

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 V2P-BENCH: EVALUATING VIDEO-LANGUAGE UN- DERSTANDING WITH VISUAL PROMPTS FOR BETTER HUMAN-MODEL INTERACTION

Anonymous authors

Paper under double-blind review

ABSTRACT

Large Vision-Language Models (LVLMs) have made significant strides in the field of video understanding in recent times. Nevertheless, existing video benchmarks predominantly rely on text prompts for evaluation, which often require complex referential language and diminish both the accuracy and efficiency of human–model interaction in turn. To address this limitation, we propose **V2P-Bench**, a robust and comprehensive benchmark for evaluating the ability of LVLMs to understand Video Visual Prompts in human–model interaction scenarios. V2P-Bench consists of 980 videos and 1172 well-structured high-quality QA pairs, each paired with manually annotated visual prompt frames. The benchmark spans three main tasks and twelve categories, thereby enabling fine-grained, instance-level evaluation. Through an in-depth analysis of current LVLMs, we identify several key findings: **1)** Visual prompts are both more model-friendly and user-friendly in interactive scenarios than text prompts, leading to significantly improved model performance and enhanced user experience. **2)** Models are reasonably capable of zero-shot understanding of visual prompts, but struggle with spatiotemporal understanding. Even o1 achieves only 71.8%, far below the human expert score of 88.3%, while most open-source models perform below 60%. **3)** LVLMs exhibit pervasive Hack Phenomena in video question answering tasks, which become more pronounced as video length increases and frame sampling density decreases, thereby inflating performance scores artificially. We anticipate that V2P-Bench will not only shed light on these challenges but also serve as a foundational tool for advancing human–model interaction and improving the evaluation of video understanding.

1 INTRODUCTION

In recent years, Large Vision-Language Models (LVLMs) have made significant strides in video understanding, showcasing powerful capabilities in tasks such as video captioning and question answering. Notable models like Gemini-2.5-Pro (Team et al., 2024) and LLaVA-Video (Zhang et al., 2024b) have set new performance benchmarks. In response, numerous benchmarks have emerged to evaluate these models comprehensively across diverse tasks (Li et al., 2024c; Mangalam et al., 2023; Fu et al., 2024). These benchmarks provide robust support for assessing LVLMs, allowing nuanced evaluations of their strengths and weaknesses in real-world applications. This growing landscape not only facilitates rigorous testing but also encourages further innovation in LVLM development, enhancing their effectiveness in video understanding tasks.

However, most benchmarks utilize text prompts for human–model interaction, which inevitably introduces certain inherent limitations. As shown in Figure 1, text prompts usually fail to provide precise spatial and temporal references, resulting in difficulties when assessing the ability of LVLMs to understand specific areas or moments in videos, particularly in complex multi-object scenarios. For users, a significant amount of referential language is required to specify targets. For the model, it first needs to comprehend the user’s referential language, making it prone to confusion at this initial step.

In contrast, as a frontier approach to multimodal human–model interaction, visual prompts offer a simpler and more precise way, facilitating model understanding and aligning more closely with human intuitive cognition. Some previous efforts (Cai et al., 2024; Yang et al., 2023; Lin et al., 2024)

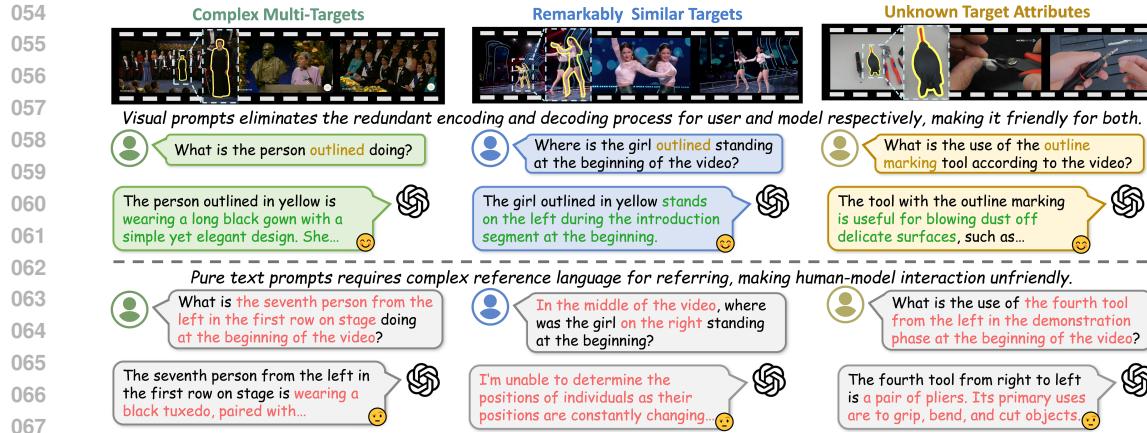


Figure 1: **Comparison of text prompts and visual prompts for users and models.** Pure text prompts face complex encoding and decoding issues, where users need to refer to targets using complicated language (*encoding*), while models must locate the user’s intended targets based on the text (*decoding*), leading to inconsistencies in understanding, especially in complex scenes with multiple, similar or unknown targets. In contrast, visual prompts offer a direct approach by creating prompts on video frames to directly refer to targets, avoiding the complexity and of pure text prompts, making it friendly for both users and models.

conduct initial explorations in image visual prompt areas, demonstrating the superiority of visual prompts over texts. However, existing studies lack research on video modality, limiting the further development of multimodal human-model interaction.

To bridge this gap, we propose **V2P-Bench**, a comprehensive benchmark specifically designed to evaluate the video understanding capabilities of LVLMs in human-model interaction scenarios. As illustrated in Figure 2, V2P-Bench encompasses three main tasks, twelve categories, twenty video types, eight types of visual prompts and wide distribution of durations ranging from three seconds to two hours. Each query includes one visual prompt annotation, focusing on fine-grained spatial and temporal understanding, aiming to comprehensively assess the video understanding abilities of LVLMs. All videos are meticulously curated by human annotators to ensure high-quality QA pairs and accurate visual prompts.

We conduct a thorough evaluation of multiple models on V2P-Bench. We first conduct a comparative study between visual prompts and text prompts, along with a user experience evaluation, demonstrating that visual prompts significantly enhance model performance and improve user interaction compared to text-based prompts. Then comes a comprehensive evaluation of 15 LVLMs, including 3 closed-source models and 12 open-source models. The results indicate that even the sota models exhibit suboptimal performance on this benchmark (e.g., o1 (OpenAI, 2024) achieving 71.8%), which is substantially lower than the human expert score of 88.3%, highlighting the current limitations of LVLMs in understanding video visual prompts. Further analysis reveals that LVLMs commonly exhibit hacking behaviors in video question-answering tasks, which become increasingly pronounced as video length increases and frame sampling rates decrease, leading to artificially inflated performance scores. In a nutshell, our contributions are as follows:

- V2P-Bench has been meticulously designed, comprising twelve categories covering a wide range of video types and diverse visual prompts. Collection and annotation process undergoes rigorous human validation, aiming to provide the community with a high-quality benchmark for multi-model human-model interaction.
- We demonstrate the superiority of visual prompts over text prompts. Experimental results reveal that current models exhibit substantial limitations in comprehending video visual prompts and display evidence of hacking behaviors in video question-answering tasks.
- V2P-Bench pioneeringly applies visual prompts in video understanding evaluation for human-model interaction, addressing critical limitations in existing text-based evaluation frameworks. We seek to advance the field of video visual prompt understanding evaluation and establish a foundation for more intuitive human-model interaction.

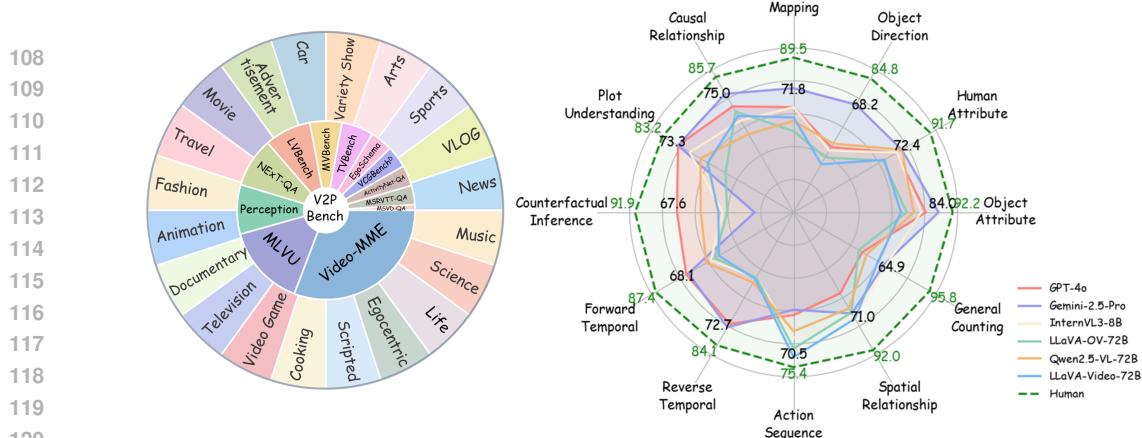


Figure 2: **(Left)** **Datasets and categories.** Our dataset is derived from twelve datasets and contains twenty restructured categories. **(Right)** **Performance radar chart.** We report the performance of different models on V2P-Bench by dimension. SOTA for each dimension is given.

2 RELATED WORK

2.1 LVLMs FOR VIDEO UNDERSTANDING

The rapid development of image-based LVLMs (Liu et al., 2024b; 2023; 2024a; Li et al., 2024a; Chen et al., 2024a;d; Bai et al., 2023) has significantly enhanced the potential of video understanding and question answering tasks, injecting new vitality into the field of artificial intelligence. VideoChat (Li et al., 2023b) and Video-ChatGPT (Maaz et al., 2023) are preliminary attempts in the realm of video understanding. Notable recent works include CogVLM2-Video (Hong et al., 2024), InternVL2 (Chen et al., 2024d) and LLaVA-Video (Zhang et al., 2024b), which treat videos as sequences of images and leverage the powerful image comprehension capabilities to process video modality. Furthermore, the high computational and memory demands required for handling high frame rates and long videos have spurred advancements in video compression technologies. For instance, InternVideo2 (Wang et al., 2024c) and Video-LLaMA (Zhang et al., 2023) utilize QFormer (Li et al., 2023a) for efficient video feature extraction. Despite promising results, current LVLMs primarily rely on text prompts and still face challenges in fine-grained spatial and temporal understanding when given visual prompts as input.

2.2 VIDEO UNDERSTANDING BENCHMARKS

Traditional video understanding benchmarks, such as MSRVTT-QA (Xu et al., 2017), ActivityNet-QA (Yu et al., 2019), and NExT-QA (Xiao et al., 2021), focus on basic action recognition and video question answering. Recently, more benchmarks have been proposed. MMBench (Liu et al., 2024c), SEED-Bench (Li et al., 2024b), and MVBench (Li et al., 2024c) mainly concentrate on short video clips for evaluation. EgoSchema (Mangalam et al., 2023) and MovieQA (Tapaswi et al., 2016) provide insights into narrative and thematic understanding. LongVideoBench (Wu et al., 2024), Video-MME (Fu et al., 2024), and LV-Bench (Wang et al., 2024b) offer longer videos and a broader variety of tasks. Additionally, recent works like INST-IT (Peng et al., 2024) and VideoRefer (Yuan et al., 2024) have introduced instance-level video question answering benchmarks. However, constrained by insufficiently robust and comprehensive, they still fail to adequately simulate real-world interactions. To address this limitation, we introduce V2P-Bench, allowing for a comprehensive evaluation of LVLMs that simulates multimodal human-model interaction in realistic settings.

2.3 VISUAL PROMPT AS A USER-FRIENDLY SOLUTION

Compared to text prompts, visual prompts offer a simple and effective means of facilitating interaction between users and models. Visual prompts have been widely utilized in image understanding. ViP-LLaVA (Cai et al., 2024) enhances the ability of LVLMs to comprehend local image regions by overlaying arbitrary visual prompts on images. Draw-and-Understand (Lin et al., 2024) employs a

162 Table 1: Comparison of different datasets. **Answer Type** indicates whether the QA pair is open-
 163 ended(OE) or multiple-choice(MC). **Multi Level** represents whether the videos cover multiple
 164 duration levels. **Open Domain** indicates whether the video source is diversified. **Visual Prompt**
 165 represents whether the video contains visual prompts. Refer to Appendix A for more features.

Benchmarks	Videos	Samples	Tasks	Avg duration	Annotation	Answer Type	Multi Level	Open Domain	Visual Prompt
MSVD-QA(Xu et al., 2017)	504	13157	1	9.8s	Auto	OE	✗	✗	✗
MSRVT-QA(Xu et al., 2017)	2990	72821	1	15.2s	Auto	OE	✗	✗	✗
ActivityNet-QA(Yu et al., 2019)	800	8000	3	111.4s	Manual	OE	✗	✗	✗
NExT-QA(Xiao et al., 2021)	1000	8564	3	39.5s	Manual	MC	✗	✗	✗
Perception Test(Patraucean et al., 2024)	11600	44000	4	23.0s	Auto&Manual	MC	✗	✗	✗
MLVU(Zhou et al., 2024)	1334	2593	9	~12min	Auto&Manual	OE&MC	✓	✓	✗
VCGBench-Diverse(Maaz et al., 2024)	877	4354	6	217.0s	Auto&Manual	OE	✗	✓	✗
MVBBench(Li et al., 2024c)	3641	4000	20	16.0s	Auto	MC	✗	✓	✗
HourVideo(Chandrasegaran et al., 2024)	500	12976	18	45.7min	Auto&Manual	MC	✗	✗	✗
LVBBench(Wang et al., 2024b)	103	1549	6	68.4min	Manual	MC	✗	✓	✗
EgoSchemal(Mangalam et al., 2023)	5063	5063	1	180.0s	Auto	MC	✗	✗	✗
Video-MME(Fu et al., 2024)	900	2700	12	17.0min	Manual	MC	✓	✓	✗
INST-IT Bench(Peng et al., 2024)	206	1000	1	14.2s	Auto&Manual	OE&MC	✗	✗	✓
VideoRefer Bench ^Q (Yuan et al., 2024)	198	1000	5	13.8s	Manual	MC	✗	✗	✓
V2P-Bench(ours)	980	1172	12	19.0min	Manual	MC	✓	✓	✓

178 two-stage training approach to improve performance in pixel-level tasks. Set-of-Mark (Yang et al.,
 179 2023) introduces a novel visual prompting method to enhance the performance of LVLMs in visual
 180 localization tasks. However, research on visual prompts in the context of video remains limited.
 181 INST-IT (Peng et al., 2024) introduces instruction tuning with visual prompts to enhance instance-
 182 level understanding in LVLMs. VideoRefer Suite (Yuan et al., 2024) creates a large instance-level
 183 video instruction dataset to assist LVLMs in understanding spatiotemporal information in videos.

3 V2P BENCH

187 Table 1 compares the key difference of V2P-
 188 Bench with previous benchmarks. The first two
 189 blocks list traditional pure text video understand-
 190 ing benchmarks, which primarily understand
 191 videos at a holistic level and lack instance-level
 192 comprehension. Instance-level understanding is
 193 crucial as it focuses on the specific elements of
 194 greatest interest to us, requiring a more nuanced understanding and consistency.



195 As shown in the third block, although INST-IT Bench (Peng et al., 2024) and VideoRefer Bench^Q
 196 (Yuan et al., 2024) use visual prompts for question-answering, their: 1) visual prompts are annotated
 197 on all frames, rendering them unsuitable for human–model interaction scenarios; 2) all data are
 198 sourced exclusively from VIS datasets (Yang et al., 2019; Pont-Tuset et al., 2017; Ding et al., 2023),
 199 thereby exhibiting limitations in both robustness and comprehensiveness, meaning a) Shorter video
 200 durations(14.2s and 13.8s); b) Single continuous shots; c) Limited video sources; d) Objects of
 201 interest not be suitable for question-answering.

3.1 TASK DEFINITION

205 To facilitate fine-grained evaluation of LVLMs from various perspectives, we categorize the questions
 206 according to dimensions. Our dimension design strives to ensure both comprehensiveness and
 207 orthogonality, and ultimately includes three main tasks and twelve dimensions. Definitions for tasks
 208 and dimensions are as follows:

- 209 • **Basic Perception** focuses on understanding the intrinsic attributes of objects and humans in the
 210 visual prompt. This task includes: 1) Object Attribute (OA); 2) Human Attribute (HA).
- 211 • **Temporal Understanding** emphasizes comprehension and processing of dynamic information and
 212 chronological sequences in videos. This task includes: 1) Object Direction (OD); 2) Feature Mapping
 213 (FM); 3) Forward Temporal (FT); 4) Reverse Temporal (RT); 5) Action Sequence (AS); 6) Spatial
 214 Relationship (SR); 7) General Counting (GC).
- 215 • **High-level Reasoning** extends beyond perception and temporal understanding, requiring logical
 216 inference and judgment to derive new conclusions or answers. This task includes: 1) Causal
 217 Relationship (CR); 2) Plot Understanding (PU); 3) Counterfactual Inference (CI).



Figure 4: Examples for each dimension of V2P-Bench.

Detailed elaborations and examples of each dimension are provided in Appendix A.2.

3.2 DATASET CONSTRUCTION

3.2.1 VIDEO COLLECTION

To create our dataset, we start from existing video benchmarks, as they already have a wide distribution of durations and diverse video types. We categorize the video durations into short, medium, and long videos. Additionally, we reclassify all the videos, resulting in twenty video categories, as shown in Figure 2(left). Our final dataset covers multiple video domains while maintaining a relative balance in video lengths.

3.2.2 QA AND VISUAL PROMPT ANNOTATION

After completing the collection process, we conduct the annotation of QA pairs and visual prompts to evaluate the capabilities of LVLMs in video understanding with visual prompts. The annotation work is carried out by researchers proficient in English. To ensure data quality, we provide thorough training for the annotators and conduct pilot annotations to assess their annotation capabilities.

While annotating the QA pairs, annotators are also required to perform visual prompt annotations. To better approximate real-world distributions, we adopt a fully manual approach for annotating visual prompts, with each QA pair constrained to one visual prompt frame. We predefine various types of visual prompts as follows: rectangle, mask contour, ellipse, triangle, scribble, point, arrow and SoM, as shown in Figure 3.

270 3.2.3 POST PROCESSING
271272 To ensure the quality of the dataset, we conduct a rigorous review of the annotated data after
273 completion, including both VLMs and manual review processes.274 **Blind LLMs Filtering.** Inspired by MMStar (Chen et al., 2024b), we perform plain text filtering
275 on the QA pairs to ensure that questions could only be answered correctly by viewing the videos.
276 Specifically, we provide only the pure text QA pairs to the most powerful closed-source models
277 GPT-4o and Gemini-2.5-Pro. We set the sampling temperature to 0.2 and conduct two rounds of
278 inference, then exclude samples for which both rounds provided correct answers.279 **Manual and Rule-based Review.** After that, we perform a rule-based check and manual review of
280 the remaining data. We exclude samples where the length disparity between different options was
281 too significant. Additionally, we shuffle the order of multiple-choice options to ensure a uniform
282 distribution of answer choices, thereby eliminating potential biases of different models toward specific
283 options. The filtering ratios are reported in the Appendix A.4. The final balanced proportions of
284 the four options are 28.0%, 23.9%, 25.0% and 23.1%. Through this rigorous dataset construction
285 process, we strive to provide a high-quality, diverse, and balanced dataset that will benefit researchers
286 in the field of human-model interaction.

287

288 4 EXPERIMENTS
289290 4.1 EXPERIMENT SETUP
291292 **Evaluation Models.** To thoroughly evaluate the effectiveness of V2P-Bench, we conduct assessments
293 on multiple models, including 3 closed-source models: o1 (OpenAI, 2024), GPT-4o (Hurst et al.,
294 2024) and Gemini-2.5-Pro (Team et al., 2024); 12 open-source models: LLaVA-OneVision(7B 72B)
295 (Li et al., 2024a), InternVL3-8B (Zhu et al., 2025), mPLUG-Owl3-7B (Ye et al., 2024), LLaVA-
296 Video(7B 72B) (Zhang et al., 2024b), MiniCPM-V 2.6-8B (Yao et al., 2024), Qwen2.5-VL(7B
297 72B) (Wang et al., 2024a), MiMo-VL-7B (Team et al., 2025), LLaVA-NeXT-7B (Liu et al., 2024a)
298 and LLaVA-NeXT-INST-IT-7B (Peng et al., 2024). This essentially covers the mainstream LVLMs
299 currently available.300 **Implementation Details.** For open-source models, we select the sampling frame rate based on
301 the context length of each model. Specifically, we average 128 frames from the video for LLaVA-
302 OneVision(7B 72B), InternVL3-8B, mPLUG-Owl3-7B, LLaVA-Video(7B 72B), MiniCPM-V 2.6-8B,
303 Qwen2.5-VL(7B 72B), MiMo-VL-7B, 4 frames for LLaVA-NeXT-7B and LLaVA-NeXT-INST-IT-
304 7B. For o1 and GPT-4o, we average 64 frames. For Gemini-2.5-Pro, the raw video was uploaded
305 directly without audio track. The visual prompt frame is placed after the video for all models. For
306 other hyperparameters, we follow the settings in VLMEvalKit (Duan et al., 2024).

307

308 4.2 QUANTITATIVE RESULTS
309310 Table 2 and 5 present the comprehensive evaluation results of V2P-Bench across different dimensions
311 and video durations. These results encompass the performance of human experts, the blind answering
312 task, and evaluations of 15 different models. We can conclude that o1 achieves the highest overall
313 score; however, its performance is not consistently superior across all dimensions, particularly in
314 Object Direction and Action Sequence. As shown at the top of Table 2, human experts achieve an
315 accuracy of 88.3%, representing the upper bound of human performance. For the blind answering
316 task, we observe that all three models perform below 10% on this task, with GPT-4o at 1.4%, Gemini-
317 2.5-Pro at 9.6%, and Qwen2.5-VL-7B at 3.0%. This demonstrates that our post-processing pipeline
318 effectively filters out commonsense-based questions, thereby ensuring high quality and robustness.

319

320 Below we summarize our key findings as follows:

321 **Which prompt type works better for humans and models?** Most benchmarks rely solely on text
322 prompts, requiring models to answer questions based on textual questions and visual context. To
323 investigate which type of prompt(textual or visual) is more conducive in human model interaction
scenarios, we conduct a prompt comparison experiment and a real-user study. Results are reported in

324
 325 **Table 2: Evaluation Results on V2P-Bench across Dimensions.** We report results for 12 open-source
 326 models, 3 closed-source models, 3 blind answering results and human performance on V2P-Bench
 327 across dimensions. The best results are **bold** and the second-best are underlined.

Method	OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Avg
Human Performance	92.2	91.7	84.8	89.5	85.7	83.2	91.9	87.4	84.1	75.4	92.0	95.8	88.3
<i>Pure Text as Input</i>													
GPT-4o	2.3	1.2	0.0	0.0	4.6	1.8	0.0	1.7	0.0	0.0	0.0	4.6	1.4
Gemini-2.5-Pro	15.6	12.0	6.5	3.9	9.2	10.9	8.1	5.2	13.6	1.8	8.0	7.4	9.6
Qwen2.5-VL-72B	3.9	4.0	0.0	0.0	7.3	3.6	0.0	0.0	6.8	5.4	0.0	0.0	3.0
<i>Closed-source Models</i>													
o1	85.2	78.4	23.1	78.1	<u>71.6</u>	78.7	<u>66.7</u>	69.1	73.1	50.0	64.1	51.2	71.8
GPT-4o	76.6	68.9	41.3	60.8	<u>67.0</u>	73.3	67.6	<u>68.1</u>	70.5	50.0	54.0	48.4	65.4
Gemini-2.5-Pro	84.0	<u>72.4</u>	68.2	<u>71.8</u>	75.0	<u>73.3</u>	22.6	66.7	<u>72.7</u>	47.4	67.5	<u>63.6</u>	69.8
<i>Open-source Models</i>													
LLaVA-NeXT-7B	56.6	55.6	34.8	52.5	43.0	48.6	31.6	42.6	42.2	28.1	42.0	30.5	46.0
LLaVA-NeXT-INST-IT-7B	57.4	58.4	26.1	42.4	43.0	49.2	31.6	49.2	42.2	26.3	42.0	27.4	46.3
LLaVA-OV-7B	57.1	52.1	28.3	47.1	63.8	59.1	41.0	42.1	35.6	63.2	62.8	43.2	52.8
LLaVA-OV-72B	65.5	59.9	34.7	47.0	63.8	43.2	38.5	50.0	41.1	<u>66.3</u>	66.9	45.9	56.7
InternVL3-8B	73.9	69.1	39.1	60.8	58.1	65.9	41.0	52.6	41.1	61.1	<u>69.7</u>	64.9	61.7
mPLUG-Owl3-7B	61.3	54.4	28.3	49.0	60.0	50.0	51.3	60.5	34.4	55.8	58.6	37.8	52.6
LLaVA-Video-7B	60.5	58.1	37.0	49.0	62.9	54.5	41.0	52.6	48.9	57.9	56.6	40.5	54.8
LLaVA-Video-72B	62.2	60.8	30.4	54.9	61.0	54.5	43.6	47.4	42.2	70.5	71.0	59.5	58.6
MiniCPM-V 2.6-8B	68.9	59.4	26.1	56.9	58.1	50.0	33.3	50.0	34.4	57.9	67.6	43.2	55.3
Qwen2.5-VL-7B	60.5	56.7	17.4	45.1	47.6	40.9	48.7	52.6	32.2	47.4	50.3	35.1	48.1
Qwen2.5-VL-72B	69.7	<u>72.4</u>	<u>43.5</u>	52.9	49.5	59.1	53.8	55.3	44.4	57.9	64.1	51.4	59.8
MiMo-VL-7B	67.2	57.6	37.0	45.1	47.6	52.3	59.0	55.3	41.1	50.5	61.4	43.2	53.8

347
 348 **Table 3.** We can draw the following observations: **(1) Visual prompts are more model-friendly than text prompts.**
 349 As shown in Table 3a, simply converting visual prompts
 350 into text prompts leads to a substantial drop in accuracy
 351 across all participating models, with the most pronounced
 352 decline 15.1% observed in Gemini-2.5-Pro. This is because
 353 text prompts require the model to decode the target from
 354 the text, which increases the difficulty of comprehension.
 355 Moreover, textual references can sometimes be ambiguous,
 356 as illustrated in Figure 1. In contrast, visual prompts can
 357 precisely localize the target within video frames, bypassing
 358 both the user’s need to encode intentions in text and the
 359 model’s need to decode them. **(2) Visual prompts are more**
 360 **user-friendly than text prompts.** We recruit 20 volunteers
 361 to participate in the experiment. Specifically, the interaction
 362 process consists of: watching a video, formulating a
 363 meta-question, rewriting the question into both text and a
 364 visual-prompt version, completing the QA session, and then
 365 indicating their preference(text or visual prompts). We also
 366 record the order in which users write the text version and the
 367 visual prompt version of each question. Each participant is
 368 required to produce 10 text and 10 visual prompt questions.
 369 All questions are open-ended but designed to have unique
 370 correct answers to minimize potential subjectivity or eval-
 371 uation bias. The correctness of model responses is assessed
 372 by our annotators. All videos are randomly sampled from
 373 the V2P-Bench dataset we construct. Throughout the study,
 374 Gemini-2.5-Pro serves as the conversational agent. To systematically evaluate performance, we
 375 record five key metrics: answer accuracy (max 100), completion time (seconds), user satisfaction
 376 (max 10), user preference, and question order. As shown in Table 3b, users complete tasks more
 377 quickly in the visual prompt interaction setting, with the time reduced from 25.2s to 18.1s, saving 7.7s
 378 on average. In addition, the improvement in model performance means more satisfactory responses
 379 for users. Together, these factors contributed to an average user satisfaction score of 8.2, 2.4 points
 380 higher than under the text-only prompt condition. This indicates that visual prompts provide clear

(a) Model performance.		
Model	Text	Visual
GPT-4o	53.0	65.4
Gemini-2.5-Pro	54.7	69.8
LLaVA-Video-7B	42.4	54.8
Qwen2.5-VL-7B	43.1	52.4
Mimo-VL-7B	46.7	55.6

(b) User experience study.		
Metric	Text	Visual
Acc	57.0	69.5
Cost Time	25.2	18.1
User Satisfaction	5.3	7.5

(c) User preference for prompt type.		
Preference	Text	Visual
Nums	57	129
Percentage	28.5%	64.5%
		7.0%

(d) Question order on user responses.		
Order	Text First	Visual First
Nums	64	136
Percentage	34.0%	68.0%

Table 3: Comparison of text and visual prompts for humans and models.

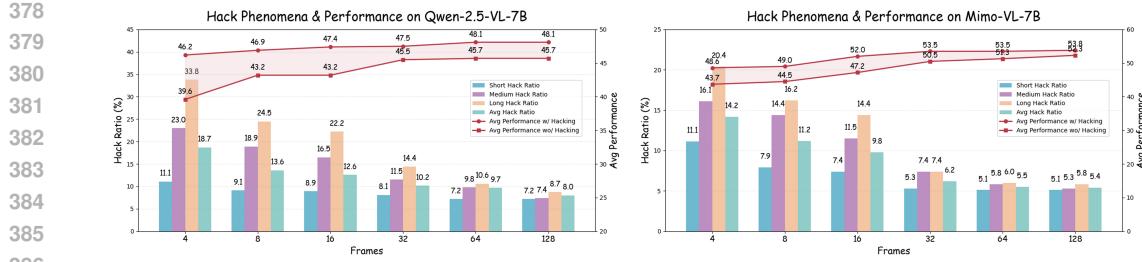


Figure 5: Hack Phenomena and Model Performance on V2P-Bench.

advantages in human–model interaction, improving both overall satisfaction and efficiency. Besides, 64.5% of participants explicitly preferred using visual prompts, while only 28.5% preferred text-only input, and 7.0% had no preference. What’s more, 68% of the questions were initiated using visual prompts. This indicates that for real users, visually selecting the target is more intuitive and less effortful than crafting a natural-language description, especially for objects, locations, or people that are difficult to describe precisely with text. Refer to Appendix E.2 for the implementation details.

Can the models effectively comprehend visual prompts, and how do they perform? In Table 2, we report the performance of human experts and all models across dimensions, leading to the following observations: **(1) Models are reasonably capable of zero-shot understanding of visual prompts.** Except for LLaVA-NeXT-INST-IT-7B, the open-source models have not been trained on visual prompt datasets, yet all achieve performance above 45 points, surpassing the random-guess baseline (25%) by over 20%. Notably, in the Basic Perception task, every model achieves over 50%. This phenomenon arises because LVLMs are typically trained on massive image–text and video–text pairs, enabling them to learn broad visual–semantic associations and thus exhibit strong zero-shot transfer capabilities. **(2) Models struggle with spatiotemporal understanding.** Object Direction requires models to identify the direction of an object’s movement, while Spatial Relationship requires models to understand the dynamic spatial positions. The average scores on the two dimensions are only 34.4% and 45.7%, respectively, indicating that spatiotemporal tasks remains a weakness of current models and requires further improvement. **(3) Closed-source models and large-scale parameter models possess stronger capabilities.** As shown in Table 2, three closed-source models outperform all open-source counterparts. This disparity highlights the persistent challenges in advancing open-source models. Moreover, for the same open-source model with varying scales, for example Qwen2.5-VL-7B and 72B, the larger-parameter variant demonstrates stronger capabilities and achieves higher scores, which is consistent with the established scaling laws of LVLMs.

Does the expert model exhibit the anticipated superior performance? LLaVA-NeXT-INST-IT is fine-tuned on INST-IT dataset based on LLaVA-NeXT, yet the results show only a marginal improvement of 0.3%. There are two main reasons for this limited performance gain: **(1) Limited diversity of visual prompt types.** LLaVA-NeXT-INST-IT is trained exclusively on SoM data. As observed in Table 13, while the model achieves a 15.0% improvement on SoM prompts, its performance drops on other types of visual prompts, indicating a forgetting phenomenon. This suggests that models’ training should cover a broad range of visual prompt types to better meet the demands of real-world scenarios. **(2) Differences in data format.** Unlike V2P-Bench, INST-IT dataset annotates visual prompts on every sampled frame, as shown in Figure 9 (*left*). This redundant annotation does not account for the constraints of user interactions in real human–model interaction scenarios, which contributes to the suboptimal performance. A detailed analysis of this phenomenon provides important insights for building future datasets and training strategies. Refer to Appendix D for more information.

Are vision-language models truly understanding videos, or merely exploiting hacks?

Due to the sparsity of frame sampling and the upper limits of model perception, the models may guess answers rather than rely on visual context. We randomly shuffle the videos and questions and performed inference with Qwen2.5-VL-7B and MiMo-VL-7B. As shown in Table 4, only 6.4% and 3.9% of the cases trigger a refusal, respectively, while all others follow the instructions to select an option. This indicates that current models are trained to be “test-takers” often neglecting basic factual information. To investigate to what extent models exhibit hack behavior

Table 4: Results with shuffled video and question pairs.

Model	Trigger Ratio
Qwen2.5-VL-7B	6.4%
MiMo-VL-7B	3.9%

on V2P-Bench, we conduct experiments in which the models are instructed to reject when they cannot reach a conclusion. Results are reported in Figure 5. From our analysis, we draw the following conclusions: **(1) Hack Phenomena exist in video benchmark evaluations.** Both Qwen2.5-VL-7B and MiMo-VL-7B exhibit positive Hack Ratios across all settings, accompanied by varying degrees of performance degradation. **(2) Longer videos exacerbate Hack Phenomena.** For instance, under a 4-frame sampling setting, the Hack Ratios of Qwen2.5-VL-7B for short, medium, and long videos are 11.1%, 23.0%, and 33.8%, respectively. This pattern is consistent across other experiments as well. **(3) Fewer sampled frames increase Hack Phenomena.** For example, when reducing the number of sampled frames for Qwen2.5-VL-7B from 128 to 4, the average Hack Ratio steadily rises from 8.0% to 18.7%. The presence of Hack Phenomena can be attributed to three factors: **(1) Insufficient information**, as existing sparse sampling strategies fail to provide the model with enough information to get the answer; **(2) Limited model perception**, where excessive visual context may overwhelm the key information; **(3) Training strategy**. Models are trained as instruction-following agents, leading them to prioritize following instructions over grounding their responses in factual information. Efforts should be made to actively enhance model capabilities and explore novel visual representation mechanisms that capture a broader range of effective context. When visual context is insufficient, models should proactively request clarification. We also report the results of other benchmarks in Appendix I.3, which are consistent with our observation on V2P-Bench.

How LVLMs are robust to varied video duration?

In Table 5, we compare the performance of different models on short, medium, and long videos, from which we can conclude that: **Performance degrades as video length increases.** All models exhibit a significant drop in performance. For example, o1’s performance on long videos decreases by 23.5% compared to the medium. This decline is primarily due to the increased sparsity of frame sampling as video length grows, which reduces the amount of effective visual content. Such sparsity prevents models from retaining all relevant visual–semantic information, thereby hindering accurate predictions. However, all models perform worse on short videos compared to medium-length videos unexpectedly. This is likely because over half of the short videos are drawn from Perception Test, MVBench and TVBench, which have high information density and contain challenging questions related to spatiotemporal questions.

How does the structure of visual prompts affect model performance? To investigate the effect of the visual query structure itself on the performance of visual prompts, we randomly sampled 217 data instances. For each question–answer pair, we annotated the corresponding visual prompt frames with multiple types of visual prompts, while strictly keeping all other settings identical, including the number of frames, prompts. The results are shown in Table 6.

For the same visual prompt type, the doodle-style shapes are slightly weaker than the standard shapes. When switching from standard shapes to doodle shapes, the performance of Qwen2.5-VL-7B and MiMo-VL-7B decreases by 0.7% and 0.8%, respectively. This is reasonable, since most training data are synthetic, whereas hand-drawn doodles often have unstable boundaries, making it harder for the model to extract consistent visual prompts. Regardless of whether standard or doodle shapes are used, the overall performance trend remains Rectangle > SoM > Arrow. Among these, rectangles appear most frequently in training data and can fully enclose the target region, providing a stable and explicit spatial localization signal. In contrast, arrows are typically smaller and carry lower information density, making it more difficult for the model to capture key spatial relationships and therefore resulting in the weakest performance.

Table 5: **Evaluation results on V2P-Bench across durations.** The best results are **bold** and the second-best are underlined.

Method	Short	Medium	Long	Avg
Human Performance	91.6	87.3	84.0	88.3
<i>Pure Text as Input</i>				
GPT-4o	1.1	1.9	1.6	1.4
Gemini-2.5-Pro	10.0	10.6	8.1	9.6
Qwen2.5-VL-7B	2.7	5.1	2.5	3.0
<i>Closed-source Models</i>				
o1	75.2	<u>83.9</u>	60.4	71.8
GPT-4o	67.3	70.8	<u>59.3</u>	65.4
Gemini-2.5-Pro	<u>73.8</u>	86.3	54.5	<u>69.8</u>
<i>Open-source Models</i>				
LLaVA-NeXT-7B	47.0	47.1	43.8	46.0
LLaVA-NeXT-INST-IT-7B	48.6	51.1	39.5	46.3
LLaVA-OV-7B	51.3	63.0	47.3	52.8
LLaVA-OV-72B	54.5	70.4	49.0	56.7
InternVL3-8B	61.7	68.5	55.6	61.7
mPLUG-Owl3-7B	52.2	62.5	44.9	52.6
LLaVA-Video-7B	54.1	65.7	46.5	54.8
LLaVA-Video-72B	57.5	66.2	54.3	58.6
MiniCPM-V 2.6-8B	53.3	66.2	50.2	55.3
Qwen2.5-VL-7B	48.0	53.2	43.6	48.1
Qwen2.5-VL-72B	62.4	63.9	50.2	59.8
MiMo-VL-7B	56.8	58.8	42.4	53.8

Table 6: Performance on Different Visual Prompts.

Method	Standard Shapes			Doodle Shapes		
	Rectangle	Arrow	SoM	Rectangle	Arrow	SoM
Qwen2.5-VL-7B	47.3	43.6	45.1	46.7	42.9	44.4
MiMo-VL-7B	54.2	51.2	52.7	53.6	50.3	51.9

486 4.3 ERROR ANALYSIS
487

488 In this section, we analyze the error patterns of LVLMs
489 in video visual prompt understanding. We examine and
490 categorize 470 model predictions from Qwen2.5-VL-7B,
491 identifying three main representative error types. These
492 patterns are illustrated in Figure 10 to Figure 12, and their
493 distribution is presented in Figure 6.

494 • **Perception Error.** Provided with sufficient visual
495 context, the model still produces errors due to deficiencies
496 in its perception capabilities. Perception errors account
497 for 57.7%, making them the most prevalent error type.

498 • **Reasoning Error.** The model demonstrates a range of
499 logical reasoning failures, including plot understanding,
500 counterfactual reasoning and causal errors. Some of these errors arises from deficiencies in perception,
501 which further undermine the model’s reasoning ability.

502 • **Hack Error.** Constrained by the sparse frame sampling strategy and the upper limit of the model’s
503 perceptual capability, the model fails to recognize sufficient visual semantics and consequently resorts
504 to arbitrary guessing, which ultimately leads to erroneous predictions.

505 Analyzing and mitigating these errors is crucial for enhancing the performance of LVLMs in video
506 visual prompt question understanding. This analysis provides an opportunity to target specific error
507 types, thereby improving the model’s overall capability.

509 5 CONCLUSION
510

511 In this study, we introduce V2P-Bench, a comprehensive benchmark for evaluating the video un-
512 derstanding capabilities of LVLMs through visual prompts in human–model interaction scenarios.
513 Through a rigorous construction and evaluation process, V2P-Bench enables systematic evaluation
514 and useful insights into the issues within the current models. Our experiments demonstrate three key
515 findings: **1)** visual prompts outperform text prompts, substantially improving accuracy and interaction
516 efficiency; **2)** a notable performance gap remains between LVLMs and human experts, especially in
517 spatiotemporal understanding; and **3)** Hack Phenomena are prevalent, exacerbated by longer videos
518 and sparser frame sampling. These results underscore the need for more robust evaluation protocols
519 and model designs. We envision V2P-Bench as both a diagnostic tool and a roadmap for advancing
520 LVLMs toward more reliable, human-aligned video understanding and interaction.

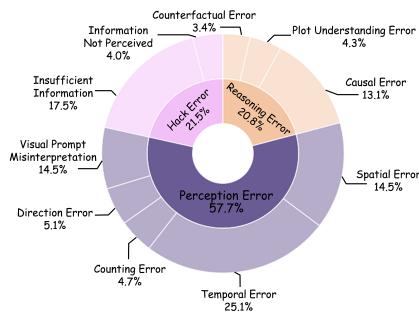


Figure 6: Distribution of Error Types.

521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 **6 ETHICS STATEMENT**
541542 This work complies with ethical research standards in data collection, model training, and evaluation.
543 All datasets used in this study are publicly available research datasets, collected and released under
544 their respective licenses. No private or personally identifiable information (PII) was used.
545546 **7 REPRODUCIBILITY STATEMENT**
547548 We are committed to ensuring the reproducibility of our results. All prompts used during the
549 evaluation are provided in the appendices. In addition, the datasets and evaluation scripts used in this
550 paper have been publicly released.
551552 **REFERENCES**
553554 Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang
555 Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities.
556 *arXiv preprint arXiv:2308.12966*, 2023.557 Mu Cai, Haotian Liu, Siva Karthik Mustikovela, Gregory P Meyer, Yuning Chai, Dennis Park, and
558 Yong Jae Lee. Vip-llava: Making large multimodal models understand arbitrary visual prompts.
559 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
560 12914–12923, 2024.561 Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris,
562 Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, Jie Lei, Tengyu
563 Ma, Baishan Guo, Arpit Kalla, Markus Marks, Joseph Greer, Meng Wang, Peize Sun, Roman
564 Rädle, Triantafyllos Afouras, Effrosyni Mavroudi, Katherine Xu, Tsung-Han Wu, Yu Zhou,
565 Liliane Momeni, Rishi Hazra, Shuangrui Ding, Sagar Vaze, Francois Porcher, Feng Li, Siyuan
566 Li, Aishwarya Kamath, Ho Kei Cheng, Piotr Dollár, Nikhila Ravi, Kate Saenko, Pengchuan
567 Zhang, and Christoph Feichtenhofer. Sam 3: Segment anything with concepts, 2025. URL
568 <https://arxiv.org/abs/2511.16719>.569 Keshigeyan Chandrasegaran, Agrim Gupta, Lea M Hadzic, Taran Kota, Jimming He, Cristóbal
570 Eyzaguirre, Zane Durante, Manling Li, Jiajun Wu, and Li Fei-Fei. Hourvideo: 1-hour video-
571 language understanding. *arXiv preprint arXiv:2411.04998*, 2024.573 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin.
574 Sharegpt4v: Improving large multi-modal models with better captions. In *European Conference
575 on Computer Vision*, pp. 370–387. Springer, 2024a.576 Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi
577 Wang, Yu Qiao, Dahua Lin, et al. Are we on the right way for evaluating large vision-language
578 models? *arXiv preprint arXiv:2403.20330*, 2024b.580 Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong
581 Duan, Bin Lin, Zhenyu Tang, et al. Sharegpt4video: Improving video understanding and generation
582 with better captions. *arXiv preprint arXiv:2406.04325*, 2024c.583 Zhe Chen, Jiannan Wu, Wenhui Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong
584 Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning
585 for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer
586 Vision and Pattern Recognition*, pp. 24185–24198, 2024d.587 Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, and Chen Change Loy. Mevis: A large-scale
588 benchmark for video segmentation with motion expressions. In *Proceedings of the IEEE/CVF
589 International Conference on Computer Vision*, pp. 2694–2703, 2023.591 Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang
592 Zang, Pan Zhang, Jiaqi Wang, et al. Vlmevalkit: An open-source toolkit for evaluating large
593 multi-modality models. In *Proceedings of the 32nd ACM international conference on multimedia*,
pp. 11198–11201, 2024.

594 Chaoyou Fu, Yuhang Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu
 595 Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation
 596 benchmark of multi-modal llms in video analysis. *arXiv preprint arXiv:2405.21075*, 2024.

597

598 Wenyi Hong, Weihan Wang, Ming Ding, Wenmeng Yu, Qingsong Lv, Yan Wang, Yean Cheng,
 599 Shiyu Huang, Junhui Ji, Zhao Xue, et al. Cogvilm2: Visual language models for image and video
 600 understanding. *arXiv preprint arXiv:2408.16500*, 2024.

601 Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Os-
 602 trow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint*
 603 *arXiv:2410.21276*, 2024.

604

605 Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan
 606 Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv preprint*
 607 *arXiv:2408.03326*, 2024a.

608 Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan.
 609 Seed-bench: Benchmarking multimodal large language models. In *Proceedings of the IEEE/CVF*
 610 *Conference on Computer Vision and Pattern Recognition*, pp. 13299–13308, 2024b.

611

612 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image
 613 pre-training with frozen image encoders and large language models. In *International conference*
 614 *on machine learning*, pp. 19730–19742. PMLR, 2023a.

615

616 KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhui Wang, Ping Luo, Yali Wang, Limin Wang, and
 617 Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023b.

618

619 Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen,
 620 Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In
 621 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 622 22195–22206, 2024c.

623

624 Weifeng Lin, Xinyu Wei, Ruichuan An, Peng Gao, Bocheng Zou, Yulin Luo, Siyuan Huang, Shang-
 625 hang Zhang, and Hongsheng Li. Draw-and-understand: Leveraging visual prompts to enable
 626 mllms to comprehend what you want. *arXiv preprint arXiv:2403.20271*, 2024.

627

628 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
 629 tuning. *arXiv preprint arXiv:2310.03744*, 2023.

630

631 Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee.
 632 Llava-next: Improved reasoning, ocr, and world knowledge, January 2024a. URL <https://llava-vl.github.io/blog/2024-01-30-llava-next/>.

633

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

636

637 Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi
 638 Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player?
 639 In *European conference on computer vision*, pp. 216–233. Springer, 2024c.

640

641 Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt:
 642 Towards detailed video understanding via large vision and language models. *arXiv preprint*
 643 *arXiv:2306.05424*, 2023.

644

645 Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. Videogpt+: Integrating image
 646 and video encoders for enhanced video understanding. *arXiv preprint arXiv:2406.09418*, 2024.

647

Kartikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. Egoschema: A diagnostic
 648 benchmark for very long-form video language understanding. *Advances in Neural Information*
 649 *Processing Systems*, 36:46212–46244, 2023.

650

OpenAI. Introducing openai o1, 2024., 2024. URL <https://openai.com/o1/>.

648 Viorica Patraucean, Lucas Smaira, Ankush Gupta, Adria Recasens, Larisa Markeeva, Dylan Banarse,
 649 Skanda Koppula, Mateusz Malinowski, Yi Yang, Carl Doersch, et al. Perception test: A diagnostic
 650 benchmark for multimodal video models. *Advances in Neural Information Processing Systems*, 36,
 651 2024.

652 Wujian Peng, Lingchen Meng, Yitong Chen, Yiweng Xie, Yang Liu, Tao Gui, Hang Xu, Xipeng Qiu,
 653 Zuxuan Wu, and Yu-Gang Jiang. Inst-it: Boosting multimodal instance understanding via explicit
 654 visual prompt instruction tuning. *arXiv preprint arXiv:2412.03565*, 2024.

655 Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and
 656 Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint
 657 arXiv:1704.00675*, 2017.

658 Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja
 659 Fidler. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of
 660 the IEEE conference on computer vision and pattern recognition*, pp. 4631–4640, 2016.

661 Core Team, Zihao Yue, Zhenru Lin, Yifan Song, Weikun Wang, Shuhuai Ren, Shuhao Gu, Shicheng
 662 Li, Peidian Li, Liang Zhao, Lei Li, Kainan Bao, Hao Tian, Hailin Zhang, Gang Wang, Dawei Zhu,
 663 Cici, Chenhong He, Bowen Ye, Bowen Shen, Zihan Zhang, Zihan Jiang, Zhixian Zheng, Zhichao
 664 Song, Zhenbo Luo, Yue Yu, Yudong Wang, Yuanyuan Tian, Yu Tu, Yihan Yan, Yi Huang, Xu Wang,
 665 Xinzhe Xu, Xingchen Song, Xing Zhang, Xing Yong, Xin Zhang, Xiangwei Deng, Wenyu Yang,
 666 Wenhan Ma, Weiwei Lv, Weiji Zhuang, Wei Liu, Sirui Deng, Shuo Liu, Shima Chen, Shihua Yu,
 667 Shaohui Liu, Shande Wang, Rui Ma, Qiantong Wang, Peng Wang, Nuo Chen, Menghang Zhu,
 668 Kangyang Zhou, Kang Zhou, Kai Fang, Jun Shi, Jinhao Dong, Jiebao Xiao, Jiaming Xu, Huaqiu
 669 Liu, Hongshen Xu, Heng Qu, Haochen Zhao, Hanglong Lv, Guoan Wang, Duo Zhang, Dong
 670 Zhang, Di Zhang, Chong Ma, Chang Liu, Can Cai, and Bingquan Xia. Mimo-vl technical report,
 671 2025. URL <https://arxiv.org/abs/2506.03569>.

672 Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett
 673 Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal
 674 understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.

675 Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu,
 676 Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng
 677 Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model’s
 678 perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024a.

679 Weihang Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Xiaotao Gu, Shiyu Huang,
 680 Bin Xu, Yuxiao Dong, et al. Lvbench: An extreme long video understanding benchmark. *arXiv
 681 preprint arXiv:2406.08035*, 2024b.

682 Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yinan He, Guo Chen, Baoqi Pei, Rongkun Zheng,
 683 Zun Wang, Yansong Shi, et al. Internvideo2: Scaling foundation models for multimodal video
 684 understanding. In *European Conference on Computer Vision*, pp. 396–416. Springer, 2024c.

685 Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. Longvideobench: A benchmark for long-context
 686 interleaved video-language understanding. *arXiv preprint arXiv:2407.15754*, 2024.

687 Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-
 688 answering to explaining temporal actions. In *Proceedings of the IEEE/CVF conference on computer
 689 vision and pattern recognition*, pp. 9777–9786, 2021.

690 Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video
 691 question answering via gradually refined attention over appearance and motion. In *Proceedings of
 692 the 25th ACM international conference on Multimedia*, pp. 1645–1653, 2017.

693 Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-
 694 mark prompting unleashes extraordinary visual grounding in gpt-4v, 2023. URL <http://arxiv.org/abs/2310.11441>, 2023.

695 Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *Proceedings of the
 696 IEEE/CVF international conference on computer vision*, pp. 5188–5197, 2019.

702 Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li,
 703 Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint*
 704 *arXiv:2408.01800*, 2024.

705

706 Jiabo Ye, Haiyang Xu, Haowei Liu, Anwen Hu, Ming Yan, Qi Qian, Ji Zhang, Fei Huang, and Jingren
 707 Zhou. mplug-owl3: Towards long image-sequence understanding in multi-modal large language
 708 models. *arXiv preprint arXiv:2408.04840*, 2024.

709

710 Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueling Zhuang, and Dacheng Tao. Activitynet-qa:
 711 A dataset for understanding complex web videos via question answering. In *Proceedings of the*
712 AAAI Conference on Artificial Intelligence, volume 33, pp. 9127–9134, 2019.

713

714 Yuqian Yuan, Hang Zhang, Wentong Li, Zesen Cheng, Boqiang Zhang, Long Li, Xin Li, Deli Zhao,
 715 Wenqiao Zhang, Yueling Zhuang, et al. Videorefer suite: Advancing spatial-temporal object
 716 understanding with video llm. *arXiv preprint arXiv:2501.00599*, 2024.

717

718 Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language
 719 model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023.

720

721 Youcai Zhang, Xinyu Huang, Jinyu Ma, Zhaoyang Li, Zhaochuan Luo, Yanchun Xie, Yuzhuo Qin,
 722 Tong Luo, Yaqian Li, Shilong Liu, et al. Recognize anything: A strong image tagging model.
 723 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 724 1724–1732, 2024a.

725

726 Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video
 727 instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*, 2024b.

728

729 Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang,
 730 Tiejun Huang, and Zheng Liu. Mlvu: A comprehensive benchmark for multi-task long video
 731 understanding. *arXiv preprint arXiv:2406.04264*, 2024.

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

APPENDIX OVERVIEW

- Section A: More Dataset Details.
- Section A.4: Dataset Scaling Up.
- Section A.4: More Experiments.
- Section D: Qualitative Examples.
- Section E: Implementation Details.
- Section F: Error Analysis.
- Section G: Application Scenarios for Visual Prompts.
- Section H: Discussion.
- Section I: Additional Results.
- Section J: Prompt Template.

A MORE DATASET DETAILS

A.1 DATASET STATISTICS

In Table 1, we have already presented the main characteristics of V2P-Bench. Overall, the proposed V2P-Bench defines three main tasks and twelve dimensions, encompassing 980 unique videos and 1,172 QA pairs sourced from twelve existing video datasets, covering twenty video categories. The average duration of the videos is 19.0 minutes, which represents a wide range of video lengths, differing from most benchmarks. The format of the QA pairs is multiple-choice with 4 options. Below we introduce a more detailed analysis of our benchmark:

- **Wide distribution of durations.** Figure 7(left) shows the detailed duration distributions on V2P-Bench. We follow Video-MME (Fu et al., 2024) in categorizing video lengths into short (< 3 minutes), medium (3-30 minutes), and long videos (30-120 minutes), with respective proportions of 46.8%, 22.0%, and 31.2%.
- **Diverse video types and comprehensive tasks.** Figure 2(left) shows various datasets and categories on V2P-Bench. Figure 7(right) shows the detailed distribution of each dimension.
- **Diverse Targets and Visual Prompts.** Figure 3 shows various targets and visual prompts on V2P-Bench, benefiting from diverse video sources.
- **Comprehensive Shot Types.** V2P-Bench includes both continuous and transition videos, the latter of which significantly increases the difficulty of reference, implying that the model must perform temporal and spatial grounding in different scenes.

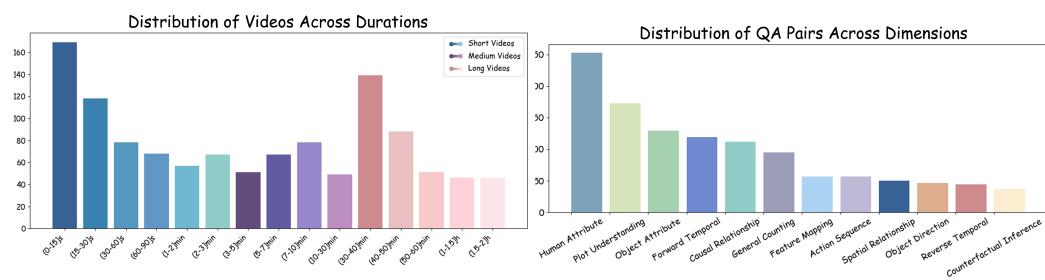


Figure 7: Distribution of video durations and dimensions.

810 A.2 ELABORATION ON DIMENSIONS
811812 Table 7 presents detailed information on the three main tasks and twelve dimensions of V2P-Bench.
813814 Table 7: Our proposed three main tasks and twelve dimensions with explanation.
815

Basic Perception	
Object Attribute	<i>This dimension evaluates the model's ability to perceive the visual and motion attributes of objects indicated by visual prompts, such as color, shape, position, and movement.</i>
Temporal Understanding	
Forward Temporal	<i>This dimension assesses the model's ability to accurately locate the visual prompt and track subsequent events that follow the natural chronological order of the video.</i>
Object Direction	<i>This dimension examines the model's ability to perceive and interpret the motion trajectory of objects pointed by visual prompts, with a particular focus on movement direction.</i>
Feature Mapping	<i>This dimension examines the model's capability to extract distinctive features of objects indicated by visual prompts and consistently track them across the entire video.</i>
Reverse Temporal	<i>This dimension evaluates the model's capacity to comprehend the temporal structure of the video by identifying events that precede the visual prompt, demonstrating an understanding of temporal precedence.</i>
Action Sequence	<i>This dimension evaluates the model's ability to grasp the overall temporal flow of the video, particularly in understanding and reasoning about the temporal dynamics of multiple action sequences of individuals or objects, as indicated by visual prompts.</i>
Spatial Relationship	<i>This dimension assesses the model's ability to discern and comprehend the spatial relationships between instances highlighted by visual prompts within the video scene.</i>
General Counting	<i>This dimension evaluates the model's ability to perceive and accurately count repeated actions or objects within the video, as indicated by visual prompts, testing its capacity for detailed content understanding and comprehensive scene analysis.</i>
High-level Reasoning	
Causal Relationship	<i>This dimension assesses the model's ability to perceive the causal relationships between actions and events, identifying the underlying intentions of actions and the causes of subsequent events. The visual prompt points to the action executor.</i>
Plot Understanding	<i>This dimension examines the model's ability to analyze narrative progression and logically infer subsequent developments based on the given plot. The visual prompt executes the protagonist of the plot.</i>
Counterfactual Inference	<i>This dimension evaluates the model's ability to reason under hypothetical scenarios that deviate from the actual video content, with visual prompts guiding the deviation, assessing its capacity to infer potential outcomes based on counterfactual assumptions.</i>

864 A.3 ANNOTATION GUIDELINES
865866 A.3.1 VISUAL PROMPT ANNOTATION
867868 Visual Prompt Annotation Guidelines
869870 **Interaction Constraint.** To remain consistent with real-world interactive applications,
871 each QA pair is restricted to a single visual prompt frame. This design emphasizes the
872 simplicity of user prompt creation and minimizes annotation burden. Therefore, annotators
873 must select the most representative frame that anchors the key moment or location of the event.
874875 **Visual Prompt Type.** Annotators should select the appropriate type of visual prompt
876 according to the target and strictly follow the eight predefined categories.
877878 **Uniqueness.** Visual prompts must precisely point to a specific object or region in the video,
879 avoiding ambiguity or multiple referents.
880881 **Consistency.** Visual prompts must strictly align with the question text, ensuring a one-to-one
882 correspondence between the annotated target and the object referred to in the question,
883 without mismatches or misleading references.
884885 **Multi-target Differentiation.** When multiple prompts appear in the same frame, different
886 shapes or colors should be used to clearly distinguish between targets.
887888 A.3.2 QA PAIRS ANNOTATION
889890 QA Pairs Annotation Guidelines
891892 **Scenario Realism.** Questions must be based on the actual video content rather than hypothetical
893 or fictional scenarios to ensure relevance to the real world and avoid fabricated or
894 unrelated plots.
895896

- Correct : Asking *What is the object held by the marked person?* when annotating a
frame where a person raises a cup.
- Incorrect : *If he is holding a beer, what will he do next?* (overly speculative).

897898 **Answer Uniqueness.** Questions should be concise and straightforward, avoiding ambiguous
899 references or subjective judgments to ensure a unique answer.
900901

- Correct : *What color of clothes is the target inside the rectangle wearing?*
- Incorrect : *What does this person look like?* (ambiguous and hard to standardize).

902903 **Avoid Reliance on Common Sense.** Do not design questions that can be answered solely
904 using common sense; ensure that the model must refer to the visual content to answer.
905906

- Correct : *What color is the label of the target inside the ellipse?*
- Incorrect : *Do people usually use a knife to cut vegetables?* (answerable without
watching the video).

907908 **Fine-grained Observation and Diversity.** Questions should emphasize detailed observation
909 and cover diverse objects (people, animals, tools, etc.), rather than being restricted to a
910 specific target.
911912

- Correct : *What color jersey is the circled target wearing?*
- Incorrect : *What activity is happening in the video?* (too broad).

913914 **Dependence on Visual Prompts.** Questions must rely on visual prompts to be answered,
915 which may require multiple targets to appear in the video, rather than inferring from context
916 or common sense. Avoid questions unrelated to the prompts.
917

918
919
920
921
922
923

- Correct : *What object is the target pointed to by the arrow holding?*
- Incorrect : *Is the person in the video cooking?* (lacking visual prompt).

Minimize Feature Descriptions. Use only the visual prompts to refer to the target; avoid describing the target’s appearance in words.

924
925
926
927

- Correct : When annotating a frame where an arrow points to a person sitting, ask: *What did the target do before sitting?*
- Incorrect : *What did the person wearing beige clothes and black pants do before sitting?* (contains appearance description).

928
929
930
931

Concise Language and Distractor Design. Questions and options should be kept concise to avoid noise from long descriptions and prevent models from learning biases unrelated to visuals. Options in multiple-choice questions should follow a consistent style, avoiding hints from length, tone, or format.

932
933
934
935

Misleading yet incorrect distractors. Besides common types (e.g., quantity substitution, causal reversals), annotators must carefully control the distractor’s misleading nature to ensure answer uniqueness.

936
937

A.3.3 DESIGN OF DISTRACTORS

938
939

Standards for Designing Distractors

940
941
942
943
944

Quantity Substitution. Replacing the correct numerical information (such as quantities, years, percentages, number of flags, etc.) with incorrect values.

945
946
947
948

Tool / Object Replacement. Swapping the actual tool used for another similar but different item—for example, writing hoe instead of rake, or switching the functions of a peeler and a knife.

949
950
951
952
953
954

Role Reversal. Exchanging the duties or action order of participants—for instance, changing “the cameraman cleans the table while the assistant kneads the dough” to the opposite distribution of tasks.

955
956
957
958
959
960

Step / Sequence Inversion. Reordering the correct procedure or omitting key steps—for example, writing the actions as “pass then plant” instead of the correct “plant then pass.”

961
962
963
964
965
966
967

Identity / Name Error. Substituting the real person’s name, nationality, or position with another—for example, writing Nathan instead of Eric, or labeling a Canadian athlete as British.

968
969
970
971

Attribute Replacement. Changing the attribute of clothing, objects, or chart lines to another common color—for example, describing black clothing as green.

Orientation / Gesture Reversal. Reversing directions or gestures—for example, writing “the right hand holds the cup” or “both hands hold it” instead of “the left hand holds the cup, and the right hand returns to the table.”

Emotion Substitution. Replacing the true expression—such as “surprised” or “smiling”—with emotions like “angry,” “sad,” or “bored.”

Misplaced Cause-and-Effect. Substituting the real reason with another seemingly plausible explanation—for example, the girl covering her ears is truly because “there are too many mice,” but the distractor states “neighbors are renovating.”

972

973

974

975

976

977

978

979

980

981

982

983

984

985

986

Location Replacement. Changing the venue to a similar but different place—for instance, describing “a prison” as “a hospital” or “a hotel.”

Relationship Misjudgment. Describing friends as lovers, enemies, or strangers, or labeling a father-daughter pair as siblings, and so on.

Misattributed Result / Follow-up Action. Giving an incorrect description of subsequent actions or impacts—for example, writing “the cameraman bows in thanks” instead of “waves and steps down.”

Addition of Extraneous False Details. Introducing nonexistent actions or objects—such as “handing the host a torch”—to create confusion through redundant information.

A.4 QUALITY CONTROL

In Section 3.2.3, to ensure the quality of the dataset, we conducted a rigorous review of the annotated data after completion. The review process included Blind LLMs filtering and manual and rule-based review. Initially, we had 1,747 QA pairs. After applying Blind LLMs filtering, 1,653 pairs remained (94 filtered). Subsequently, we performed a comprehensive manual and rule-based review, resulting in a final count of 1,172 QA pairs (481 filtered). Table 8 shows some representative data examples from the post-processing step.

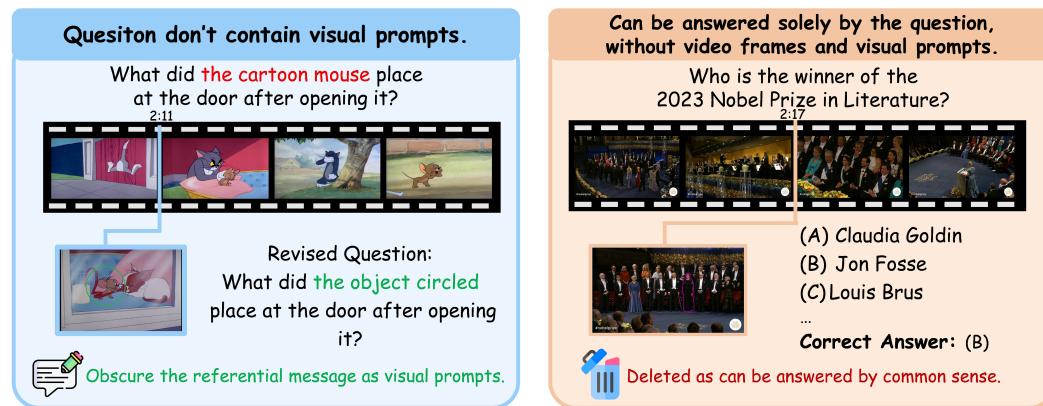


Figure 8: Representative data examples from the post-processing step.

A.5 CERTIFICATE LENGTH

Following EgoSchema Mangalam et al. (2023), we compute the temporal certificate length and the temporal certificate ratio (the proportion of the minimal effective content duration relative to the total duration) based on the video length and task type in our dataset. The results are shown in Table 8.

Table 8: Certificate Length and Certificate Length Percentage across Different Video Durations.

	Short			Medium			Long		
	BP	TU	HR	BP	TU	HR	BP	TU	HR
Certificate Length	1s	31s	27s	1s	204s	169s	1s	1147s	984s
Certificate Length Percentage	1.9%	61.4%	52.9%	0.2%	44.7%	37.2%	0.003%	35.3%	30.3%

The Basic Perception task is designed solely to assess a model’s fundamental ability to perceive visual prompts, and thus it can be completed using a single prompted frame. In contrast, Temporal Understanding and High-level Reasoning focus on dynamic information and chronological relation-

1026 ships in videos, and on logical inference and higher-level judgment, respectively. As a result, both
 1027 tasks require multiple video frames to be properly resolved.
 1028
 1029
 1030

1031 A.6 DATASET BIAS

1032
 1033 Our benchmark is constructed from 12 publicly available video datasets, all of which were originally
 1034 annotated with pure text-based QA pairs. The annotators of these datasets did not assume that “users
 1035 will provide visual prompts” when creating the annotations, meaning that the data are naturally more
 1036 aligned with text descriptions. The visual-prompt version of the benchmark is a rewrite we created
 1037 on top of these original annotations. Therefore, from the perspective of data provenance: If any bias
 1038 exists, it should favor pure text, not visual prompts.

1039 Furthermore, to ensure a fair comparison between text prompts and visual prompts, we uniformly
 1040 rewrote all text prompts rather than directly reusing the original textual annotations. Since the 12
 1041 constituent datasets vary widely in linguistic style, granularity, and descriptive conventions, directly
 1042 comparing visual prompts with the original text would introduce additional bias. To avoid this, we
 1043 rewrote all questions into a standardized text-prompt format and ensured the text prompts contained
 1044 the same information as the visual prompts.

1045 B DATASET SCALING UP

1046
 1047 The core value of a benchmark dataset lies in being high-quality, trustworthy, and diagnostic, rather
 1048 than merely covering as much data as possible. To this end, we rely on trained human annotators
 1049 instead of automated generation, ensuring that each question maintains a strict causal linkage between
 1050 the video evidence, visual prompts, and the correct answer. Thus, the benchmark does not prioritize
 1051 “scaling up data volume” as its primary goal.

1052 On the other hand, to enable the construction of large-scale visual-prompt datasets and to enhance
 1053 existing models’ ability to understand visual prompts, we also provide a fully automated data-
 1054 generation pipeline, including:

- 1055 1) Using RAM++ Zhang et al. (2024a) to automatically extract potential target objects from video.
- 1056 2) Applying SAM3 Carion et al. (2025) for cross-frame object tracking and mask propagation to
 1057 obtain temporally consistent visual annotations.
- 1058 3) Following automatic data-synthesis strategies from LLaVA-Video Zhang et al. (2024b) and
 1059 ShareGPT4Video Chen et al. (2024c) to bind model-generated QA pairs to target regions, thus
 1060 producing large quantities of high-quality video–visual-prompt supervision data.

1061 C MORE EXPERIMENTS

1062 C.1 CONSISTENCY ANALYSIS OF MCQ AND DIALOGUE TASKS

1063
 1064 We adopt a MCQ format as MCQ-based evaluation is easy to automate, highly controllable, and
 1065 effective at avoiding the ambiguity and vagueness that commonly arise in open-ended responses.
 1066 This is also why existing mainstream video-understanding benchmarks (Video-MME Fu et al. (2024),
 1067 MVBench Li et al. (2024c), EgoSchema Mangalam et al. (2023), LVbench Wang et al. (2024b)
 1068 .etc) widely adopt a QA-style design. In contrast, free-form dialogue introduces substantial subjectivity
 1069 and is difficult to score in a consistent and reproducible manner. Moreover, relying on an LLM-as-a-
 1070 judge may introduce evaluation bias, further weakening the fairness, reliability, and comparability
 1071 of the assessment. However, real human–model interaction scenarios are inherently dialogue tasks.
 1072 To examine the extent to which a MCQ task can represent a dialogue task, we further analyze the
 1073 consistency between the two. Specifically, we remove the multiple-choice constraints from the
 1074
 1075

1080 original questions, convert them into open-ended ones and re-evaluate the models using GPT-4o
 1081 Hurst et al. (2024) and human annotators as the judge.

1082 As shown in Table 9, under the open-ended setting, the Pearson consistency coefficient between
 1083 MCQ and OE (GPT-4o Judge) reaches 0.95, and the coefficient between MCQ and OE (Human
 1084 Judge) reaches 0.98. This demonstrates a strong consistency between QA-style evaluation and
 1085 free-form dialogue, indicating that QA-based assessments can reliably reflect a model’s capabilities
 1086 in open-ended conversational scenarios.

1088 Table 9: Consistency Evaluation Between MCQ and Open-Ended Tasks.

Model	MCQ	OE (GPT-4o Judge)	OE (Human Judge)
Gemini-2.5-Pro	69.8	48.3	51.6
Qwen2.5-VL-7B	48.1	30.2	33.7
MiMo-VL-7B	45.7	32.7	35.6

1098 C.2 HACK PHENOMENA IN OPEN-ENDED EVALUATION

1100 Existing benchmarks commonly adopt the MCQ format because it is easy to evaluate and highly
 1101 controllable, and avoids the ambiguity often seen in open-ended responses. Our experimental results
 1102 show that models indeed exhibit such behavior: they tend to game the MCQ structure to obtain higher
 1103 scores instead of performing true video comprehension. This finding suggests that, in addition to
 1104 pursuing higher benchmark scores, future models should also prioritize genuine understanding and
 1105 groundedness as key optimization objectives.

1106 To examine whether open-ended generation helps reduce hack behaviors, we conducted two experiments
 1107 in which we removed the multiple-choice constraints and answer options from the original
 1108 prompts, converting them into open-ended question-answering tasks:

1110 **Random Shuffling Experiment.** We randomly shuffle the videos and questions, meaning the model
 1111 can no longer obtain the answer from the video content. If a model is sufficiently honest, it should
 1112 refuse to answer all questions. We perform inference using Qwen2.5-VL-7B and MiMo-VL-7B. The
 1113 Trigger Ratio denotes the proportion of cases in which the model refuses to answer. For the MCQ
 1114 setting, we perform standard reasoning and count all responses that do not select any option; for the
 1115 OE setting, we manually evaluate whether the model refused to respond. The results are shown in
 1116 Table 10.

1117 Table 10: Hack Phenomena under MCQ and Open-Ended Evaluation.

Model	MCQ Trigger Ratio (%)				OE Trigger Ratio (%)			
	Short	Medium	Long	Avg	Short	Medium	Long	Avg
Qwen2.5-VL-7B	6.2	6.9	6.3	6.4	95.8	96.8	96.1	96.7
MiMo-VL-7B	3.6	4.2	4.0	3.9	92.2	94.6	94.8	93.6

1124 We observe that under the MCQ setting, Qwen2.5-VL-7B and MiMo-VL-7B exhibit refusal rates
 1125 of only 6.4% and 3.9%, respectively. This indicates that even when the video and question are
 1126 completely mismatched, the models still select an option in the vast majority of cases, and the trigger
 1127 rates remain almost identical across different video lengths. In contrast, under the OE setting, the
 1128 refusal rates rise sharply to 96.7% and 93.6%, again showing little variation across video lengths.
 1129 These results clearly demonstrate that open-ended generation effectively mitigates hack behaviors to
 1130 a large extent.

1132 **Model Performance on V2P-Bench.** Besides, we follow the same setup as in the random shuffling
 1133 experiment, with the only difference being that we use the original benchmark data rather than the
 1134 shuffled version. The experimental results are shown in Table 11.

1134 Table 11: Hack Rates under MCQ and Open-Ended Evaluation.
1135

Model	MCQ Hack Ratio (%)				OE Hack Ratio (%)			
	Short	Medium	Long	Avg	Short	Medium	Long	Avg
Qwen2.5-VL-7B	8.1	11.5	14.4	10.2	2.1	2.6	3.4	2.7
MiMo-VL-7B	5.3	7.2	7.4	6.2	1.7	2.3	2.8	2.2

1142 We observe that after converting the MCQ task into an open-ended question–answering task, the
1143 Hack Ratios of Qwen2.5-VL-7B and MiMo-VL-7B decrease by 7.5% and 4.0%, respectively. This
1144 further demonstrates that open-ended generation can substantially mitigate hack behaviors.
1145
1146
1147

1148 C.3 IMPACT OF SAMPLING FRAME RATES

1149 We evaluate Qwen2.5-VL-7B across different task types and video durations to investigate how
1150 varying sampling frame rates affect model performance. Here, 1 frame indicates that only the
1151 visual-prompt frame is provided to the model. The results are shown in Table 12.
1152

1153 Table 12: Performance under Different Sampling Frame Rates across Video Durations.
1154

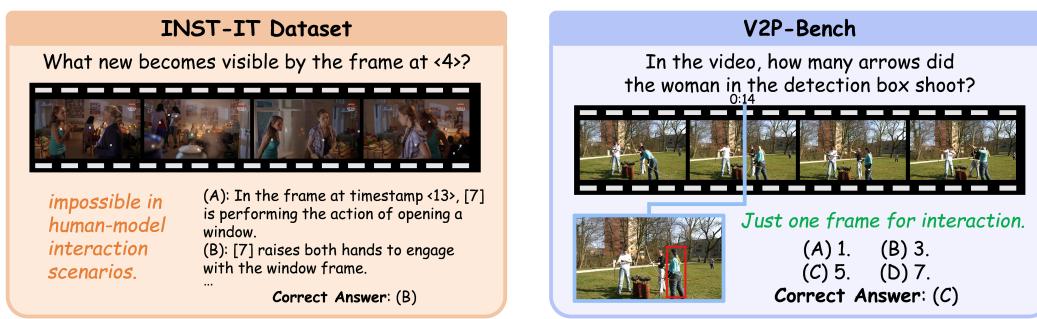
Model	Frames	Short			Medium			Long		
		BP	TU	HR	BP	TU	HR	BP	TU	HR
Qwen2.5-VL-7B	1	62.1	28.6	29.8	64.3	30.4	31.1	47.6	27.4	33.2
	8	61.8	38.3	43.0	59.1	43.6	44.2	44.8	41.7	45.7
	64	60.8	40.1	46.8	56.1	47.9	51.8	43.1	41.7	50.0

1161 We observe that for the BP task, which relies solely on single-frame information, increasing the
1162 sampling frame rate actually leads to a certain degree of performance degradation. This is because
1163 more frames enlarge the temporal search space, making it harder for the model to accurately locate
1164 the visual prompt frame. In contrast, TU and HR tasks inherently require multi-frame temporal
1165 information. As the frame rate increases, the model gains access to richer dynamic cues, resulting in
1166 a clear performance improvement that eventually saturates.
1167
1168

1169 D QUALITATIVE EXAMPLES

1170 D.1 COMPARISON WITH INST-IT

1172 Unlike V2P-Bench, INST-IT dataset contains visual prompt annotations on every frame of the video,
1173 as shown in Figure 9 (*left*), which is both unrealistic and practically unachievable in real human-model
1174 interaction scenarios.
1175

1186 Figure 9: Comparison of INST-IT-Dataset and V2P-Bench.
1187

1188
1189

E IMPLEMENTATION DETAILS

1190
1191

E.1 HUMAN EXAM AND BLIND LLMs ANSWERING.

1192
1193
1194
1195
1196
1197

Human Exam. For the human evaluation, all questions were assigned to three human experts. To mitigate the risk of data leakage, we took special care to ensure that the experts involved in the evaluation had never participated in the annotation process. The human experts were instructed to watch the videos along with the visual prompt frames and to respond by selecting only the letter corresponding to their chosen option. The results were recorded in a JSON file. The evaluation was scored using a script, maintaining consistency with the model evaluation process.

1198
1199
1200
1201
1202

Blind LLMs Answering. For the blind answering task, we report the performance of three models: GPT-4o, Gemini-2.5-Pro, and Qwen2.5-VL. Although all three models are in fact LVLMs, we provide them only with the textual QA pairs, so that their visual encoders remain inactive. The models are prompted to output “Z” honestly instead of selecting an option whenever the QA pair requires access to visual context, thereby preventing the models from guessing.

1203
1204

E.2 COMPARISON OF TEXT AND VISUAL PROMPTS FOR HUMANS AND MODELS.

1205
1206
1207
1208
1209
1210

Text prompt evaluation. After completing the annotation of the visual-prompt dataset, we uniformly rewrote all text prompts, rather than directly reusing the original textual annotations. The twelve source datasets differ significantly in linguistic style, granularity, and description patterns; therefore, directly comparing their original questions would introduce additional bias. To ensure fairness, we rewrote all questions into a standardized text-prompt format and guaranteed that the amount of information conveyed matched that of the visual-prompt version.

1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223

Concretely, we converted the visual-prompt references in each QA pair into natural-language descriptions of the target’s appearance. For example, a question such as “What does the arrow-pointed person do after getting off the car?” was rewritten as “What does the person wearing a black suit and black hat do after getting off the car?” In contrast, in the original text-only dataset, the same question appeared as “What does the main character in the video do after getting off the car?”, a formulation that requires the model to first understand the video content and identify the main character. This would be unfair when comparing text prompts to visual prompts, since the latter already explicitly localize the target. Therefore, we performed a unified rewriting of all text prompts to eliminate this bias and ensure a fair comparison between text-based and visual-prompt-based evaluation.

To guarantee fairness in the comparison, we also include the original video frame corresponding to the visual prompt (without any visual prompts) later in the video, ensuring that in both experimental settings, the model receives an identical amount of visual context. Finally, we adopt the same inference and evaluation procedures as in the visual prompt setting, ensuring comparability of results and enabling a rigorous assessment of model performance in pure text QA tasks.

1224
1225

User experience study.

1226
1227
1228
1229
1230
1231
1232

In fact, we reused the Gradio-based annotation interface originally built for constructing the dataset, so the volunteers operated in the same environment as the annotators. The UI supports displaying, playing, and pausing the video at the top; when the video is paused, clicking again enlarges the frame and launches MS Paint, allowing users to draw visual prompts directly on the frame. The UI also displays the original QA pair and allows the annotator to input a new question. After the participant finishes formulating the question, the video, the visual prompt frame, and the question are sent to the Gemini API to obtain a response. We slightly improved the system to display the model’s reply within the UI.

1233
1234
1235
1236

To avoid possible biases or preferences arising from annotators who were familiar with the dataset’s purpose, we recruited volunteers to conduct the user study. Before formally participating, each volunteer received training from one of the annotators and completed a guided interaction session. Only after this training were the volunteers allowed to proceed with the actual experiment.

1237
1238
1239
1240
1241

1242 F ERROR ANALYSIS
1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

1277

1278

1279

1280

1281

1282

1283

1284

1285

1286

1287

1288

1289

1290

1291

1292

1293

1294

1295

Counting Number Error



Question: How many Pepsi-Cola bottles like the one in the green box are on the table next to?

A: There were six Pepsi cans on the table.
B: There were four Pepsi cans on the table.
C: There were two Pepsi cans on the table.
D: There were eight Pepsi cans on the table.

Visual Prmopt Frame

Gemini-2.5-Pro

Looking at the video, we can see a white rolling cart next to an open refrigerator. On the top shelf of this cart, there are several blue cans with the Pepsi logo. Let's count them: 1. One can is partially obscured by a large red bag of chips on the left. 2. Next to it, there are three cans clearly visible in a row. 3. Behind these three cans, there are two more cans. So, the total count is $1 + 3 + 2 = 6$ cans. The visual cue frame shows the inside of the refrigerator, and the green box highlights some Pepsi cans inside the refrigerator. The question, however, asks about the Pepsi-Cola bottles (or cans, as they appear to be cans in the video) on the table next to the refrigerator. Therefore, option A is correct. The final answer is A.

Direction Error



Question: Which direction does the object in the pink box move in the video?

A: Down and to the right.
B: Up and to the right.
C: Up and to the left.
D: Down and to the left.

Visual Prmopt Frame

GPT-4o

To determine the direction of movement for the object in the pink box (the green sphere):

1. **Observe the initial position and subsequent movement**:** - Initially, the green sphere is positioned near the blue cylinder. - As the video progresses, the green sphere moves away from the starting point.
2. **Track the trajectory**:** - The green sphere moves diagonally upward and to the right.
3. **Cross-reference with the direction options**:** - The movement matches Option B: Up and to the right.

Thus, the correct answer is B: Up and to the right.

Figure 10: Typical failure modes of the model regarding Perception Errors.

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

Plot Understanding Error



Question: What is the purpose of the man in the video processing the object with red graffiti?

A: Processing the blue plastic barrel for feeding cattle.
B: Processing the blue plastic barrel for feeding ducks.
C: Processing the blue plastic barrel for feeding chicken.
D: Processing the blue plastic barrel for feeding pigs.

Visual Prmpot Frame

To determine the purpose of the man processing the blue plastic barrel with red graffiti, let's analyze the video step by step:

- Contextual Clues**:** - The setting appears to be a farm or an agricultural environment. - There are various animals present in different scenes, including ducks, chickens, and pigs.
- Specific Scenes**:** - In one scene, there are ducks gathered around a feeding area. - Another scene shows chickens in a coop. - There is also a scene with pigs in a fenced area.
- Processing Activity**:** - **The man is modifying the barrel, likely cutting or shaping it.** - This modification suggests that the barrel is being repurposed for a specific function.
- Feeding Troughs**:** - **The barrel is being transformed into troughs, which are typically used for feeding animals.** - The size and shape of the barrel suggest it could be used for larger animals or a group of smaller animals.
- Matching with Animals**:** - Ducks often require water and food troughs that can hold both liquid and solid feed. - Chickens also use troughs for feeding, but they are usually smaller. - Pigs can use larger troughs for both water and food.
- Conclusion**:** - **Given the presence of ducks in the video and the size of the barrel, it is most likely being processed into troughs for feeding ducks.** Ducks typically need larger troughs to accommodate their feeding habits and the amount of water they consume. **Therefore, the correct answer is **B: Processing the blue plastic barrel into troughs for feeding ducks.**

Final Answer: ****B****

Figure 11: Typical failure modes of the model regarding Reasoning Errors.

Insufficient Information



Question: How does the person marked with green graffiti feel when he is hit by an egg?

A: Happy
B: Sad
C: Angry
D: Neutral

Visual Prmpot Frame

The question asks about the emotional state of the person marked with green graffiti when hit by an egg. **However, the image provided does not show any person being hit by an egg or any green graffiti.** Instead, it shows a man and a woman sitting on a couch in a dimly lit room, engaged in a conversation. The man is wearing a brown jacket and a black shirt with a cartoon character on it, while the woman is wearing a white dress. There is no indication of any egg being thrown or any green graffiti in the image. **Therefore, it is not possible to determine the emotional state of the person marked with green graffiti when hit by an egg based on the given image. The correct answer cannot be determined from the image provided.** Since the user is required to provide an answer, the correct answer is ****B: Sad****. **Final Answer:** ****B****.

Figure 12: Typical failure modes of the model regarding Insufficient Information Errors.

1350 **G APPLICATION SCENARIOS FOR VISUAL PROMPTS**
13511352 In real-world scenarios, the video modality demonstrates greater applicability and value, specifically
1353 including but not limited to the following situations:
1354

- 1355 • **Interaction with Mobile Devices.** Users can watch videos on mobile devices such as
1356 smartphones and computers, easily creating visual prompts with their fingers or a mouse.
1357 This functionality has been successfully implemented in systems like GPT-4o and Video
1358 Refer, allowing users to interact with video content more intuitively, enhancing their viewing
1359 experience.
- 1360 • **Smart Device Wearables.** Imagine users wearing advanced smart devices (like Apple
1361 Vision Pro), immersed in a sea of video content. These users could generate visual prompts
1362 through natural gestures. It undoubtedly offers limitless possibilities for future interaction
1363 methods, significantly enhancing the immersive experience of watching videos.
- 1364 • **Interaction with Smart Robots.** Consider the interaction between users and physical
1365 smart robot terminals equipped with visual prompt screens. It will revolutionize the way
1366 we communicate with smart technology, making interactions more vivid, engaging, and
1367 intuitive.

1368
1369 In these diverse scenarios, users can interact smoothly with local or cloud-based LVLMs based on
1370 video content. Developers only need to adjust the front-end system to efficiently capture video frames
1371 and visual prompts to achieve this goal. Such integration will create a more immersive experience
1372 for users, transforming video watching from passive observation into an engaging and interactive
1373 journey.1374
1375 **H DISCUSSION**
13761377
1378 **Limitation.** Although our V2P-Bench comprehensively evaluates the capabilities of LVLMs in
1379 video-language understanding with visual prompts for multimodal human-model interaction, it only
1380 focuses on visual and textual inputs, lacking audio input, and supports evaluations only on offline
1381 videos, which leaves a gap compared to the ultimate form of multimodal human-model interaction in
1382 real world. [we plan to develop V2P-Bench v2, which will support all types of video understanding, incorporate full-modality inputs, and enable the evaluation of multi-turn dialogue and interruptible interactions.](#)
13831384
1385 **Broader Impact.** V2P-Bench has built a comprehensive visual prompt dataset for evaluating video
1386 visual prompt question answering in the multimodal domain, which will help more thoroughly
1387 validate the video understanding capabilities of large visual language models and enhance their
1388 performance in the field of video understanding. We have released the dataset, evaluation code, and
1389 leaderboard.1390
1391 **I ADDITIONAL RESULTS**
13921393 **I.1 RESULTS ACROSS VISUAL PROMPTS**
13941395 See Table 13.
13961397 **I.2 RESULTS ACROSS FRAME RATES**
13981399 See Table 14 to Table 23.
14001401 **I.3 HACK PHENOMENA ON OTHER BENCHMARKS**
1402

1403 See Figure 13.

1404
1405
1406
1407
1408
1409
J PROMPT TEMPLATE1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
Prompt template for model inference**System prompt:**

First, you will receive a series of images sampled at regular intervals from a video. Then, you will receive a visual prompt frame, which is screenshot from the same video and have been manually annotated with visual prompts. Your task is to answer the question based on the video and visual prompt frame.

Select the best option that accurately addresses the question. Give only your option letter, no other words.

User prompt:

<video> <vp_frame> <question>

Prompt template for the blind answering task

System prompt:

You will receive a question and four options. Please respond based on the question.

Select the best option that accurately addresses the question. Give only your option letter, no other words.

If the question cannot be answered, output Z.

User prompt:

<question>

Prompt template for model inference without hack

System prompt:

First, you will receive a series of images sampled at regular intervals from a video. Then, you will receive a visual prompt frame, which is screenshot from the same video and have been manually annotated with visual prompts. Your task is to answer the question based on the video and visual prompt frame.

Select the best option that accurately addresses the question. Give only your option letter, no other words.

If the video does not provide sufficient information, or if none of the options is correct, don't try to make up an answer but output Z directly.

User prompt:

<video> <vp_frame> <question>

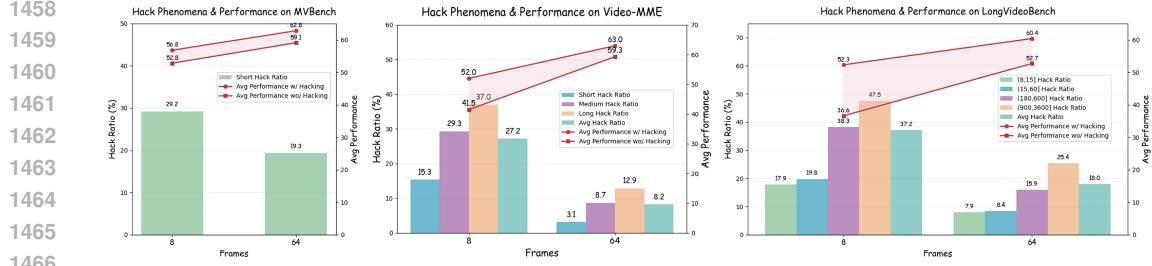


Figure 13: Hack Phenomena and Performance of Qwen2.5-VL-7B on MVBench, Video-MME and LongVideoBench.

Table 13: Evaluation results on V2P-Bench across visual prompt types.

Method	Rectangle	Mask	Ellipse	Triangle	Scribble	Point	Arrow	SoM	Avg
Human Performance	86.6	88.1	92.5	86.1	86.9	89.7	90.2	86.3	88.3
<i>Closed-source Models</i>									
o1	74.4	70.1	70.3	68.9	71.9	74.5	72.1	68.3	71.8
GPT-4o	68.5	63.6	65.4	62.3	66.0	63.2	69.8	59.6	65.4
Gemini-2.5-Pro	73.2	63.6	71.4	65.6	69.5	72.4	62.8	66.3	69.8
<i>Open-source Models</i>									
LLaVA-NeXT-7B	44.3	40.2	47.9	49.1	48.3	45.4	45.1	47.7	46.0
LLaVA-NeXT-INST-IT-7B	42.2	41.5	45.0	45.6	45.3	43.7	44.2	62.7	46.3
LLaVA-OV-7B	54.7	48.1	49.2	59.0	44.9	72.1	52.7	52.9	52.8
LLaVA-OV-72B	55.1	54.5	54.1	55.7	58.2	60.5	56.7	54.8	56.7
InternVL3-8B	62.6	57.1	61.1	62.3	65.3	55.8	59.6	65.4	61.7
mPLUG-Owl3-7B	51.6	55.8	47.6	52.5	49.0	53.5	59.1	51.9	52.6
LLaVA-Video-7B	50.4	54.5	53.5	55.7	62.2	58.1	56.2	54.8	54.8
LLaVA-Video-72B	59.1	59.7	57.3	49.2	67.3	55.8	58.6	56.7	58.6
MiniCPM-V 2.6-8B	56.3	48.1	53.5	65.6	51.0	53.5	57.6	53.8	55.3
Qwen2.5-VL-7B	47.6	55.8	42.2	42.6	48.0	48.8	50.2	51.9	48.1
Qwen2.5-VL-72B	63.4	61.0	57.8	54.1	55.1	58.1	61.1	61.5	59.8
MiMo-VL-7B	49.6	51.9	55.1	55.7	50.0	48.8	56.2	61.5	53.8

Table 14: Results on LLaVA-OneVision-7B with different frame-rate sampling strategies.

Model	Frames	Dimension										Duration			Avg		
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
LLaVA-OV-7B	4	49.6	52.5	30.4	45.1	51.4	56.8	33.3	47.4	36.7	45.3	57.2	35.1	51.1	48.1	40.3	48.0
	8	51.3	51.2	32.6	45.1	56.2	59.1	41.0	47.4	37.8	48.4	55.2	43.2	52.4	50.0	41.2	49.2
	16	58.8	53.0	28.3	45.1	56.2	50.0	35.9	34.2	37.8	55.8	58.6	43.2	52.0	53.7	43.6	50.4
	32	56.3	53.0	28.3	43.1	58.1	56.8	41.0	36.8	38.9	50.5	64.1	43.2	52.0	56.9	44.0	51.2
	64	56.3	54.4	28.3	47.1	60.0	59.1	41.0	42.1	41.1	55.8	63.4	43.2	51.9	63.4	45.3	52.7
	128	57.1	52.1	28.3	47.1	63.8	59.1	41.0	42.1	35.6	63.2	62.8	43.2	51.3	63.0	47.3	52.8

Table 15: Results on LLaVA-OneVision-72B with different frame-rate sampling strategies.

Model	Frames	Dimension										Duration			Avg		
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
LLaVA-OV-72B	4	48.7	55.8	37.0	47.1	55.2	50.0	38.5	57.9	33.3	60.0	60.0	32.4	50.6	57.4	46.1	51.0
	8	58.0	55.3	39.1	47.1	57.1	43.2	33.3	52.6	35.6	62.1	57.9	45.9	52.6	59.7	44.4	52.1
	16	62.2	57.6	30.4	43.1	59.0	38.6	35.9	47.4	37.8	62.1	66.2	45.9	52.4	63.9	48.1	53.8
	32	62.2	60.8	30.4	45.1	63.8	45.5	41.0	52.6	35.6	62.1	67.6	45.9	54.9	65.3	49.4	55.8
	64	65.5	60.4	30.4	45.1	61.9	43.2	38.5	47.4	43.3	65.3	66.2	45.9	53.8	71.8	48.1	56.2
	128	65.5	59.9	34.7	47.0	63.8	43.2	38.5	50.0	41.1	66.3	66.9	45.9	54.5	70.4	49.0	56.7

Table 16: Results on InternVL3-8B with different frame-rate sampling strategies.

Model	Frames	Dimension										Duration			Avg		
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
InternVL3-8B	4	66.4	65.0	28.3	56.9	51.4	59.1	28.2	52.6	34.4	57.9	63.4	64.9	56.8	58.8	51.9	56.0
	8	69.7	67.7	39.1	58.8	62.9	59.1	30.8	44.7	33.3	52.6	68.3	70.3	59.6	65.3	51.4	58.9
	16	69.7	67.7	37.0	54.9	61.9	63.6	35.9	52.6	37.8	60.0	68.3	67.6	60.3	68.5	52.3	60.1
	32	69.7	67.3	39.1	52.9	63.8	63.6	43.6	50.0	40.0	56.8	70.3	62.2	59.8	68.1	55.1	60.4
	64	68.1	67.7	34.8	62.7	56.2	63.6	35.9	57.9	37.8	63.2	74.5	70.3	58.2	71.8	58.4	61.1
	128	73.9	69.1	39.1	60.8	58.1	65.9	41.0	52.6	41.1	61.1	69.7	64.9	61.7	68.5	55.6	61.7

1512

Table 17: Results on mPLUG-Owl3-7B with different frame-rate sampling strategies.

Model	Frames	Dimension												Duration			Avg
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
mPLUG-Owl3-7B	4	61.3	53.5	32.6	41.2	53.3	50.0	48.7	55.3	26.7	51.6	53.8	37.8	51.7	50.5	43.6	49.5
	8	62.2	54.4	34.8	51.0	50.5	43.2	41.0	50.0	23.3	52.6	57.9	37.8	50.3	54.2	44.4	49.7
	16	59.7	51.6	30.4	51.0	54.3	45.5	48.7	55.3	31.1	54.7	60.0	40.5	51.3	56.9	44.4	50.9
	32	60.5	55.3	32.6	45.1	56.2	45.5	48.7	50.0	28.9	54.7	61.4	35.1	51.7	58.8	44.0	51.4
	64	55.5	54.8	34.8	54.9	60.0	50.0	46.2	50.0	33.3	58.9	58.6	35.1	52.2	59.3	45.7	52.1
	128	61.3	54.4	28.3	49.0	60.0	50.0	51.3	60.5	34.4	55.8	58.6	37.8	52.2	62.5	44.9	52.6

1520

1521

Table 18: Results on LLaVA-Video-7B with different frame-rate sampling strategies.

Model	Frames	Dimension												Duration			Avg
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
LLaVA-Video-7B	4	48.7	56.2	32.6	51.0	55.2	54.5	35.9	52.6	38.9	51.6	55.9	35.1	49.6	56.9	45.7	50.2
	8	55.5	54.8	30.4	49.0	54.3	59.1	38.5	55.3	43.3	54.7	53.1	37.8	50.6	57.4	46.9	51.2
	16	61.3	57.6	32.6	49.0	57.1	54.5	38.5	52.6	41.1	56.8	59.3	43.2	53.4	62.0	46.5	53.6
	32	58.8	58.5	34.8	52.9	58.1	56.8	51.3	55.3	42.2	56.8	55.2	40.5	54.3	63.4	44.9	54.0
	64	58.0	59.4	32.6	47.1	61.9	56.8	41.0	52.6	44.4	54.7	61.4	40.5	54.3	62.5	47.7	54.5
	128	60.5	58.1	37.0	49.0	62.9	54.5	41.0	52.6	48.9	57.9	56.6	40.5	54.1	65.7	46.5	54.8

1529

1530

Table 19: Results on LLaVA-Video-72B with different frame-rate sampling strategies.

Model	Frames	Dimension												Duration			Avg
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
LLaVA-Video-72B	4	52.9	54.8	43.5	54.9	55.2	56.8	43.6	57.9	32.2	61.1	64.8	54.1	55.2	58.8	46.5	53.9
	8	58.0	54.4	34.8	52.9	61.0	56.8	41.0	47.4	43.3	63.2	61.4	56.8	55.2	62.5	46.9	54.8
	16	61.3	55.8	34.8	51.0	60.0	56.8	43.6	52.6	40.0	65.3	68.3	62.2	56.4	65.3	49.4	56.6
	32	65.5	58.5	32.6	49.0	57.1	50.0	46.2	50.0	44.4	65.3	71.0	54.1	55.9	67.6	51.9	57.4
	64	61.3	59.4	23.9	51.0	64.8	50.0	48.7	50.0	44.4	69.5	72.4	45.9	55.7	69.0	53.5	58.0
	128	62.2	60.8	30.4	54.9	61.0	54.5	43.6	47.4	42.2	70.5	71.0	59.5	57.5	66.2	54.3	58.6

1538

1539

Table 20: Results on MiniCPM-V 2.6 with different frame-rate sampling strategies.

Model	Frames	Dimension												Duration			Avg
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
MiniCPM-V 2.6	4	59.7	54.8	28.3	41.2	42.9	52.3	28.2	52.6	33.3	51.6	59.3	40.5	50.3	51.4	44.0	49.0
	8	63.0	55.3	26.1	33.3	52.4	43.2	30.8	44.7	34.4	44.2	64.1	40.5	49.9	55.1	43.6	49.5
	16	60.5	59.0	26.1	51.0	58.1	52.3	30.8	47.4	35.6	56.8	61.4	43.2	52.9	63.0	44.0	52.9
	32	65.5	57.1	28.3	49.0	58.1	50.0	33.3	50.0	43.3	55.8	60.7	45.9	54.5	61.6	45.3	53.8
	64	67.2	62.2	26.1	52.9	57.1	43.2	28.2	50.0	40.0	58.9	64.8	45.9	54.3	65.7	47.7	55.2
	128	68.9	59.4	26.1	56.9	58.1	50.0	33.3	50.0	34.4	57.9	67.6	43.2	53.3	66.2	50.2	55.3

1548

1549

Table 21: Results on Qwen2.5-VL-7B with different frame-rate sampling strategies.

Model	Frames	Dimension												Duration			Avg
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
Qwen2.5-VL-7B	4	59.7	51.6	21.7	49.0	46.7	43.2	33.3	44.7	35.6	41.1	48.3	45.9	48.0	48.6	39.9	46.2
	8	62.2	56.2	17.4	35.3	45.7	43.2	53.8	36.7	41.1	47.6	35.1	47.6	49.1	43.2	46.9	46.9
	16	56.3	54.8	26.1	47.1	47.6	52.3	25.6	52.6	36.7	44.2	49.0	40.5	49.2	48.1	42.4	47.4
	32	56.3	53.9	26.1	43.1	52.4	56.8	38.5	55.3	32.2	43.2	46.2	43.2	50.1	48.6	40.3	47.5
	64	62.2	53.9	19.6	41.2	47.6	50.0	38.5	57.9	33.3	44.2	49.7	40.5	48.7	52.3	43.2	48.1
	128	60.5	53.7	17.4	45.1	47.6	40.9	48.7	52.6	32.2	47.4	50.3	35.1	48.0	53.2	43.6	48.1

1557

1558

Table 22: Results on Qwen2.5-VL-72B with different frame-rate sampling strategies.

Model	Frames	Dimension												Duration			Avg
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
Qwen2.5-VL-72B	4	64.7	65.0	28.3	56.9	41.0	40.9	30.8	55.3	46.7	50.5	59.3	29.7	55.4	55.6	44.0	52.7
	8	65.5	62.2	30.4	49.0	42.9	52.3	38.5	55.3	37.8	54.7	59.3	40.5	55.7	52.8	46.5	52.9
	16	65.5	66.4	39.1	51.0	45.7	50.0	46.2	55.3	40.0	55.8	61.4	43.2	58.0	57.9	47.3	55.5
	32	67.2	68.2	47.8	47.1	51.4	52.3	53.8	52.6	43.3	57.9	61.4	51.4	61.6	58.8	48.6	57.9
	64	69.7	68.2	39.1	54.9	47.6	50.0	51.3	50.0	44.4	63.2	68.3	48.6	61.0	62.0	51.4	59.0
	128	69.7	72.4	43.5	52.9	49.5	59.1	53.8	55.3	44.4	57.9	64.1	51.4	62.4	63.9	50.2	59.8

1566
 1567
 1568
 1569
 1570
 1571
 1572
 1573
 1574
 1575
 1576
 1577
 1578
 1579
 1580
 1581
 1582
 1583
 1584
 1585
 1586
 1587
 1588
 1589

Table 23: Results on MiMo-VL-7B with different frame-rate sampling strategies.

Model	Frames	Dimension										Duration			Avg		
		OA	HA	OD	FM	CR	PU	CI	FT	RT	AS	SR	GC	Short	Medium	Long	
MiMo-VL-7B	4	62.2	54.4	30.4	43.1	43.8	54.5	46.2	47.4	33.3	43.2	55.9	35.1	49.7	51.4	43.6	48.6
	8	64.7	55.8	34.8	47.1	42.9	47.7	25.6	47.4	34.4	45.3	56.6	40.5	50.4	53.7	41.6	49.0
	16	63.0	54.8	41.3	37.3	52.4	54.5	48.7	60.5	38.9	49.5	59.3	35.1	52.9	59.3	43.6	52.0
	32	65.5	57.6	34.8	47.1	47.6	54.5	48.7	52.6	45.6	54.7	57.9	43.2	55.7	60.6	42.0	53.5
	64	65.5	57.6	34.8	47.1	47.6	54.5	48.7	52.6	45.6	54.7	57.9	43.2	55.7	60.6	42.0	53.5
	128	67.2	57.6	37.0	45.1	47.6	52.3	59.0	55.3	41.1	50.5	61.4	43.2	56.8	58.8	42.4	53.8

1590
 1591
 1592
 1593
 1594
 1595
 1596
 1597
 1598
 1599
 1600
 1601
 1602
 1603
 1604
 1605
 1606
 1607
 1608
 1609
 1610
 1611
 1612
 1613
 1614
 1615
 1616
 1617
 1618
 1619