

SCP: Spatial Causal Prediction in Video

Supplementary Material

Contents

A Limitations	12
A.1 Limitations in Spatial Pattern Coverage . . .	12
A.2 Inherent Limits of Scene Variety	12
A.3 Data Scale Bounded by Quality Assurance .	12
A.4 Scope and Limitation of Causality in SCP . .	12
B Benchmark Construction Details	12
B.1. Video Sources	12
B.2. QA Candidates Generation Details	13
B.3. Manual Annotation Details	13
B.4. Quality Control	14
B.5. Statistics	14
B.6. Human Performance	14
C Task Definition	14
D Experimental Details	15
D.1. Model Setup	15
D.2 How Well Do Current Models Perform? . . .	15
D.2.1. Baseline Experiment.	15
D.3 What Affects Spatial Causal Prediction? . . .	15
D.3.1. Perception vs. Reasoning Decomposition.	15
D.3.2. Single-Frame vs. Multi-Frame Reasoning.	16
D.3.3. Vision Causal Perception.	17
D.3.4. Visible Range Comparison.	17
D.3.5. Is Visual Information Necessary? . .	17
D.4 How to Improve Spatial Causal Prediction? .	17
D.4.1. Model Scale-up Effect Analysis. . .	17
D.4.2. CoT Reasoning.	19
D.4.3. Self-Think Reasoning.	19
D.4.4. Perception Enhancement Strategy. . .	19
D.4.5. Causal Prediction Enhancement Strategy.	20
E Benchmark Comparison	21
F Supplementary Results	22
F.1. Extended Causal Consistency Experiment . .	22
F.2. Physical Commonsense Probing Experiment	22
F.3. Error Case Analysis	23

A. Limitations

Although SCP-Bench provides a rigorous and carefully structured evaluation of spatial causal prediction, it still entails several potential limitations simply because the domain itself is inherently complex and difficult to capture in full.

A.1. Limitations in Spatial Pattern Coverage

SCP-Bench spans a wide range of spatial patterns, yet the space of possible configurations in real-world videos is effectively unbounded. Some highly specialized or rare spatial structures are therefore naturally beyond the scope of the current benchmark.

A.2. Inherent Limits of Scene Variety

Although the dataset includes a broad collection of scenes, real-world environments exhibit far greater variability than any finite benchmark can capture. Certain niche or atypical scenarios inevitably remain under-represented.

A.3. Data Scale Bounded by Quality Assurance

SCP-Bench adopts a controlled data scale, as each item must pass rigorous checks on perspective, temporal clarity, and spatial precision. Meeting these requirements demands substantial validation effort, which keeps the dataset compact while preserving consistency and high fidelity.

A.4. Scope and Limitation of Causality in SCP

SCP-Bench examines models' SCP capability through inference over unseen spatial state transitions under partial observability. Explicit causal annotations, such as interventions, counterfactuals, or causal graphs, are not yet incorporated and remain for future work.

B. Benchmark Construction Details

B.1. Video Sources

To ensure that SCP-Bench provides comprehensive and reliable evaluation coverage, we carefully considered both viewpoint diversity and scene variability when selecting video sources. The main data sources include:

- **Ego-Exo4D** [21] provides synchronized egocentric and multiple aligned exocentric viewpoints across activities such as sports, household repairs, and everyday human interactions. Its multi-view alignment is essential for constructing multi-perspective questions in SCP-Bench.
- **HD-EPIC** [36] contributes high-resolution egocentric recordings focused primarily on kitchen environments.

These videos offer well-controlled, densely annotated footage but exhibit limited scene diversity due to their constrained capture settings.

- **YouTube-8M** [1] supplies large-scale, in-the-wild content spanning thousands of scene categories and diverse physical contexts. Its broad coverage helps enrich SCP-Bench with varied object configurations, dynamic motions, and unconstrained environments.
- **ActivityNet** [9] contains long-form, human-centric activity videos with rich temporal structure and complex multi-object interactions. These properties make it a valuable source for identifying segments with clear causal dynamics and predictable spatial evolution.

From these large-scale datasets, we perform careful, multi-stage filtering to extract clips that satisfy the requirements of spatial causal prediction. This ensures that the final benchmark maintains high-quality visual content, diverse physical scenarios, and balanced coverage across scene categories, perspectives, causal directions, and question types.

B.2. QA Candidates Generation Details

To construct a large pool of QA candidates, we designed a fully automated generation pipeline with a two-stage filtering procedure that narrows data from long videos to segments and then to short clips, since spatial causal reasoning questions typically concern only specific portions of a video. QA candidates are generated around these clips to ensure precise alignment with the underlying events and their temporal boundaries.

► **Segments Selecting.** Starting from event-level structure, we first collect the segment annotations provided by the original datasets and then use GPT-5 to propose additional plausible segments directly from the videos. To further control the quality of the selected segments, GPT-5 is prompted to (i) generate a detailed description for each candidate segment and (ii) filter these candidates based on multiple criteria, including spatial dynamics, event coherence, and content clarity. This procedure yields segments that feature well-defined events and distinct spatial changes.

► **Clips Selecting.** From the refined set of segments, we next identify the specific portions that exhibit meaningful spatial evolution. GPT-5 is prompted to propose candidate clip boundaries for each segment under the constraint that every clip must present a clear and coherent trajectory of spatial change, while avoiding rapid shot transitions, severe motion blur, and other forms of visual degradation. To ensure reliability, each proposed clip is then independently re-evaluated by GPT-5 using only the clip itself. Clips displaying weak spatial variation, ambiguous dynamics, or unstable visual quality are discarded. This process yields a curated collection of high-quality clips that are well-suited for generating SCP QA pairs.

► **QA Generation.** Given each selected clip, we design

a prompting scheme that guides GPT-5 to generate reliable spatial causal prediction questions. Because spatial reasoning is inherently sensitive to *perspective*, every question is required to explicitly specify its viewpoint, either a camera view or the perspective of a salient object, to avoid ambiguous interpretations. In line with our definition of spatial causal prediction, each question must also clearly identify the *unseen temporal region* that the model needs to infer from the visible portion of the clip. The *spatial question component* then specifies the core phenomenon being tested and is constrained to one of our eight spatial reasoning categories. To further guarantee answer determinism, the prompting template enforces precise spatial references, such as relative positions, motions, and object interactions, so that each question is tightly grounded in visual evidence and admits a single, verifiably correct answer.

Using this pipeline, GPT-5 produces a large set of high-quality spatial causal prediction QA candidates that serve as the foundation for SCP-Bench.

B.3. Manual Annotation Details

Following automatic QA generation, human annotators perform multiple rounds of selection, validation, and refinement to obtain the final version of SCP-Bench. To support this workflow, we develop 3 annotation tools that help annotators through (i) filtering QA candidates and defining cutpoints, (ii) validating retained items, and (iii) refining correctable items during annotation.

Filtering and Cutpoint Finding. Annotators first review the automatically generated candidates and determine which items satisfy the SCP task requirements. A key step in this stage is cutpoint identification, where annotators designate the timestamp separating the visible portion of the clip from the unseen region. These cutpoints ensure that the answer is not directly observable in the visible portion, yet remains inferable through spatial causal reasoning grounded in that visible evidence. As illustrated in Fig. 19, the filtering tool allows annotators to inspect each clip alongside its candidate questions, decide whether a question should be retained, and assign the corresponding cutpoint. The tool additionally provides a split-view interface, enabling annotators to preview the visible and unseen parts of the clip independently. For each retained question, GPT-5 subsequently generates additional distractor options to complete the multiple-choice set.

Validation and Optimization. Next, annotators use the validation tool (cf. Fig. 20) to simulate the actual SCP setting: each question must be answered strictly from the visible portion of the clip. This process verifies whether the question is answerable, viewpoint-consistent, and aligned with the intended spatial causal reasoning. Items that fail

validation are flagged along with a specific reason, e.g., “unclear perspective”, “answerable from visible region”, or “ambiguous spatial reference”. If a problematic item is deemed correctable, annotators use the repairing tool (Fig. 21) to edit its question wording, perspective specification, temporal boundary, or answer set. Through repeated rounds of validation and targeted refinement, annotators progressively resolve inconsistencies and remove invalid items, ultimately yielding the finalized SCP-Bench.

B.4. Quality Control

To ensure the overall quality and reliability of SCP-Bench, we employ both a rigorous automated pipeline and a carefully managed human annotation process. In the **automated stage**, we leverage GPT-5, one of the strongest commercially available models, and implement a multi-stage filtering workflow that enforces strict constraints on segment selection, clip extraction, and QA generation, thereby ensuring the initial candidate pool is of high quality. In the **human annotation stage**, all annotation tasks are conducted by seven well-trained annotators who have undergone extensive instruction on the SCP task definition, selection rules, and evaluation criteria. As detailed in Sec. B.3, we design a suite of specialized annotation tools that support efficient inspection, consistent decision-making, and systematic refinement of QA items. During the **validation stage**, annotators must provide brief justifications for their decisions, documenting why an item is accepted, rejected, or requires revision. For items that remain ambiguous, additional annotators re-examine the case and engage in discussion to resolve disagreements. If consensus cannot be reached, a voting procedure is used to determine whether the item should be retained. To maintain consistent quality over time, we also conduct periodic acceptance checks and random audits of annotator decisions, ensuring that annotation standards are uniformly upheld throughout the process.

B.5. Statistics

Through this carefully designed construction pipeline, we obtain SCP-Bench with 2,500 QA pairs across 1,181 video clips. The benchmark spans four major attributes: question type, causal direction, perspective setting, and scene category, as shown in Table 4.

B.6. Human Performance

Spatial causal reasoning is an intuitive capability that humans routinely exercise in daily life. To clearly quantify the gap between current models and human cognition, we conduct a human evaluation on SCP-Bench. Participants are presented with the same task setting as the models: only the visible portion of each clip and the corresponding question are provided, and they must select the most appropriate answer from the multiple-choice options. The aggregated

Category	Video Clip	QA Instance
• Scene Category		
Sports Related	467	964
Life Records	289	691
Driving-view Related	122	308
Animal Related	85	143
Artistic Performances	64	95
Factory/Machine Related	48	149
Others	106	150
• Perspective		
Ego	205	548
Exo	943	1546
Exo-Exo	14	80
Ego-Exo	14	44
Exo-Ego	12	66
• Causal Direction		
Forward	1053	2037
Backward	368	463
• Question Type		
Planning	417	491
Relation	641	874
Relative Distance	322	376
Spatial State	270	278
Appearance Order	162	167
Relative Speed	153	155
Counting	117	117
Relative Size	42	42

Table 4. Distribution of benchmarks over the scene category, perspective, Causal direction, and question type. These attributes are combined to form the final dataset, which consists of 1,181 video clips and 2,500 QA pairs.

accuracies provide an estimate of human performance on SCP-Bench. To ensure the validity and reliability of the human results, we recruit seven participants and enforce a strict “first-exposure” rule, i.e., each participant only answers questions they had never seen before, preventing any prior knowledge from influencing the outcomes.

C. Task Definition

We define 8 task categories to cover the key variations in spatial causal structure:

- **Appearance Order**: infers which object or action becomes relevant earlier in the unfolding sequence.
- **Counting**: tracks cardinality changes driven by underlying spatial dynamics.
- **Planning**: probes short-horizon action tendencies and anticipated movement choices.
- **Relation**: concerns the temporal variation of spatial configurations between objects.
- **Relative Distance**: focuses on how the spatial distance between objects changes over time.

- **Relative Size:** examines the comparative object scale as perceived through the evolving viewpoint.
- **Relative Speed:** concerns how objects move at different speeds and how those speeds change over time.
- **Spatial State:** evaluates how an object’s posture or interaction state is expected to change.

D. Experimental Details

D.1. Model Setup

Our full evaluation process covers 29 models spanning proprietary, open-source, and spatial-reasoning-specialized systems, enabling a broad and comprehensive assessment of current MLLMs. For proprietary models, we include the latest frontier releases, GPT-5 [40], Gemini-2.5-Pro [14], Gemini-2.5-Flash [14], and Claude-Sonnet-4.5 [3], which represent the most advanced commercially available multimodal reasoning models. For open-source models, we evaluate a diverse set of leading vision-language systems with different model sizes, including Qwen3-VL [54], Qwen3-Omni [53], InternVL3.5 [45], MiniCPM-V-4.5 [62], DeepSeek-VL2 [52], NVILA [32], LLaVA-OneVision [2, 29], and LLaVA-NeXT-Video [66]. We additionally include Spatial-MLLM [49] and SpaceR [35], which are explicitly designed to enhance spatial and relational reasoning. All models are evaluated using their default inference configurations to ensure consistent and fair comparison. For models that cannot directly process video inputs, we adopt a standardized preprocessing strategy: sampling one frame per second and feeding the frames sequentially according to their temporal order. In the following sections, we provide a detailed description of the in-depth analysis settings and present the corresponding experimental results.

D.2. How Well Do Current Models Perform?

D.2.1. Baseline Experiment.

To test the current performance of MLLM on SCP-Bench, we strictly follow the setting in our SCP task design:

- **Baseline:** For single-view questions, the model is given the visible part of a video clip, a question, and options. For multi-view questions, an image indicating the perspective to answer is additionally provided.

The exact prompt templates used in our experiments are shown below for clarity:

Baseline Prompt (Single-View)

Based on the provided video clip, please answer the following question. Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)
Question: {question}

Options: {options}

Baseline Prompt (Multi-View)

You are given a video and an image from the question perspective. Based on the events in the video clips, answer the question as if you are the camera in the REFERENCE IMAGE (not the video camera). Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

D.3. What Affects Spatial Causal Prediction?

D.3.1. Perception vs. Reasoning Decomposition.

This experiment examines the relative importance of visual perception and spatial causal reasoning, aiming to identify the key bottleneck in SCP performance. We evaluate Qwen3-VL-8B, InternVL3.5-8B, and MiniCPM-V-4.5 under the following evaluation settings:

- **Gold Video (Pure Perception):** The model directly observes the unseen part of the video clip (the gold video), removing the need for causal reasoning.
- **Dense Caption Only (Pure Reasoning):** Replace video input with its dense caption generated by Tarsier [44], removing the need for video perception.
- **Dense Caption Auxiliary:** The model receives both the video and its dense caption, intending to reduce the perceptual burden.

As illustrated in Table 5, the results show that causal reasoning plays a more critical role in SCP performance than raw video perception. In addition, supplying dense captions alongside videos leads to further accuracy improvements, suggesting that current models still face limitations in dynamic spatial perception in videos.

Dense Caption Auxiliary Prompt (Single-View)

Based on the provided video clip, please answer the following question. Choose the most appropriate option from the given choices. A dense caption is provided to help understand the video. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Dense Caption: {dense caption}

Table 5. Evaluation results for the ‘‘Perception vs. Reasoning Decomposition’’ experiment. ‘‘Avg.’’ indicates the overall average accuracy. Red and green indicate performance gains and drops relative to the Baseline setting, respectively.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• Gold Video									
Qwen3-VL-8B	54.96 ^{↑7.44}	62.28	58.12	59.47	48.17	54.26	83.33	57.42	57.91
InternVL3.5-8B	58.64 ^{↑8.12}	68.86	57.26	64.97	51.37	60.37	80.95	60.00	58.27
MiniCPM-V-4.5	43.72 _{↓0.08}	48.50	47.01	44.40	35.47	50.80	78.57	54.19	43.53
• Dense Caption Only									
Qwen3-VL-8B	46.76 _{↓0.76}	51.50	37.61	53.97	40.16	48.94	80.95	60.00	40.29
InternVL3.5-8B	41.24 _{↓9.28}	47.31	34.19	46.23	35.01	42.02	78.57	51.61	38.85
MiniCPM-V-4.5	45.76 ^{↑1.96}	58.08	41.88	50.51	36.04	50.00	83.33	58.06	43.88
• Dense Caption Auxiliary									
Qwen3-VL-8B	47.72 ^{↑0.20}	54.49	48.72	52.75	41.76	46.01	88.10	53.55	46.04
InternVL3.5-8B	50.56 ^{↑0.04}	60.48	53.85	54.18	43.02	56.12	64.29	60.65	45.32
MiniCPM-V-4.5	44.28 ^{↑0.48}	56.29	45.30	51.73	33.64	46.28	85.71	54.84	42.09

Table 6. Evaluation results for the ‘‘Single-Frame vs. Multi-Frame Reasoning’’ experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• Cutpoint Image									
Qwen3-VL-8B	49.52 ^{↑2.00}	49.70	52.14	50.92	44.62	50.27	83.33	55.48	51.80
InternVL3.5-8B	51.92 ^{↑1.40}	67.07	52.99	52.14	44.16	57.18	76.19	61.29	50.36
MiniCPM-V-4.5	43.80 _{↓0.00}	52.69	55.56	44.60	35.35	50.80	76.19	50.97	40.29
Spatial-MLLM	42.04 ^{↑2.28}	53.46	32.73	38.90	37.30	53.55	71.43	55.69	33.33

Table 7. Evaluation results for the ‘‘Vision Causal Perception’’ experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• Reverse Video									
Qwen3-VL-8B	46.76 _{↓0.76}	52.69	52.14	46.84	43.94	46.54	90.48	48.39	42.45
InternVL3.5-8B	48.56 _{↓1.96}	56.89	56.41	49.49	42.33	53.19	64.29	59.35	43.53
MiniCPM-V-4.5	42.24 _{↓1.56}	49.10	45.30	41.96	35.13	48.67	73.81	50.97	41.37
LLaVA-OneVision-1.5-8B	45.32 _{↓0.20}	52.69	47.01	45.42	39.47	50.00	80.95	49.03	44.60

Dense Caption Auxiliary Prompt (Multi-View)

You are given a video and an image from the question perspective. Based on the events in the VIDEO CLIP, answer the question as if you are the camera in the REFERENCE IMAGE (not the video camera). Choose the most appropriate option from the given choices. A dense caption is provided to aid in understanding the video. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Dense Caption: {dense caption}

D.3.2. Single-Frame vs. Multi-Frame Reasoning.

Video contains richer spatial-temporal cues than an image. To test whether models truly exploit temporal continuity rather than treating frames independently, we compare a single-frame setting with the baseline setting, which uses multiple frames, on Qwen3-VL-8B, InternVL3.5-8B, MiniCPM-V-4.5, and Spatial-MLLM:

- **Baseline (Multi-Frame):** As described in Sec. D.2.1.
- **Cutpoint Image (Single-Frame):** Replace video input with the frame at the cutpoint.

As shown in Table 6, the single-frame input yields slightly better performance across all four models. This suggests that multi-frame inputs introduce noise rather than useful temporal cues, indicating that models do not effectively use temporal information. The single-frame gains further reveal a bias toward image-based perception and limited exploitation of dynamic spatial-temporal signals.

Single Frame Prompt (Single-View)

Based on the provided image, please answer the following question. Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Single Frame Prompt (Multi-View)

You are given a reference image (the first) and an image from the question perspective (the second image). Based on the events in the first image, answer the question as if you were the camera in the second image. Choose the most appropriate option from the

given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

D.3.3. Vision Causal Perception.

To examine whether models truly leverage temporal–spatial cues for spatial causal inference, we conduct an experiment in which the input videos are played in reverse. Specifically, for Qwen3-VL-8B, InternVL3.5-8B, MiniCPM-V-4.5, and LLaVA-OneVision-1.5-8B, we introduce the following variant:

- **Reverse Video:** Reverse the input video while keeping all other settings identical to the baseline experiment.

As shown in Table 7, model performance exhibits only minor degradation under reversal, suggesting that these systems make limited use of the underlying spatial–causal structure present in video dynamics.

D.3.4. Visible Range Comparison.

Unlike standard VideoQA settings, SCP limits models to the visible portion of a clip, making it essential to examine how different visible ranges influence model performance. We therefore evaluate three input conditions on Qwen3-VL-8B, InternVL3.5-8B, MiniCPM-V-4.5, and Spatial-MLLM:

- **Baseline (Seen Part):** As described in Sec. D.2.1.
- **Gold Video (Unseen Part):** As described in Sec. D.3.1, the model directly sees the unseen part.
- **Full Video:** The model is given the entire video clip.

As shown in Table 8, all models except MiniCPM-V-4.5 achieve their highest accuracy under the gold video setting, followed by full video, with the baseline performing the worst. MiniCPM-V-4.5, however, performs best with full video. These results indicate that when the answer is directly observable, the task places minimal demands on spatial causal reasoning, leading to higher accuracy. For most models, the full video introduces additional, potentially distracting content beyond what is strictly necessary, resulting in reduced performance compared with the gold video.

D.3.5. Is Visual Information Necessary?

Modern MLLMs exhibit strong text-based abilities, motivating an examination of whether visual information is needed for SCP. To investigate this, we consider two conditions: answering SCP questions using their internal priors and only a textual description of the video. We evaluate Qwen3-VL-8B, InternVL3.5-8B, MiniCPM-V-4.5, and LLaVA-OneVision-1.5-8B under the following settings:

- **Pure Text (Internal Priors):** Answering only the textual question without any visual information.
- **Dense Caption Only (Textual Information):** As described in Sec. D.3.1.

- **Baseline (Visual Information):** The baseline setting as described in Sec. D.2.1.

The results in Table 9 indicate that when models rely solely on their intrinsic prior knowledge (i.e., pure text input), their performance drops substantially. Although supplying dense captions leads to noticeable improvements, the gains fall short of the baseline that uses video inputs. These findings suggest that textual priors alone are insufficient for SCP, and even dense captions cannot fully capture the richness, continuity, and fine-grained cues encoded in visual observations.

PureText Prompt

Please answer the following question. Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Dense Caption Only Prompt (Single-View)

Based on the provided dense caption description, please answer the following question. Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Dense Caption: {dense caption}

Dense Caption Only Prompt (Multi-View)

You are given a video dense caption and a reference image dense caption describing another viewpoint. Based on the dense captions, answer the question as if you were the camera in the reference image. Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Video Dense Caption: {video dense caption}

Image Dense Caption: {image dense caption}

D.4. How to Improve Spatial Causal Prediction?

D.4.1. Model Scale-up Effect Analysis.

Increasing model size is a direct way to improve performance. To examine how parameter scale affects SCP, we evaluate two model families with rich size variations, Qwen3-VL and InternVL-3.5, covering 14 models as listed

Table 8. Evaluation results for the “Visible Range Comparison” experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Baseline</i>									
Qwen3-VL-8B	47.52	54.49	51.28	49.29	42.33	49.47	90.48	46.45	46.40
InternVL3.5-8B	50.52	59.88	54.70	54.79	43.82	54.52	61.90	58.71	44.96
MiniCPM-V-4.5	43.80	53.29	49.57	43.99	36.04	49.20	76.19	52.26	42.81
Spatial-MLLM	39.76	45.51	28.21	33.81	38.33	49.73	66.67	50.97	32.37
• <i>Gold Video</i>									
Qwen3-VL-8B	54.96 ^{↑7.44}	62.28	58.12	59.47	48.17	54.26	83.33	57.42	57.91
InternVL3.5-8B	58.64 ^{↑8.12}	68.86	57.26	64.97	51.37	60.37	80.95	60.00	58.27
MiniCPM-V-4.5	43.72 _{↓0.08}	48.50	47.01	44.40	35.47	50.80	78.57	54.19	43.53
Spatial-MLLM	41.12 ^{↑1.36}	37.73	24.79	34.21	32.38	47.07	66.67	49.03	28.78
• <i>Full Video</i>									
Qwen3-VL-8B	54.56 ^{↑7.04}	61.08	59.83	61.10	48.40	53.19	85.71	54.84	53.24
InternVL3.5-8B	57.00 ^{↑6.48}	67.66	60.68	66.80	47.83	57.18	76.19	61.29	55.04
MiniCPM-V-4.5	52.68 ^{↑8.88}	63.47	72.65	53.97	45.31	53.72	73.81	60.65	49.64
Spatial-MLLM	41.12 ^{↑1.36}	47.10	32.73	39.92	37.76	49.20	71.43	52.69	32.73

Table 9. Evaluation results for the “Is Visual Information Necessary?” experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Pure Text</i>									
Qwen3-VL-8B	42.00 _{↓5.52}	49.10	35.04	47.25	34.67	46.28	69.05	40.65	45.32
InternVL3.5-8B	44.72 _{↓5.80}	50.90	42.74	48.88	36.61	52.66	73.81	51.61	41.01
MiniCPM-V-4.5	43.32 _{↓0.48}	53.29	44.44	46.03	35.01	51.60	69.05	52.90	37.77
LLaVA-OneVision-1.5-8B	40.20 _{↓5.32}	43.11	41.03	42.16	33.41	47.87	73.81	38.71	41.37
• <i>Dense Caption Only</i>									
Qwen3-VL-8B	46.76 _{↓0.76}	51.50	37.61	53.97	40.16	48.94	80.95	60.00	40.29
InternVL3.5-8B	41.24 _{↓9.28}	47.31	34.19	46.23	35.01	42.02	78.57	51.61	38.85
MiniCPM-V-4.5	45.76 ^{↑1.96}	58.08	41.88	50.51	36.04	50.00	83.33	58.06	43.88
LLaVA-OneVision-1.5-8B	43.12 _{↓2.40}	49.70	29.91	50.31	36.04	44.41	69.05	53.55	42.81

Table 10. Evaluation results for the “Model Scale-up Effect Analysis” experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Qwen3-VL Series</i>									
Qwen3-VL-2B	43.04	41.92	42.74	45.01	40.85	44.41	59.52	47.10	40.65
Qwen3-VL-4B	50.24	58.68	47.86	52.55	44.51	55.05	90.48	60.00	42.09
Qwen3-VL-8B	47.52	54.49	51.28	49.29	42.33	49.47	90.48	46.45	46.40
Qwen3-VL-30B-A3B	54.16	65.27	52.14	54.79	46.22	56.65	85.71	66.45	57.19
Qwen3-VL-32B	56.84	59.88	51.28	58.66	52.63	57.98	90.48	67.10	55.04
Qwen3-VL-235B-A22B	61.04	67.07	54.70	60.90	55.03	63.03	97.62	74.84	63.31
• <i>InternVL3.5 Series</i>									
InternVL3.5-1B	35.16	29.94	43.59	37.07	30.66	39.36	38.10	49.03	31.65
InternVL3.5-2B	40.84	44.91	44.44	45.01	37.07	40.16	61.90	46.45	35.97
InternVL3.5-4B	46.12	57.49	47.86	44.60	42.79	47.61	76.19	49.68	43.17
InternVL3.5-8B	50.52	59.88	54.70	54.79	43.82	54.52	61.90	58.71	44.96
InternVL3.5-14B	50.28	59.88	57.26	54.79	42.56	54.26	69.05	63.23	42.45
InternVL3.5-30B-A3B	52.48	65.87	52.99	56.01	43.25	56.12	80.95	67.10	49.64
InternVL3.5-38B	53.56	62.28	53.85	56.01	46.34	57.98	90.48	65.81	48.20
InternVL3.5-241B-A28B	56.96	67.07	60.68	61.10	46.11	60.37	90.48	68.39	60.07

Table 11. Evaluation results for the “CoT Reasoning” experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Vanilla CoT</i>									
Qwen3-VL-8B	51.52 ^{↑4.00}	63.47	50.43	51.53	44.39	54.52	90.48	61.29	51.80
InternVL3.5-8B	50.36 _{↓0.16}	59.28	51.28	49.29	43.02	56.65	61.90	65.16	51.08
MiniCPM-V-4.5	44.88 ^{↑1.08}	55.09	40.17	48.68	35.47	53.19	73.81	59.35	39.93
LLaVA-OneVision-1.5-8B	42.88 _{↓2.64}	48.50	39.32	46.44	33.30	53.99	69.05	54.19	39.57

in Table 10. Overall, scaling up is an effective strategy for SCP. The largest Qwen3-VL-235B-A22B model achieves an improvement of 18% over the smallest Qwen3-VL-2B model. Similarly, the largest InternVL3.5-241B-A28B model outperforms the smallest InternVL-1B model

by 21.81%. However, scaling is not strictly monotonic among the smaller models. For example, Qwen3-VL-8B performs 2.72% worse than Qwen3-VL-4B. This suggests that the benefits of scaling become stable only when the increase in parameters reaches a more substantial range.

D.4.2. CoT Reasoning.

To investigate whether chain of thought (CoT) prompting can improve performance on SCP, we apply a vanilla CoT setting to Qwen3-VL-8B, InternVL3.5-8B, MiniCPM-V-4.5, and LLaVA-OV-1.5-8B:

- **Vanilla CoT:** Augment the prompt with “Think step by step” to encourage the model to generate its chain of thought reasoning.

As shown in Table 11, Qwen3-VL-8B gains about 4% in accuracy, while the other three models show no clear improvement. This indicates that vanilla CoT prompting does not consistently enhance SCP performance.

Vanilla CoT Prompt (Single-View)

Based on the provided video clip, please think step by step to answer the following question. Show your reasoning process explicitly, and at the end, provide your final choice marked as ‘Final Answer:’. Your final answer should contain only the option text, not its letter.

Question: {question}

Options: {options}

Vanilla CoT Prompt (Multi-View)

You are given a video and a reference image showing another viewpoint. Based on the events in the VIDEO CLIP, reason step by step as if you were the camera in the REFERENCE IMAGE, not the video camera. Show your reasoning process explicitly, and at the end, provide your final choice marked as ‘Final Answer:’. Your final answer should contain only the option text, not its letter.

Question: {question}

Options: {options}

InternVL3.5 Thinking Mode System Prompt

You are an AI assistant that rigorously follows this response protocol:

1. First, conduct a detailed analysis of the question. Consider different angles, potential solutions, and reason through the problem step-by-step. Enclose this entire thinking process within `<think>` and `</think>` tags.
2. After the thinking section, provide a clear, concise, and direct answer to the user’s question. Separate the answer from the think section with a new-line.

Ensure that the thinking process is thorough but remains focused on the query. The final answer

should be standalone and not reference the thinking section.

D.4.3. Self-Think Reasoning.

Besides chain-of-thought prompting, some models are trained on data with explicit reasoning traces, which encourages latent internal reasoning during inference time. To evaluate whether self-thinking improves SCP performance, we configure the thinking modes of different models as follows:

- **Self-Thinking:** each model is run with its official thinking mode. Qwen3-VL-8B uses the Qwen3-VL-8B-Thinking variant; InternVL3.5-8B enables thinking via its system prompt; MiniCPM-V-4.5 activates thinking with `enable_thinking=true`.

The results in Table 12 show that enabling self-think leads to a slight performance drop compared with the baseline. This suggests that the current self-think mechanisms do not provide useful spatial causal reasoning signals and may even introduce additional noise that harms performance.

D.4.4. Perception Enhancement Strategy.

Although perception is not the main bottleneck for SCP, previous experiments indicate that models still struggle with some dynamic spatial-temporal cues. We therefore explore 2 strategies that modestly enhance perceptual input:

- **Dense Caption Auxiliary:** As described in Sec. D.3.1.
- **Spatial Interaction Graph:** The model is guided to identify key objects and the key background elements, from which a spatial relation graph and an interaction graph are constructed. The model answers the question based on this structured spatial representation.

The results in Table 13 show that both enhancements yield only marginal improvements. This indicates that perception-level augmentation alone is insufficient to meaningfully improve SCP performance.

Spatial-Interaction Graph Guiding Prompt

Based on the provided video clip, you should `{temporal_direction}` to answer the question. `{view_text}` Please answer the question based on the following framework and think step by step to analyze through each part. Your final answer should contain only the option text, not its letter.

Framework:

Part 1. Carefully analyze the video to identify the key spatial changing objects and their affected environmental elements that need to be addressed in answering the questions. You can construct triples of the form:

Table 12. Evaluation results for the ‘‘Self-Think Reasoning’’ experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Self-Thinking</i>									
Qwen3-VL-8B	47.44 _{↓0.08}	55.69	45.30	50.10	38.79	51.86	83.33	58.71	48.20
InternVL3.5-8B	47.84 _{↓2.68}	59.28	51.28	51.12	39.36	53.46	66.67	54.19	46.40
MiniCPM-V-4.5	42.72 _{↓1.08}	49.10	48.72	40.94	35.35	50.53	69.05	50.32	43.88

Table 13. Evaluation results for the ‘‘Perception Enhancement Strategy’’ experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Dense Caption Auxiliary</i>									
Qwen3-VL-8B	47.72 _{↑0.20}	54.49	48.72	52.75	41.76	46.01	88.10	53.55	46.04
InternVL3.5-8B	50.56 _{↑0.04}	60.48	53.85	54.18	43.02	56.12	64.29	60.65	45.32
MiniCPM-V-4.5	44.28 _{↑0.48}	56.29	45.30	51.73	33.64	46.28	85.71	54.84	42.09
LLaVA-OneVision-1.5-8B	45.80 _{↑0.28}	53.89	46.15	49.08	39.13	47.61	83.33	54.84	42.81
• <i>Spatial Interaction Graph</i>									
Qwen3-VL-8B	51.76 _{↑4.24}	69.46	46.15	52.95	45.08	55.05	78.57	70.32	43.53
InternVL3.5-8B	48.80 _{↓1.72}	59.88	45.30	49.49	38.67	55.85	78.57	67.10	50.00
MiniCPM-V-4.5	43.88 _{↑0.08}	59.88	42.74	47.66	32.49	48.14	76.19	59.35	44.60
LLaVA-OneVision-1.5-8B	42.16 _{↓3.36}	52.69	44.44	43.18	33.87	47.07	76.19	49.68	43.17

Table 14. Evaluation results for the ‘‘Causal Prediction Enhancement Strategy’’ experiment.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Text Scaffold</i>									
Qwen3-VL-8B	61.23 _{↑13.71}	64.29	56.67	67.44	54.48	64.08	91.89	64.06	56.25
InternVL3.5-8B	61.72 _{↑11.20}	64.29	57.78	67.44	53.25	65.49	91.89	67.97	59.13
MiniCPM-V-4.5	60.15 _{↑16.35}	58.16	52.22	66.28	51.32	64.08	97.30	67.19	59.62
LLaVA-OneVision-1.5-8B	59.77 _{↑14.52}	55.10	53.33	66.05	50.09	62.68	89.19	71.09	62.02
• <i>Image Scaffold</i>									
Qwen3-VL-8B	49.21 _{↑1.69}	46.94	36.67	57.27	43.76	48.24	91.89	50.78	46.63
InternVL3.5-8B	51.60 _{↑1.08}	52.04	43.33	58.20	47.10	52.46	72.97	56.25	45.67
MiniCPM-V-4.5	43.64 _{↓0.16}	48.98	53.33	44.57	36.91	46.48	81.08	50.78	38.46
LLaVA-OneVision-1.5-8B	46.13 _{↑0.61}	56.12	35.56	52.42	38.84	48.24	83.78	45.31	43.75
• <i>Video Scaffold</i>									
Qwen3-VL-8B	52.63 _{↑5.11}	51.02	45.56	57.51	48.86	51.76	94.59	52.34	50.48
InternVL3.5-8B	53.38 _{↑2.86}	56.12	47.78	59.35	47.98	53.17	86.49	57.03	49.04
MiniCPM-V-4.5	44.29 _{↑0.49}	47.96	43.33	45.03	37.79	50.35	83.78	51.56	39.42
LLaVA-OneVision-1.5-8B	47.86 _{↑2.34}	51.02	45.56	51.96	41.65	52.11	81.08	46.88	44.71

```
{ "changing_objects": [ "..."],
  "affected_environmental_elements":
  [ "..."] } Only list the objects and environmental
elements that need to be addressed in answering
the question.
```

Part 2. Construct the spatial graph that describes the spatial relations among the entities, which is helpful for reasoning about the question. You can construct triples of the form: {"subject": "...", "relation": "...", "object": "..."} to describe the spatial relations and trends of changes during the video clip.

Part 3. Construct the interaction graph that describes the interaction relations among the entities, which is helpful for reasoning about the question. You can construct triples of the form: {"subject": "...", "action": "...", "object": "..."} to describe the

interaction during the video clip.

Part 4. {temporal_text}

Part 5. Based on previous thinking and constructed, think through and choose the most spatio-temporally consistent option.

Rules:

- Show your reasoning process explicitly, and at the end provide your final choice marked as 'Final Answer:'.
- Your final answer should contain only the option text, not its letter.
- Don't get caught up in repetitive reasoning.

Question: {question}

Options: {options}

D.4.5. Causal Prediction Enhancement Strategy.

To investigate whether spatial causal prediction scaffolds can mitigate the reasoning bottleneck in SCP, and which modality provides the greatest benefit, we evaluate models

Table 15. Comparison with existing benchmarks. Modality denotes the input type. Dynamic/Static indicates whether the scene content itself undergoes temporal changes, rather than mere camera movement. View Type specifies whether the questions involve single or multiple viewpoints. Perspective denotes whether the benchmark includes explicitly designed ego (first-person) and exo (third-person) perspective settings. Causal Reasoning indicates whether the benchmark requires inferring outcomes or states driven by causal dependencies. Seen/Unseen reflects whether the queried information is directly observable within the given visual content.

Benchmark	QA pairs	Modality	Dynamic/Static	View Type		Perspective		Causal Reasoning	Seen/Unseen
				Single	Multi	Ego	Exo		
3DSRBench [33]	3,772	Image	Static	✓	✗	✗	✓	✗	Seen
InternSpatial-Bench [15]	6,008	Image	Static	✓	✗	✗	✓	✗	Seen
OmniSpatial [27]	~8,400	Image	Static	✓	✗	✓	✓	✗	Seen
Spatial457 [46]	23,752	Image	Static	✓	✗	Undeclared		✗	Seen
All-Angles Bench [57]	~2,100	Image	Static	✗	✓	Undeclared		✗	Seen
EmbSpatial-Bench [16]	3,640	Image	Static	✓	✗	✓	✗	✗	Seen
MMSI-Bench [56]	1,000	Image	Static	✗	✓	Undeclared		✗	Seen
MindCube [59]	21,154	Image	Static	✗	✓	Undeclared		✗	Unseen
VSI-Bench [55]	5,130	Video	Static	✓	✗	✓	✗	✗	Seen
VLM4D [69]	~1,800	Video	Dynamic	✓	✗	✓	✓	✗	Seen
STI-Bench [30]	2,064	Video	Dynamic	✓	✗	✓	✓	✗	Seen
DSI-Bench [67]	~1,700	Video	Dynamic	✓	✗	Undeclared		✗	Seen
SCP-Bench (Ours)	2,500	Video	Dynamic	✓	✓	✓	✓	✓	Unseen

on a simplified subset of the benchmark that is single-view and requires predicting future states. The experiments are performed on Qwen3-VL-8B, InternVL3.5-8B, MiniCPM-V-4.5 and LLaVA-OneVision-1.5-8B in 3 modalities:

- **Text Scaffold:** Using GPT-5 [40], which performs best on SCP, we generate a detailed description of the future that is required to answer the question.
- **Image Scaffold:** Using Flux.1-dev-12B [28], we generate a predictive future frame directly based on the textual future description.
- **Video Scaffold:** Using Wan2.2-TI2V-5B [43], we attempt to generate a video approximation of the unseen portion by starting from the cutpoint frame and guiding the synthesis with the textual future description.

Fig. 13 shows an example using scaffolds. Each scaffold is provided as auxiliary information to support the model in answering the question. In Table 14, all four models show performance gains, with textual scaffolds yielding the largest improvement, followed by video scaffolds. Image scaffolds offer the weakest benefit. Textual scaffolds offer the most complete and explicit description of the future. In contrast, image and video scaffolds can only approximate the real scene to a limited extent because they do not capture the full textual detail. Video scaffolds, however, perform better than image scaffolds because they provide richer temporal and spatial cues.

Text Scaffold Prompt

Based on the provided video clips and the future description, please answer the following question. Choose the most appropriate option from the given

choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Future Description: {future_description}

Image Scaffold Prompt

Based on the provided video clips and a future image after the video, please answer the following question. Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

Video Scaffold Prompt

Based on the provided video clips and the generated future video clips, please answer the following question. Choose the most appropriate option from the given choices. (Only answer the choice text without any explanation)

Question: {question}

Options: {options}

E. Benchmark Comparison

As shown in Table 15, we provide a more detailed comparison between SCP-Bench and representative spatial reason-

Table 16. Results of the ‘‘Extended Causal Consistency’’ experiment, where models are evaluated by how often they misclassify the correct answer as the least plausible option. Lower values indicate stronger causal consistency. ‘‘Avg.’’ denotes the overall average error rate.

Model	Avg.↓	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Misclassifying the Correct Answer as Least Plausible</i>									
Qwen3-VL-8B	29.76	34.73	4.27	26.27	27.35	42.29	21.43	29.68	35.61
InternVL3.5-8B	30.96	32.34	8.55	26.27	32.49	38.83	33.33	28.39	33.45
MiniCPM-V-4.5	34.36	37.72	15.38	28.11	32.84	45.48	35.71	39.35	38.13
LLaVA-OneVision-1.5-8B	31.64	36.53	2.56	32.38	26.89	46.28	33.33	31.61	34.53

Table 17. Evaluation results for the ‘‘Physical Commonsense Probing’’ experiment. ‘‘Avg.’’ indicates the overall average accuracy.

Model	Avg.	Appearance Order	Counting	Planning	Relation	Relative Distance	Relative Size	Relative Speed	Spatial State
• <i>Vanilla CoT</i>									
Qwen3-VL-8B	51.52 ^{↑4.00}	63.47	50.43	51.53	44.39	54.52	90.48	61.29	51.80
InternVL3.5-8B	50.36 ^{↓0.16}	59.28	51.28	49.29	43.02	56.65	61.90	65.16	51.08
MiniCPM-V-4.5	44.88 ^{↑1.08}	55.09	40.17	48.68	35.47	53.19	73.81	59.35	39.93
LLaVA-OneVision-1.5-8B	42.88 ^{↓2.64}	48.50	39.32	46.44	33.30	53.99	69.05	54.19	39.57
• <i>Physics-aware CoT</i>									
Qwen3-VL-8B	50.80 ^{↑3.28}	61.08	52.14	50.51	45.88	57.71	73.81	58.06	43.17
InternVL3.5-8B	48.24 ^{↓2.28}	57.49	45.30	49.29	40.39	57.98	83.33	61.29	41.01
MiniCPM-V-4.5	42.52 ^{↓1.28}	56.29	43.59	46.64	34.67	46.81	66.67	52.90	35.97
LLaVA-OneVision-1.5-8B	44.04 ^{↓1.48}	47.90	44.44	44.60	36.84	50.53	80.95	56.13	42.09

ing benchmarks. Benchmarks such as OmniSpatial [27], Spatial457 [46], and EmbSpatial-Bench [16] are built around image-only settings and assess spatial reasoning in static scenes. Although they offer meaningful insights into how models interpret spatial structure, they fail to capture the evolving spatial relationships that arise in dynamic settings. All-Angles-Bench [57] and MMSI-Bench [56] extend single-view evaluation to multi-view settings, but they remain confined to image-based static scenes. Although presented in video form, VSI-Bench [55] effectively evaluates static indoor scenes rather than genuinely dynamic settings. VLM4D [69], STI-Bench [30], and DSI-Bench [67] use videos to introduce dynamic spatial reasoning, but their evaluation remains confined to interpreting the information explicitly present in the visible video frames. They disregard both the assessment of causal reasoning and the predictive analysis of unobserved portions of a spatial process, whereas spatial causal prediction addresses these aspects and better reflects the demands of everyday situations. MindCube [59] examines a model’s ability to hypothesize unseen spatial states using images and text, but it remains restricted to static scenes and lacks visual spatial dynamics.

In contrast, SCP-Bench focuses on spatial causal prediction in dynamic video settings, requiring models to infer spatial states beyond the visible part. In addition, SCP-Bench explicitly incorporates diverse perspective settings, covering both egocentric and exocentric viewpoints as well as single-view and multi-view configurations.

F. Supplementary Results

F.1. Extended Causal Consistency Experiment

To further examine models’ causal consistency in the SCP task, we design an experiment in which the model is re-

quired to select the least plausible option instead of the correct one. This setting evaluates whether the model can reliably rule out options that contradict the causal structure of the video, providing a complementary perspective on its causal reasoning stability.

As shown in Table 16, we evaluate Qwen3-VL-8B, InternVL3.5-8B, MiniCPM-V-4.5, and LLaVA-OV-1.5-8B by measuring how often they mistakenly classify the correct answer as the least plausible option, where a lower rate indicates stronger causal consistency. Although InternVL3.5-8B achieves better performance in the baseline setting, Qwen3-VL-8B produces fewer such errors in this consistency test. As illustrated in Fig. 14, Qwen3-VL-8B evaluates each option’s plausibility more independently, whereas InternVL3.5-8B often commits to a presumed correct answer and then selects the option most opposed to it. When its initial guess is incorrect, this strategy makes it more likely to mislabel the true answer as the least plausible.

F.2. Physical Commonsense Probing Experiment

To examine the role of physical commonsense in SCP performance, we introduce a physics-aware CoT prompting strategy. The prompt instructs the model to first enumerate the relevant physical principles and then apply them to answer the question. However, as shown in Table 17, physics-aware CoT does not clearly improve performance compared to vanilla CoT. Analyzing the model’s reasoning traces in Fig. 15 reveals a consistent pattern: although the model can correctly identify the appropriate physical laws, it often fails to apply them to the scene. This indicates that current models lack the ability to operationalize physical principles in context and to integrate them into reliable causal reasoning.

Q: From the camera's perspective, right after the surfer plants the paddle on the inside of the turn and the rail begins to bite near the lower-left of the frame, toward which screen side does the board's nose end up pointing as the carve completes near the center-right?
True Answer: Right

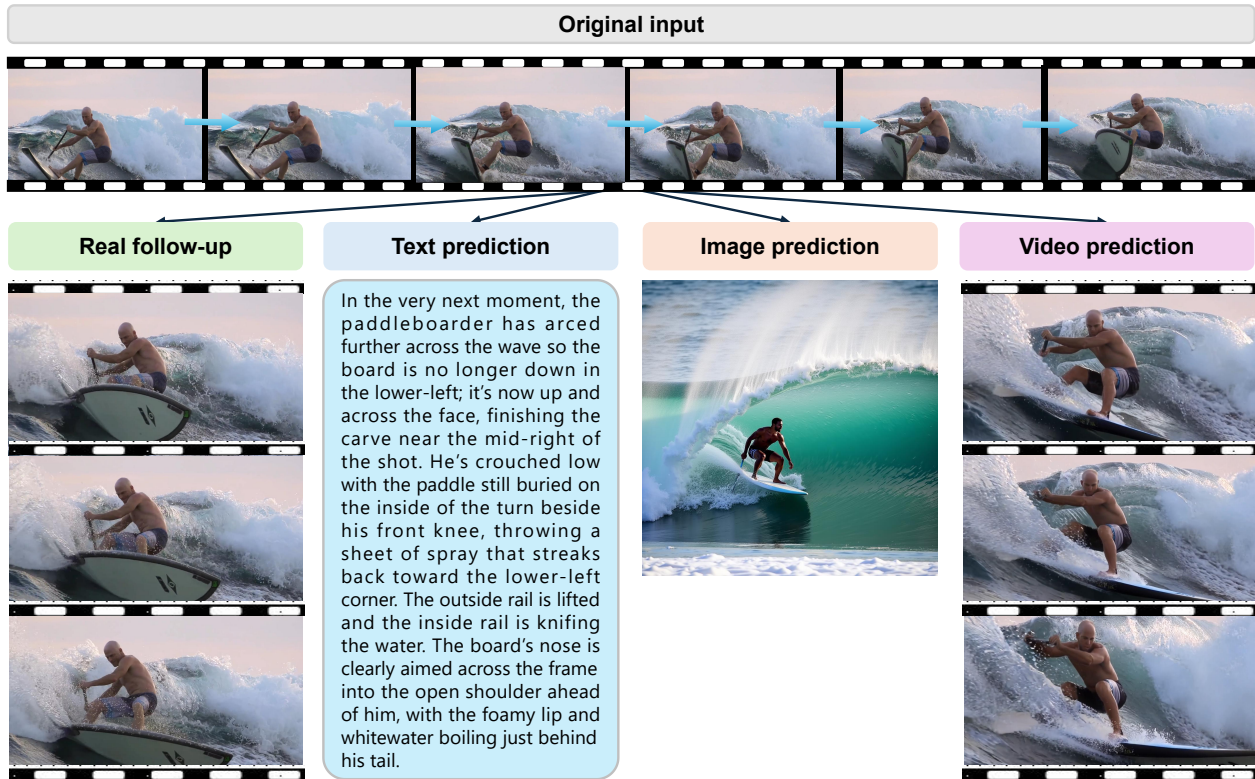


Figure 13. Example of causal prediction enhancement. Here we present a comparative analysis of a case study, showing the actual subsequent video and the causal prediction results for text, image, and video modalities.

F.3. Error Case Analysis

To further understand why models struggle on the SCP task, we analyze representative error cases and identify several underlying causes of failure:

- **Dynamic Integration Failure.** A common failure cause is the model's inability to integrate motion cues over time, leading it to rely on isolated frames or short local motion fragments instead of coherent temporal dynamics. As illustrated in Fig. 16, the first case considers only the riders' initial positions while ignoring the subsequent motion changes. The second case focuses solely on the individuals' current distances, overlooking how their relative positions evolve as they move. In the first case in Fig. 17, the model captures only the boat's local turning motion but fails to recognize its global dynamic trajectory.
- **Prior-Driven Hallucination.** The model fails to attend to detailed visual cues and instead leans too heavily on its prior assumptions. As shown in the second case of Fig. 17, although it correctly identifies the main evolving motion, it subsequently jumps to a conclusion

driven by prior assumptions, ignoring the actual scene details and producing an incorrect answer.

- **Causal Reasoning Failure.** The model fails to produce grounded causal reasoning and instead jumps to premature, unsupported conclusions. In the first case of Fig. 18, for example, it asserts that the blue car is moving faster without any justification, overlooking the causal reasoning needed to compare how the distance between cars evolves over time, which is a process humans naturally perform.
- **Cross-Modal Attribution Bias.** The model fails to integrate visual and textual modalities into a coherent line of reasoning, often overcommitting to one source and producing biased conclusions. As shown in the second case of Fig. 18, it becomes fixated on interpreting the textual mention of the player placing the ball and derives its answer from that alone, while entirely ignoring the visual evidence of the ball's actual trajectory that should have informed the reasoning.



Q: After the worker finishes several cuts within the dense cluster and begins turning away with the shears lowered, in which direction will they head next from the camera's viewpoint: toward the open left foreground, deeper into the back rows beyond the pipe, or further right?
True Answer: Toward the open left foreground.

InternVL3.5-8B:
 ... Toward the open left foreground: **The worker has been moving to the right, and there is no indication of a turn toward the left.** ...
 Most Inappropriate Option: **Toward the open left foreground** *Mistake!*

Qwen3-VL-8B :
 ... **The back rows beyond the pipe are further away and not the immediate direction of movement after finishing a section.** ...
 Most Inappropriate Option: **Go toward the back rows beyond the pipe**

Figure 14. Analysis sample for extended causal consistency experiment. The models are required to identify the least plausible option, while InternVL3.5-8B does not independently assess the plausibility of each choice but instead derives the least plausible option by reasoning backward from its own incorrect answer, revealing a lack of internal causal consistency.

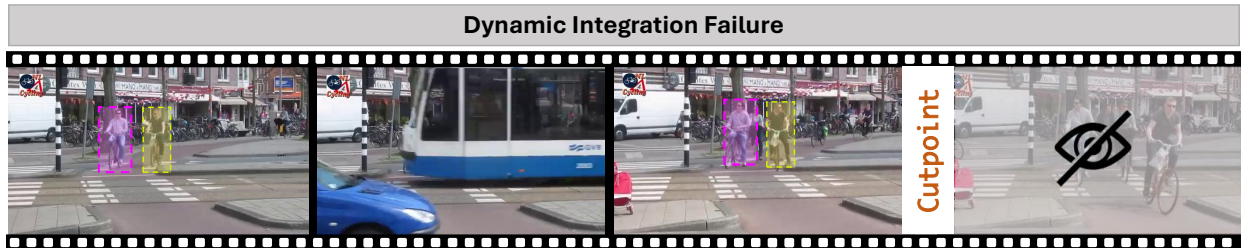


Q: From the camera positioned behind the goal, as the blue kicker's right leg swings through and the ball first leaves his foot, into which part of the goal mouth will the ball ultimately enter: left, center, or right?
True Answer: Left.

LLaVA-OneVision-1.5:
Physics Principles Identification: Line-of-sight, Perspective projection, Relative motion ... The perspective projection confirms that the ball is on a trajectory that will **pass through the right section of the goal mouth.**
 Final Answer: **Right** *Wrong!*

Minicpm-8B:
Physics Principles Identification: Line-of-sight, Perspective projection ... **Without explicit curve indication but considering direct projection, the answer is right.**
 Final Answer: **Right** *Wrong!*

Figure 15. Visualization of the physical commonsense probing experiment. To answer each question, models must first infer the underlying physical principles depicted in the video. Although current models can often recognize the appropriate physical laws, they still struggle to apply them effectively when generating correct predictions.



Q: Relative to each other, when movement resumes, who reaches the near curb island first: the dark-shirt rider on the camera-right or the gray-hoodie rider on the camera-left?

True Answer: The dark-shirt rider on the camera-right

Human explanation:

As seen in the previous footage, when the car stopped at the intersection, the dark-shirt rider on the camera-right was positioned further forward than the gray-hoodie rider on the camera-left, suggesting that the dark-shirt rider on the camera-right reaches the near curb island first.

Qwen3-VL-8B:

...The two riders both start pedaling. The gray-hoodie rider is slightly ahead of the dark-shirt rider as they both continue to move forward. ... Gray-hoodie still leads, will hit the curb island first.

Final Answer: The gray-hoodie rider on the left.

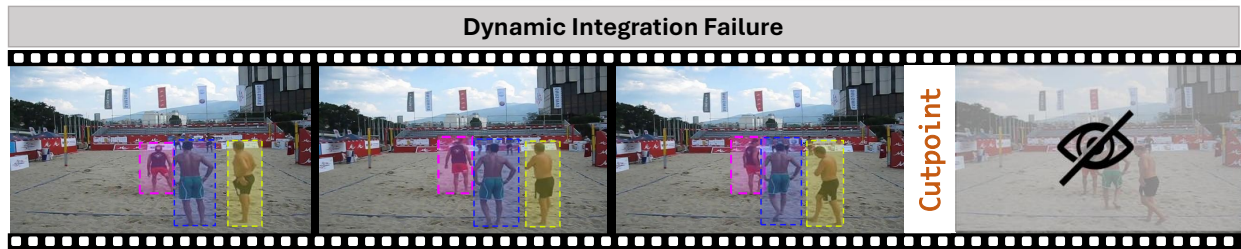
Wrong!

LLaVA-One-Vision-70B:

The dark-shirt rider on the camera-right... The gray-hoodie rider is closer to the curb island. Since the gray-hoodie rider has the initial positional lead, this slight advantage will be maintained over the short distance.

Final Answer: The gray-hoodie rider on the left.

Wrong!



Q: From the camera's perspective, starting when the three foreground men first step off the near baseline and head toward the net, who increases their distance from the baseline the fastest?

True Answer: The man in green shorts

Human explanation:

In the initial frames, the man in the black shirt is observed to be positioned farthest forward but shows minimal forward progression. Conversely, the man in the black shorts demonstrates rapid lateral movement. However, the man in green shorts subsequently maintains sustained acceleration and forward strides, ultimately achieving the fastest increase in distance from the baseline.

GPT-5:

The man in black shirt is the first to step off the baseline and strides forward quickly, creating the largest and fastest growing gap from the baseline compared to the man in green shorts and the man in black shorts, who move more slowly.

Final Answer: The man in black shirt

Wrong!

Gemini-2.5-Pro:

...If you observe their movement closely, the man in the black shirt on the left is moving at a noticeably quicker pace, almost a light jog, compared to the other two men who are walking more slowly. Because he is moving faster, he increases his distance from the baseline the fastest.

Final Answer: The man in black shirt

Wrong!

Figure 16. Cases of dynamic integration failure. The models fail to perceive dynamic changes and remain confined to static observations.

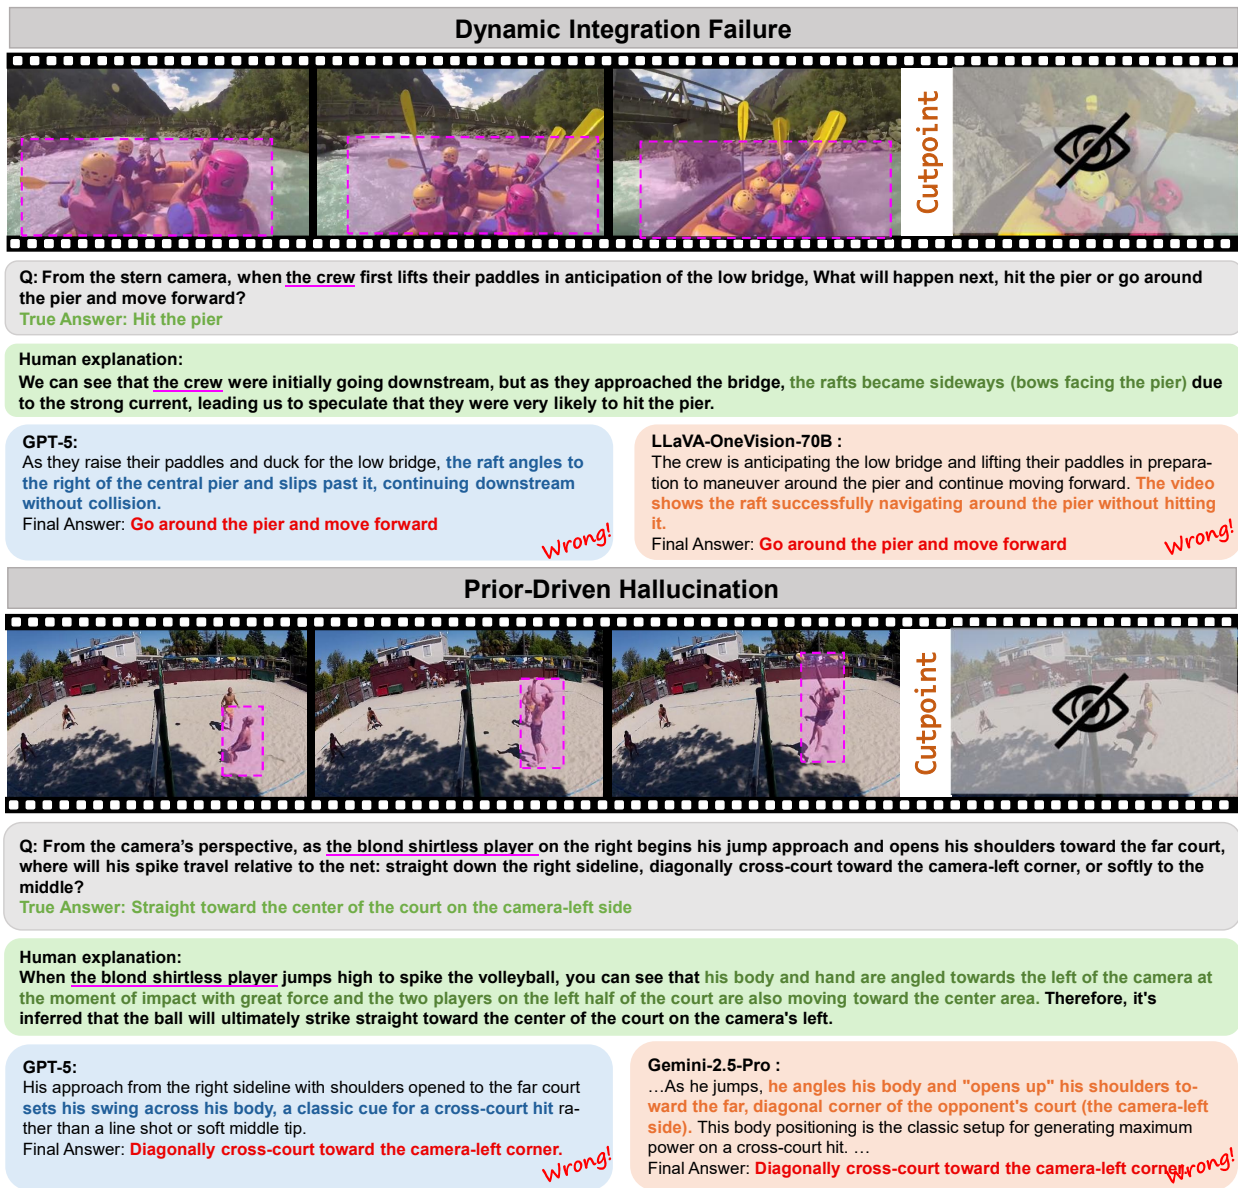


Figure 17. Cases of dynamic integration failure and prior-driven hallucination. The models struggle to perceive and infer the correct global spatial dynamics and often overlook critical scene details during reasoning due to prior-driven biases.

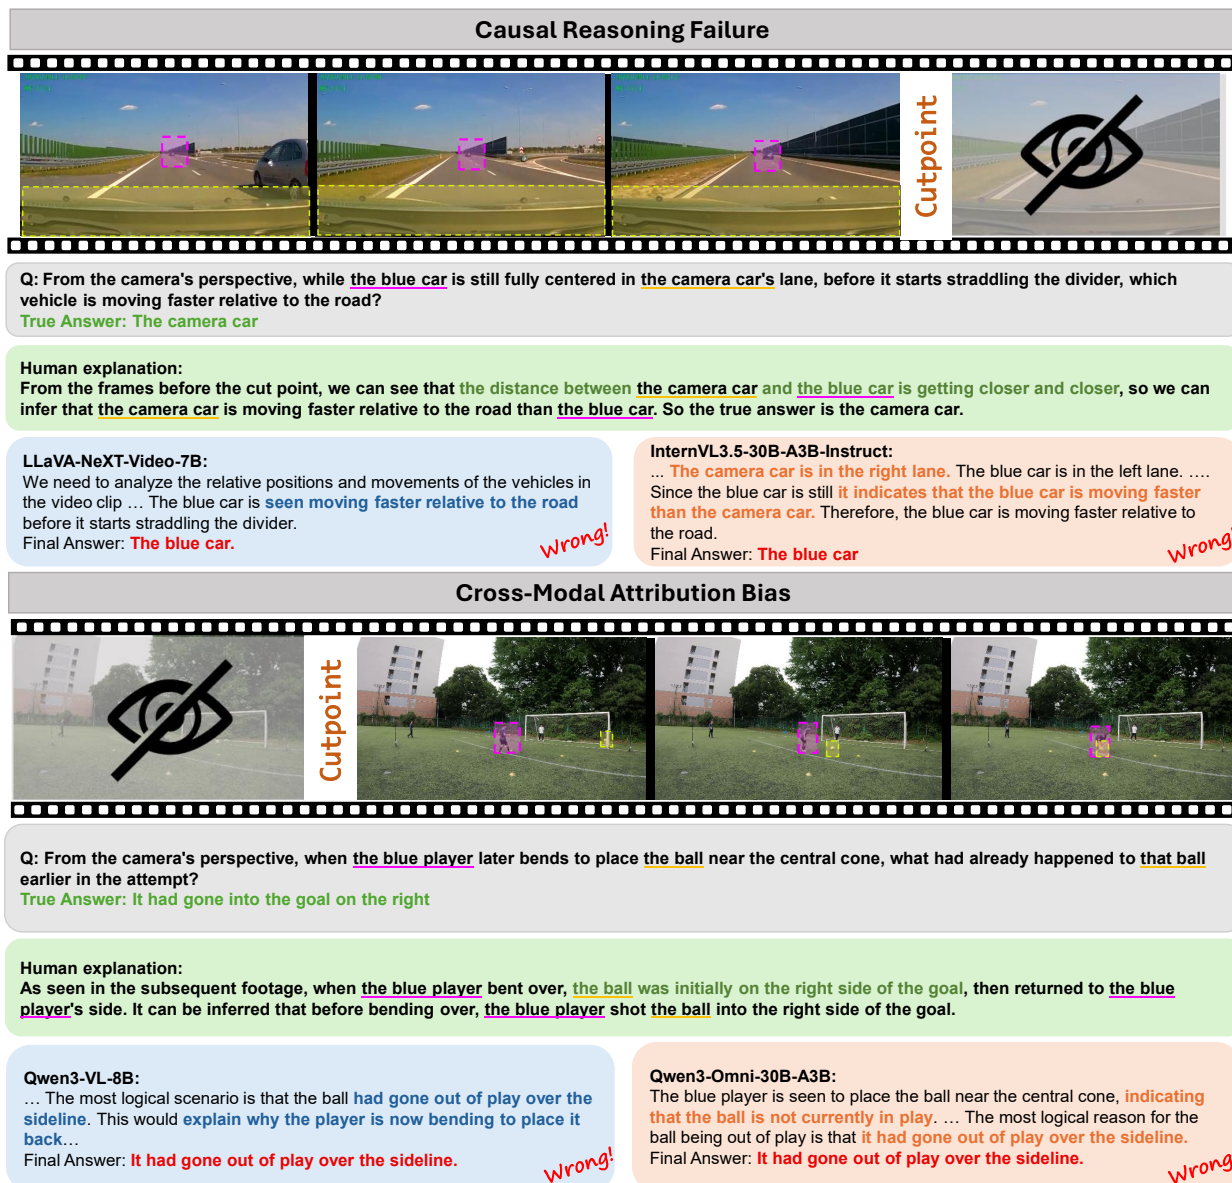


Figure 18. Cases of causal reasoning failure and cross-modal attribution bias. The models struggle to perform causal reasoning and instead produce unsupported conclusions, leading to erroneous predictions. Moreover, they may overlook information from other modalities and fall into single-modality reasoning.

Video source dir: v_pIE3KNmuwj4

2 files

Refresh Collapse file list Check QA count

v_pIE3KNmuwj4.mp4

Segment List

All Pending Reviewed Unavailable

Use the dropdown to change segment status

v_pIE3KNmuwj4_seg01
 Status: Reviewed
 Video: v_pIE3KNmuwj4
 QA Count: 8
 Last Modified: 2025-10-05 16:34:24

v_pIE3KNmuwj4_seg02
 Status: Unavailable
 Video: v_pIE3KNmuwj4
 QA Count: 8
 Last Modified: 2025-10-06 12:11:01

v_pIE3KNmuwj4_seg03
 Status: Reviewed
 Video: v_pIE3KNmuwj4
 QA Count: 8
 Last Modified: 2025-10-06 12:15:41

v_pIE3KNmuwj4_seg04
 Status: Reviewed
 Video: v_pIE3KNmuwj4
 QA Count: 7
 Last Modified: 2025-10-06 12:18:17

QA List Current: v_pIE3KNmuwj4_seg01

From the boat's coordinate system (bow forward), as the boat first begins a leftward arc to round the small palm islet on the right side of the frame, where will the wakeboarder be relative to the boat when the turn develops: port side (inside), directly astern, or starboard side (outside)?

Answer: Starboard side (outside)

Type: Planning | Direction: Forward (model sees first half) | Time: 00:45.54 - 00:52.63

Cut point: Current Time

View: v_pIE3KNmuwj4.mp4 + Select current view

Reason: In the sub-segment the boat initiates a left turn while the rope is angled to the right; the rider's momentum swings him outward to the boat's starboard side during the turn.

Edit Delete Play Select Mark Usable

From the shoreline view, when the helicopter lifts from near the grass and starts closing on the boat along the bank, which will be moving faster relative to the water and end up ahead moments later: the helicopter or the boat?

Answer: The helicopter

Type: Relative Speed | Direction: Forward (model sees first half) | Time: 00:45.54 - 00:52.63

Cut point: Current Time

View: v_pIE3KNmuwj4.mp4 + Select current view

Reason: The sub-segment shows the helicopter rapidly reducing the gap and sliding forward along the shore-parallel line while the boat's speed appears steady, implying the helicopter will overtake.

Edit Delete Play Select Mark Usable

From the aerial camera, at the instant the rider is carving wide with his trailing hand touching the water on the right half of the lake, what will he do next: continue farther outward, release the handle, or cut back toward the boat's wakes?


Answer: Cut back toward the boat's wakes

Type: Planning | Direction: Forward (model sees first half) | Time: 00:45.54 - 00:52.63

Cut point: Current Time

Video Player

Video: v_pIE3KNmuwj4
 Type: multi
 Current View: v_pIE3KNmuwj4.mp4
 Available Views: v_pIE3KNmuwj4.mp4



0:05 / 2:37

00:05.63 / 02:37.62 Play selected QA time Play first half Play second half

Shortcuts: Space Play/Pause | Enter Play selected QA time

Figure 19. Manual filtering tool demonstration. Annotators use it to filter appropriate QA candidates and determine their cutpoints.

✔ Usable: 11
 ✘ Unusable: 0
 🟢 Usable & Answered: 0
 🟡 Usable & Unanswered: 11
 Progress: 0.0%

Load Quiz File

Question 1
● Unanswered

Question 2
● Unanswered

Question 3
● Unanswered

Question 4
● Unanswered

Question 5
● Unanswered

Question 6
● Unanswered


Question 7
● Unanswered

Question 8
● Unanswered

Question 9
● Unanswered

Question 10
● Unanswered

Question 11
● Unanswered



0:00 / 7:45

Current: 00:00.00 | Segment: 01:25.00 - 01:36.00 | Cut: 01:29.57

Perspective: youtube_008.mp4 | youtube_008.mp4

▶ Play First Half

From the camera's view when the zebra crosswalk is fully visible, which appears larger: the white van at the right curb or the white truck ahead in the lane?

Relative Size | Forward | 01:25:00 - 01:30:00

The white van

The white truck ahead in the lane

Invalid Reasons Guide

- ❌ **Wrong Question Type:** Incorrect question classification, e.g., marking a counting question as a reasoning question
- ❌ **Unclear Question Perspective:** Unclear perspective selection, cannot determine observation angle
- ❌ **Question Itself Has Issues:** Abrupt time points, unclear question target
- ❌ **Inappropriate Time Interval:** Cut point needs fine-tuning, etc.
- ❌ **Inappropriate Options:** Poor quality distractors (obviously fake, multiple correct answers duplicated), incorrect original answer, etc.
- ❌ **Cannot Answer, Completely Discard:** Completely unable to answer this question for various reasons

👁 Show Answer

🚫 Mark as Unusable

Figure 20. Validation tool demonstration. Annotators use it to validate each item in strict accordance with the SCP task specification.

Total QAs: 46
v2 Version: 1

Load JSON File

P02-01 15 QA

- QA 1 v1
- QA 2 v1
- QA 3 v1
- QA 4 v1
- QA 5 v1
- QA 6 v1
- QA 7 v1
- QA 8 v1
- QA 9 v1
- QA 10 v1
- QA 11 v1
- QA 12 v1
- QA 13 v1
- QA 14 v1
- QA 15 v1

+ Add New QA

v_21Pz1qddZl 4 QA

+ Add New QA

v_6YyqM9-XgAk 1 QA

+ Add New QA

v_AH4c5vqgUic 1 QA

+ Add New QA

v_JDn5TW9WwM 3 QA

+ Add New QA

v_JTrwGfPjNtU 1 QA

+ Add New QA

v_K3Z3z8t-RIQ 1 QA

+ Add New QA

v_NjTkn2nalaac 1 QA

+ Add New QA

v_OqLUp37WQMA 2 QA

+ Add New QA

v_VlLq4bAHCOI 1 QA

+ Add New QA

v_ZcgahXg_ELw 2 QA

+ Add New QA

v_Zda5-WZHJZY 1 QA

+ Add New QA

v_7XW-BFK_ZY 1 QA

+ Add New QA

v_gbuRv8phs1Y 1 QA

+ Add New QA

v_J7vUMNM84Yo 1 QA

+ Add New QA

v_Jen7R78v5NY 1 QA

+ Add New QA

v_pQa5XFTLoY 1 QA

+ Add New QA

v_sG3jPmUXFvU 1 QA


+ Add New QA

youtube_069 6 QA

+ Add New QA

youtube_168 1 QA

+ Add New QA



2:17 / 29:47

Current: 02:17.37

Set as cut_point

Set as Question Time

Play First Half

Play Second Half

Answer

Ground Truth *

Left

Question Validity

Is Question Valid

Invalid (false)

Invalid Reason *

Please Select Invalid Reason

Please select the specific reason why the question is invalid

All Perspectives

Main Perspective

Question Perspective

Basic Information

QA ID

segment_1756344897675_ga_4

Video Name

P02-01

Question *

From the camera's perspective during the transition when the pan leaves the dish rack and is carried to the stove, in which horizontal direction across the screen does the pan predominantly move?

Question Type *

Relation

Time Direction *

forward

Start Time (MM-SS-XX)

07:44.00

End Time (MM-SS-XX)

07:47.00

Cut Point (MM-SS-XX)

07:46.38

Question Perspective Time Point (MM-SS-XX)

e.g.: 01:34.50

Options

Down

Right

Diagonal

Up

Left

+ Add Option

Copy This QA

Delete This QA

Figure 21. Repairing tool demonstration. Annotators use it to repair and optimize any attribute of the item as needed.