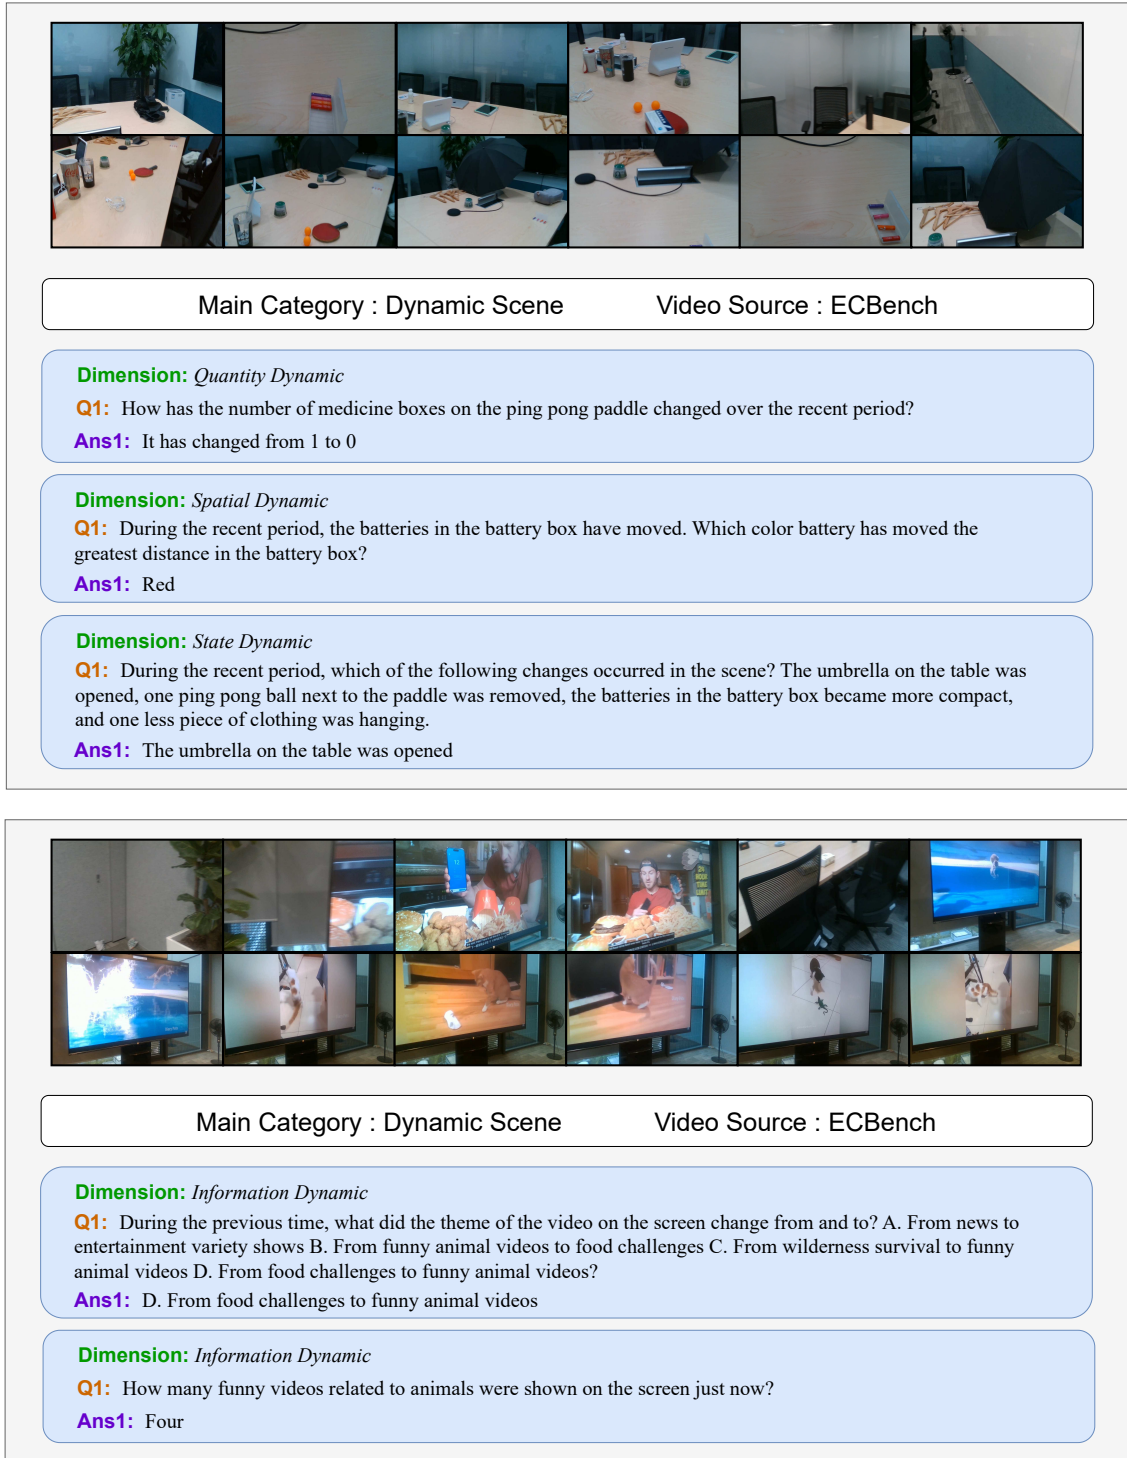


Figure 10. Visualization of question answering examples in the dynamic scene test set.

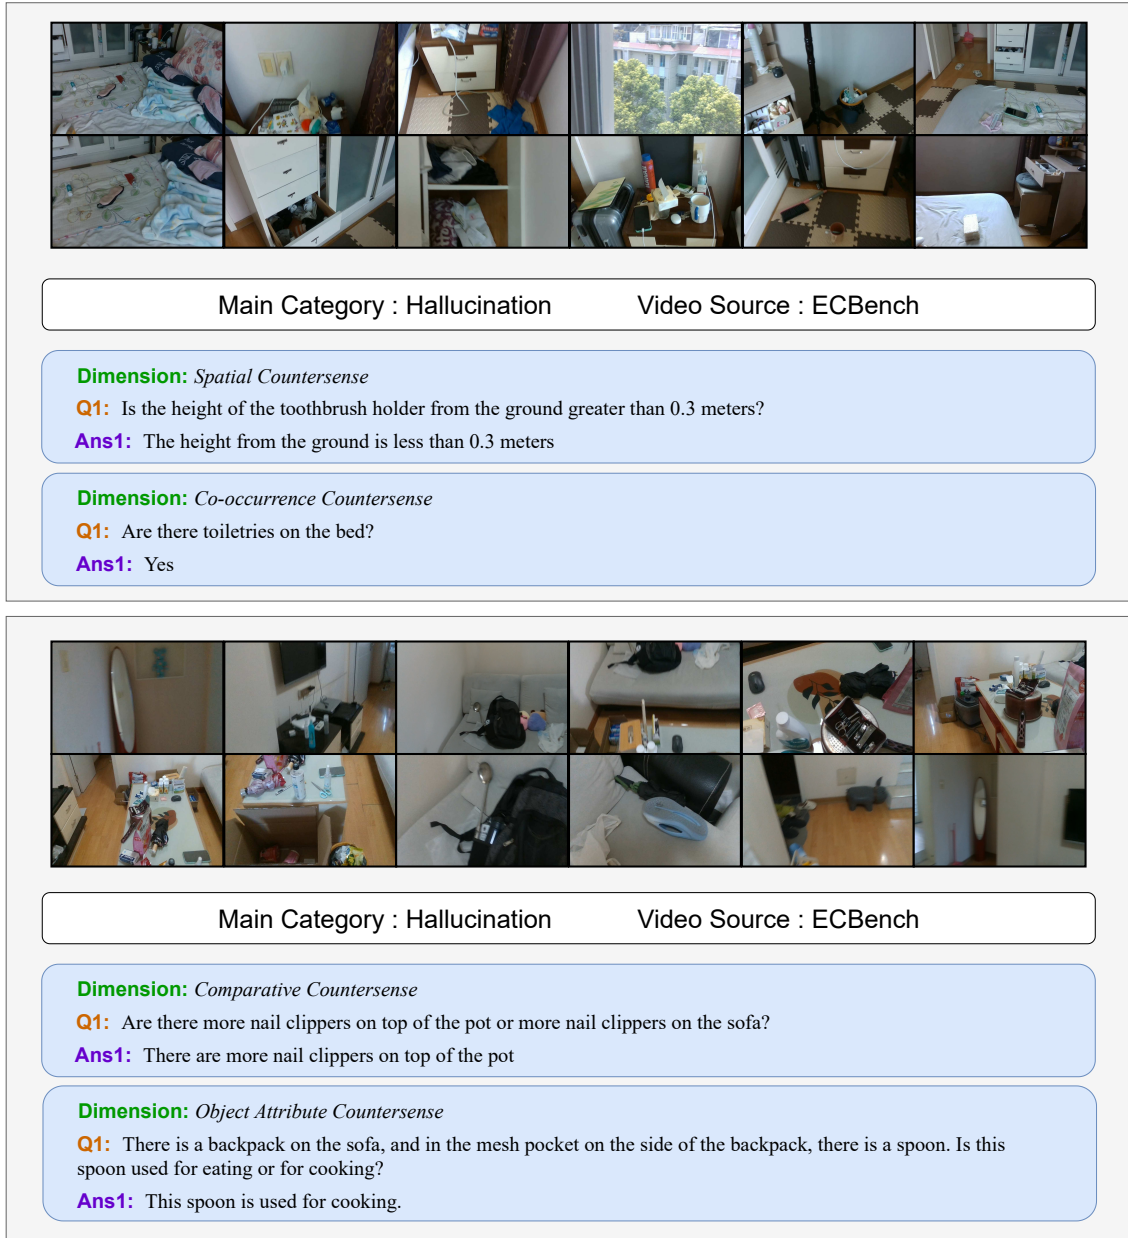


Figure 11. **Visualization of question answering examples in the hallucination test set.** The main focus is on the commonsense over-confidence category.

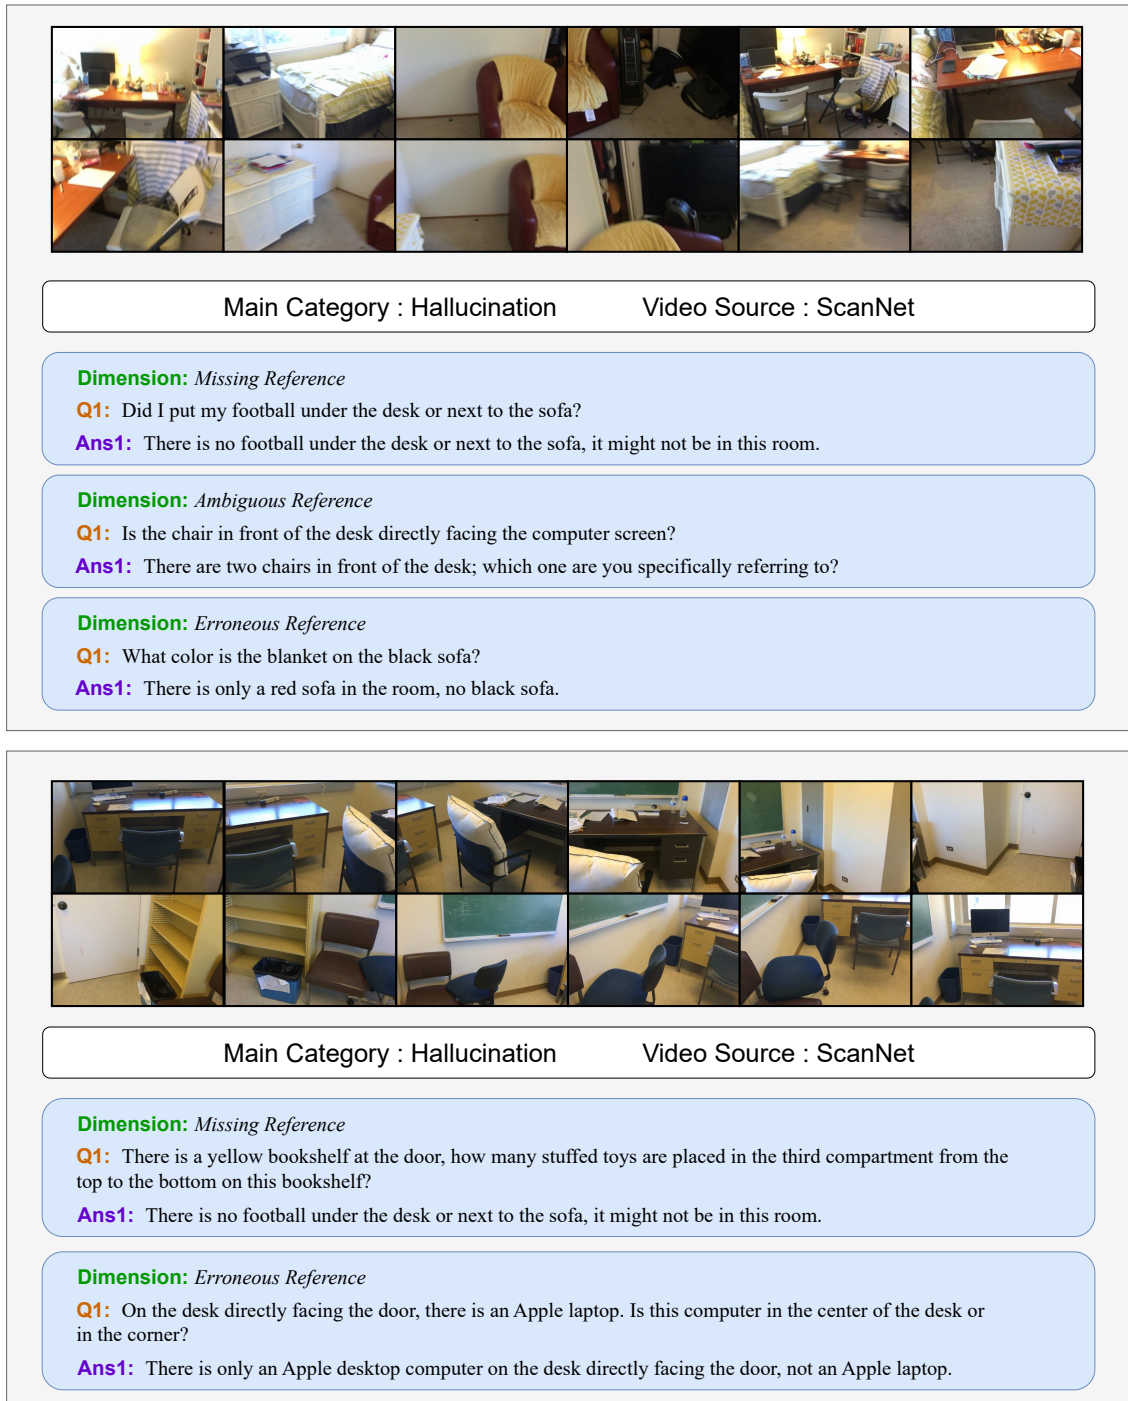


Figure 12. **Visualization of question answering examples in the hallucination test set.** The main focus is on the user-input over-confidence category.

7. Construction of Evaluation Framework

As mentioned in Sec. 3.1, we design ECEval, a fine-grain evaluation framework that considers different types of tasks. That is, we also annotate a partially correct answer for each question in addition to a completely correct answer. Specifically, we first design some open-ended question-answer pairs as examples and then instruct GPT-4o to classify each question-answer into two types: closed-ended and open-ended. For each open-ended answer, we additionally manually annotate a 0.5-point answer as the reference. When scoring each model’s responses, gpt-4o provides a continuous score between 0 and 1 based on the 1-point answer (100% correct) and the 0.5-point answer (50% correct). The prompt for LLM-assisted classification is as follows:

```
System Prompt: You are a helpful assistant.
You help me label questions that belong to
free-response questions. Please note that
if it is a multiple-choice question, a
numerical question, counting, OCR, a true/
false question, or a status judgment
question or similar questions with definite
answers, it does not belong to the free-
response question type. If you determine
that this question belongs to the free-
response question type. Please output ###
Free Response: True, otherwise ###Free
Response: False.
USER:
###Question: If you pick up the footstool
between the sofa and the TV and flip it 180
angle upside down, which objects in the
scene will be affected and how?
###Label Answer: Three handles and one
remote control will fall onto the carpet.
Assistant:
###Free Response: True
USER:
###Question: {example query}
###Label Answer: {example answer}.
Assistant:
###Free Response: {example Judgement}
....
USER:
###Question: {user query}
###Label Answer: {labeled answer}.
Assistant:
```

8. Additional Dataset Analysis

In Sec. 3.2 of the main text, we conduct a comprehensive quantitative analysis of ECBench. This section delves into a more detailed examination of the question-answer text within ECBench. In Fig. 13a, we present a word cloud depicting all the questions in ECBench. The most frequently occurring terms are predominantly spatial pronouns, such

as “left” and “front.” These words are crucial for robot-centric referentiality and are among the most commonly used vocabulary in our interactions with robots. Additionally, ECBench contains numerous object nouns, including “sofa” and “chair,” which represent the most common items within indoor environments and are frequently mentioned by users.

In Fig. 13b, we enumerate the top ten most prevalent object nouns in ECBench, with “table” appearing most frequently at 606 occurrences. Overall, ECBench encompasses 1,067 distinct object nouns, exhibiting a richness in object referentiality that far surpasses that of other embodied question-answering datasets [1, 27].

9. More Details of Experiment

For all experiments, we set the temperature to 0.2. For the GPT family models, we request the endpoint of GPT-4o-0513 and GPT-4o-mini-0513 through the Azure¹. For Qwen2VL-72B/7B and InternVL2, we use vllm² for employment. To score each model’s prediction, we uniformly employ GPT-4o-0513 to evaluate the predicted answers based on labeled answer and 0.5-point answer. We employ the same prompt for benchmarking and evaluating stages across all experiments as follows:

Benchmarking Each Model

```
System Prompt: You are moving in an indoor
environment. The image sequence is the
scene you just saw. You are now staying at
the last frame of the video. Please answer
the question with one word or one sentence,
as concise and accurate as possible.
USER: The question: {user query}
```

Scoring Open-ended Question

```
System Prompt: You are a helpful assistant.
Please score the predicted answer according
to the given question and huamn labeled
the 5-score answer, and the 3-score answer.
0 score represents completely wrong, 5
scores represents completely correct, and 3
scores represents partially correct.
Please refer to them to score the predicted
answer: [0, 1, 2, 3, 4, 5]. You need to
consider the answer from two perspectives:
accuracy and completeness. Output ###Judge:
USER:
###Question: If you pick up the footstool
between the sofa and the TV and flip it 180
degree upside down, which objects in the
scene will be affected and how?
###5 Score Answer: Three handles and one
remote control will fall onto the carpet.
```

¹<https://azure.microsoft.com/en-us/>

²<https://docs.vllm.ai/en/latest/>

direction for our future research.

Broader Impacts As an evaluation benchmark focused on the domain of embodiment, ECBench is poised to attract significant attention from researchers interested in exploring the cognitive processes of perceiving the real world through RGB-D videos. Furthermore, ECBench aims to facilitate the current LVLMs in transcending the limitations posed by online images and videos, thereby shifting their focus more towards the visual input modalities that robots encounter in the real world. The development of world models based on RGB-D video is a shared aspiration among AI researchers. We also hope that ECBench will contribute substantially to the advancement of robotic visual cognition capabilities.