

# FLAT-PACK BENCH: Evaluating Spatio-Temporal Understanding in Large Vision-Language Models through Furniture Assembly

## Supplementary Material

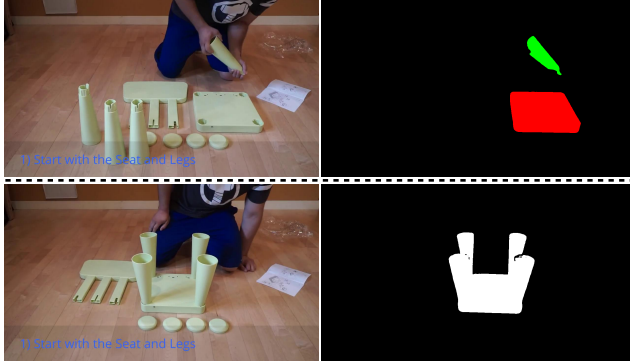


Figure S1. **Incomplete Segmentations in IMAW [15].** (Top) Segmentations are only provided for the parts that are about to be connected, i.e., the seat panel and the leg. Parts not being interacted with (the other legs and the backrest) are not annotated. (Bottom) Furthermore, instead of granular segmentations of each individual part, IMAW only has masks at the sub-assembly level.

## A. Benchmark Details

### A.1. Additional Benchmark Statistics

The dataset contains videos of 24 unique furniture items, consisting of 3-19 parts (average of 7 parts). Keyframe videos are pre-extracted at 1 FPS, whereas trimmed videos retain their original variable frame rates, averaging 28.98 FPS. The average duration of keyframe videos is about 6 minutes, compared with 2.98 minutes for trimmed videos. Based on keyframe videos and connection-event annotations derived from IMAW[15], the minimum number of frames required to answer a question is 113.7 on average (median: 69.5). Similarly, for TRACK questions, the average interval between the visual prompts is 141.1 frames (median: 76).

### A.2. Incomplete Segmentations in IMAW

In Sec. 3, we mentioned that IMAW [15] has incomplete segmentation annotations. Here, we elaborate this statement:

1. We found that segmentation annotations are only provided for parts that are in the process of being connected in a particular key frame. This limits the questions to these particular key frames, and even so, to only the parts annotated therein.
2. Furthermore, IMAW only includes information at the sub-assembly granularity, which precludes questions that one might want to ask about specific parts.

Figure S1 shows some examples of available segmentations in IMAW to illustrate these issues. Thus, we annotate our own segmentation maps as described in Sec. 3.

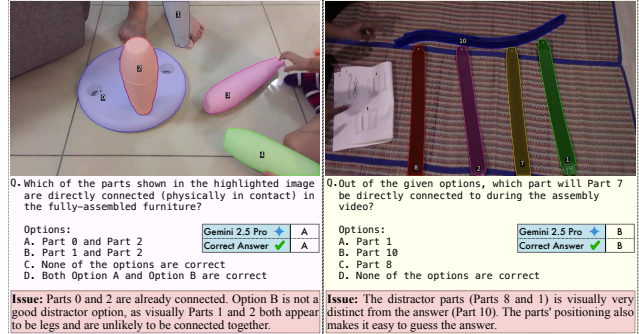


Figure S2. **Issues with Auto-generated Questions.** We show some examples to highlight the issue of auto-generated questions where the answer can be arrived at using the visual prompt alone and commonsense reasoning. Observe that we can arrive at the correct answer without requiring any temporal cues from the video.

## A.3. Manual Question Curation

**Pitfalls of Auto-generated Questions.** Figure S2 shows some examples of auto-generated questions, where the question can be solved using the visual prompt alone. In the first example (left), Part 0 is already connected to Part 2 in the prompt. Furthermore, Parts 1 and 2 are both clearly legs and unlikely to be connected given the structure of Part 0. Overall, this frame is unsuitable for a MATE question. Similarly, in the second example, the distractor parts (Parts 1 and 8) are visually quite distinct from Part 10. Their positioning, as if they are about to be inserted into Part 10 (not Part 7), also makes it easy to answer this question without any temporal context. Observe that even without the videos, the correct answer is easy to infer from the visual prompt and commonsense reasoning alone. Such examples motivated us to manually curate our questions.

**Question Templates.** Table S1 shows the full list of question templates used by our annotators. Annotators had to populate the parts highlighted in red with part IDs and provide options that they believe are challenging (parts with similar appearance, parts that are connected much later in the assembly compared to the state shown in the visual prompt, etc.) All templates require up to 4 options, except for the second template in MATE, which is a binary question.

## B. Evaluation Details

Here we provide additional experimental details for Sec. 4, including video processing and visual prompt construction.

**Video Data Processing.** All videos are resized to a height

Task	Question Template
MATE	<ul style="list-style-type: none"> <li>• Out of the given options, which part will <i>{query-part}</i> be directly connected to during the assembly video?</li> <li>• Are <i>{query-part1}</i> and <i>{query-part2}</i> connected in the fully-assembled furniture?</li> <li>• Which of the parts shown in the highlighted image are directly connected (physically in contact) in the fully-assembled furniture?</li> </ul>
TLOC	<ul style="list-style-type: none"> <li>• Which of the highlighted parts are connected (brought in direct physical contact) to <i>{query-part}</i> next starting from the assembly state shown in the image?</li> <li>• Which highlighted part is the last to connect (brought in direct physical contact) to <i>{query-part}</i> in the video?</li> <li>• Which highlighted pair of parts is connected first in the video?</li> <li>• Which highlighted pair of parts is connected last in the video?</li> <li>• Which labeled part was connected last to reach the current assembly state shown in the image?</li> <li>• Which labeled part will be connected next to the current assembly state shown in the image?</li> </ul>
TORD	<ul style="list-style-type: none"> <li>• Which of the following options lists the highlighted parts in the exact order that they become directly attached (physically in contact) to <i>{query-part}</i> in the video, from the first attachment event to the last?</li> <li>• Each option below lists a sequence of part connections. Which of these options gets the temporal ordering of the connections correct?</li> </ul>
TRACK	<ul style="list-style-type: none"> <li>• Identify the correct set of matches from Image A to Image B.</li> <li>• Identify which part in Image B matches Part <i>{query-part}</i> in Image A.</li> </ul>

Table S1. **Question Templates.** We replace the **highlighted** part in the question template from scene to scene to construct our benchmark.

of 480 px and a width of 640 px. We dump all videos at 1 FPS before passing them as input to the model.

**Video Subsampling.** During evaluation, we first decode each video at its stored frame rate and expose all frames in that stored video to the model pipeline, so that no connection event is discarded a priori. Thus, keyframe videos provide all 1 FPS keyframes, while trimmed videos provide all frames at their original variable frame rates. Any subsequent frame selection is then performed by the individual model pipeline, subject to its context-length constraints, using a common 1 FPS temporal view for both video types. The videos in our benchmark are quite long (average duration of 6 minutes for key-frame videos at 1 FPS, i.e., 360 frames on average). Typically, it is not possible to pass all the frames as input to models due to their limited context length. Hence, we uniformly subsample the frames before passing them as input. Subsampling frames uniformly from a fixed frame rate of 1 FPS is similar to the FPS-sampling strategy advocated by recent works [16, 49]. For the number of frames to subsample, we followed the official guides/cookbooks released by the models wherever possible. Below we provide some details about subsampling choices:

- **Proprietary Models:** For Gemini [28], we provide the entire video as input. For trimmed videos that exceed the duration limit of 1 hour for Gemini, we dump the videos at the minimum FPS where the duration limit is satisfied. For GPT-5 [19], we provide 499 frames (498 in case of TRACK questions), as per the API token limit, which accommodates 500 frames.
- **Open Models:** For the Qwen family of models [1, 22] (including ArrowRL [35]), we sample a maximum of 768 frames, while maintaining a context length of 20480 tokens, as suggested by their official cookbook [22]. For all other open models, we subsampled 32 frames, except Video-LLaVA [13] (8 frames), as per the setting used

within their respective training and evaluation pipelines.

Note that, for all models, during subsampling of *Concat* prompts, we modify the subsampling procedure so that the first two frames (containing the visual prompts) are retained, while maintaining the total number of frames.

**Visual Prompt Construction.** As mentioned in Sec. 4, we use three different settings for constructing the visual prompt – *Mixed-Media*, *Collage*, and *Concat*. Examples of the complete visual inputs for different prompt settings can be found on the project site<sup>3</sup>

The images that constitute the visual prompts are also resized to a resolution of 480 × 640. For TRACK questions, the jumbled image prompt (Image B in Fig. 2) is provided as the *second* image in the input for *Mixed-Media*, as the middle image in each frame for *Collage*, and as the second frame of the video for *Concat*. These details are also conveyed to the model in the task instructions provided along with the input.

**Evaluation Setup for Specialized Models.** We set up the evaluation for specialized models as follows:

- **ArrowRL** [35]: As discussed, for ArrowRL [35], we use a similar evaluation setup to Qwen 2.5-VL, as that is the base model from which it was created.
- **PerceptionLM** [5]: For PerceptionLM [5], we subsample 32 frames from the video, similar to the evaluation settings described in the paper. We evaluate PerceptionLM only on *Collage* and *Concat* prompts because it lacks support for *Mixed-Media* prompts.
- **VideoRefer** [39]: For evaluating VideoRefer [39], instead of using our prompting strategies, we follow its inference recipe by directly giving the part segmentations and the frame indices for the visual prompts as input to the model.
- **GenS** [38]: For GenS [38], we found that the original model can only handle up to 256 frames. Hence, to eval-

<sup>3</sup>[flat-pack-bench.github.io/visual-prompts/](https://flat-pack-bench.github.io/visual-prompts/)

uate longer videos, we split the video into chunks of 256 frames and obtained a relevance score for frames in each chunk. We filtered the frames in each chunk using a score threshold of 4 and provided the remaining frames as input to the base model (Gemini 2.5 Pro).

**Task Instructions.** In addition to the video, the visual prompts, and the question, the input contains task instructions that describe the input visual data to the model and specify the output format that the model should follow. Figures S3 to S5 show the task instructions that we use in our evaluations. When asking TRACK questions, we make adjustments to the input descriptions, informing the model that there are two visual prompts instead of one. Note that the task instructions are not a part of the benchmark. We welcome users of the benchmark to try out different phrasings of the instructions. For reproducibility, we release the source code used in our evaluations at: [github.com/justachetan/flat-pack-bench](https://github.com/justachetan/flat-pack-bench).

**Final Prompt.** The final prompt used in our evaluations has the following structure: [Video Frames] [Visual Prompt] [Jumbled Visual Prompt] [Task Instructions] [Question]. The [Jumbled Visual Prompt] is exclusive to TRACK questions.

**Human Evaluation.** Figure S6 shows the instructions for the human participants who attempted our benchmark both the standard task and the image-only task (Sec. 5.3). In the image-only task, we asked the participants to select an answer even if they were unsure.

In addition to the in-house human evaluation, we also conducted a broader human evaluation study on Prolific [21] with a smaller set of 186 questions. Each participant was given a video and 4-5 questions related to the video. We collected at least 3 responses (up to 5) per question and computed the final answer using majority voting. Figure S7 shows the UI of the annotation tool that was shown to the participants. We used a similar tool for our in-house annotations. Results are shown in App. C. We provide the instructions page that the Prolific participants received before beginning annotations on our project website: [flat-pack-bench.github.io/assets/prolific-instructions/](https://flat-pack-bench.github.io/assets/prolific-instructions/).

## C. Additional Results

**Human Performance.** For the human performance study (Sec. 4), we recorded a micro-average unanimous response rate of 80%, indicating that our questions are clear and consistently understood. Task-wise, TORD attains the highest agreement (88%), followed by MATE (86%), and TRACK (77%), and lastly TLOC (70%). Errors largely stemmed from confusing similar-looking parts over longer videos.

For the image-only human-performance results (Tab. 4), we observed a micro-average agreement rate of 25%. Task-

Table S2. **Human evaluation summary on FLAT-PACK BENCH.** We show the performance of Prolific and our in-house participants, along with some reference models on the same subset of questions. Agreement is measured by unanimous response rate.

Cohort / Model	Micro Avg.	TORD	TLOC	TRACK	MATE
<i>Human Evaluators</i>					
Prolific Participants	72.58	75.51	88.88	65.47	73.07
Agreement	27.95	28.57	37.04	21.43	38.46
In-house Participants	98.92	100.00	100.00	97.62	100.00
Agreement	84.94	91.83	70.37	85.71	84.61
<i>Reference Models</i>					
InternVL3-78B	44.09	44.90	51.85	41.67	42.31
Qwen2.5-VL-72B	45.70	46.94	29.63	48.81	50.00

wise breakdowns for image-only variant were: 17% for TORD, 36% for TLOC, 27% for TRACK, and 20% for MATE.

The results of the broader human evaluation study discussed in Sec. B are given in Tab. S2. We also compare these results with the performance of in-house participants on the same subset of questions. As we can see, Prolific participants have a low agreement rate, but their performance is still significantly better than the reference models. This shows that while the questions are generally solvable with high accuracy by humans, they certainly require high cognitive effort.

**Zero-shot Results.** In Sec. 4 we stated that we evaluated all models on every combination of video type (key-frame, trimmed) and visual prompt type (Mixed-Media, Concat, and Collage). Tab. 2 showed the results for the best settings for each model. We provide the complete results in Tab. S6. An interactive version of the table is also available on our project website<sup>4</sup>. To obtain a more robust estimate of performance, we use *bootstrapping* to compute the 95% confidence interval (by sampling 50 videos with replacement for 100k trials). If a setting is not shown, it was not evaluated due to cost constraints (e.g., GPT-5 [19] was assessed only on mixed-media prompts for key-frame videos) or because it does not support a prompt format (e.g., Mixed-Media prompts were not supported by some models, such as LLaVA-Video [42]).

**Correlation with Video Difficulty and Duration.** To test whether model performance correlates with human-perceived difficulty, we manually annotated the videos on factors like the number of cut shots, the motion complexity of parts, etc. Figure S8 shows the instructions given to the manual annotators. We collected two annotations per-video and averaged the scores to obtain a difficulty score for each video. However, we found that for a standard open LVLM (Qwen2.5-VL 72B) human-perceived difficulty scores do not predict model performance (Spearman’s  $\rho=-0.20$ ,  $p=0.17$ ). Model performance also showed little correlation with video duration Spearman’s  $\rho=-0.21$ ,  $p=0.13$ ).

<sup>4</sup>[flat-pack-bench.github.io/results/](https://flat-pack-bench.github.io/results/)

## Task Instructions for Mixed-Media Prompts

You are a furniture-assembly expert. You are given:

1. A video of a furniture assembly in progress.
2. A single labeled frame from the same video, displaying numeric IDs on each visible furniture part. Call this a visual prompt.

### For TRACK Questions

2. Two labeled frames from the same video, called Image A and Image B, each displaying numeric IDs on visible furniture parts. Note: The numeric IDs in Image A and Image B are not necessarily the same; the same ID may refer to different parts in each image.

3. A multiple-choice question (with its list of answer options) that refers to both the video and the image.

Some additional assumptions to keep in mind:

- Two furniture parts are "connected" if they are directly attached (physically in contact), like they would be in the final assembly.
- Parts simply touching each other physically, but not in their final assembly position, are not considered "connected".
- While deciding the order of "connection" events between two parts, consider the order in which they were connected last as their actual connection, i.e., if two parts are connected, then disconnected and then connected again, the last connection is the one that matters.
- In the course of the video, some parts may be attached physically to each other before they are fixed in their final assembly position by screwing, gluing, or otherwise securing them. In such cases, the parts are considered "connected" from the moment they are physically attached, not when they are fixed. Hence, when answering questions about the order of connections, consider the physical attachment as the connection event, even if the parts are not yet secured to each other.

Your responsibilities:

- Watch and interpret the assembly steps in the video.
- Examine the labeled image frame to understand part relationships.
- Read the question and all answer choices.
- Determine the correct option.
- Respond **\*\*only\*\*** with a JSON object containing a single key, "answer", whose value is the letter of your chosen option (e.g., "A", "B", "C").

**\*\*No additional text, explanations, or formatting -- just the JSON string.\*\***

Now answer the following question:

Figure S3. **Task Instructions for the LVLMs.** We provide the models with task instructions that describe the input format, and additional information such as how we define `connection` events. When describing the inputs, we make adjustments for TRACK questions by mentioning that there are two visual prompts instead of one. The task instructions can be further broken down into 3 sections, describing the **Inputs**, **Assumptions**, and **Instructions** (colored accordingly above).

## D. Analysis Details

### D.1. Self-Explanations

We have provided some additional examples of self-explanations, expanding upon the results discussed in Sec. 5.4 on our project website. We also provide the full breakdown of error categories for the generated rationales discussed in Sec. 5.4 in Fig. S9.

### D.2. Experimental Details for TVA

We discussed an agentic baseline TVA in Sec. 5.5 and reasons for its poor performance. In this section, we provide the experimental details used to generate our results.

**Contact-Reasoning Issues.** For our contact-reasoning experiment, we use the task instructions from the image-only prompting experiments in App. D.4. We used two versions of the question template:

1. *In this task, "connected" means the two parts are in direct physical contact, in the same way they will be when the furniture is fully assembled (not merely near each other or partially aligned). Based on the given image, are {query\_part1} and {query\_part2} connected?*
2. *Are {query\_part1} and {query\_part2} connected (physically in contact) in the shown image?*

Table S3 shows the results. We observed poor performance across all settings, indicating that LVLMs struggle to understand even simpler concepts like physical contact.

### Task Instructions for Collage Prompts

- You are a furniture-assembly expert. You are given a video input where each frame consists of two images side-by-side:
- Right side: A frame from the furniture assembly video showing the assembly process.
  - Left side: A single labeled frame from the video on the right side, fixed for the entire video, displaying numeric IDs on each visible furniture part.

The left-side labeled frame remains constant throughout the video.

You will be asked multiple-choice questions referring to the assembly video and the labeled parts shown on the left.

#### For TRACK Questions

You are a furniture-assembly expert. You are given a video input where each frame consists of three images side-by-side:

1. Right side: A frame from the furniture assembly video showing the assembly process.
2. Left side: A labeled frame from the video on the right side, fixed for the entire video, displaying bright numeric IDs on each visible furniture part. Call this Image A.
3. Center: Another labeled frame from the video on the right side, fixed for the entire video, also displaying bright numeric IDs on each visible furniture part. Call this Image B.

Both Image A and Image B remain constant throughout the video.

Note: The numeric IDs in Image A and Image B are not necessarily the same; the same ID may refer to different parts in each image.

You will be asked multiple-choice questions referring to the assembly video and the labeled parts shown on the left and central images.

• • •

Your responsibilities:

1. Use the right-side video frames to observe the assembly steps.
2. Use the fixed left-side labeled frame to identify and relate the furniture parts mentioned in the question.

#### For TRACK Questions

2. Use the fixed labeled frames on the left side (Image A) and in the center (Image B) to identify and relate the furniture parts mentioned in the question.

3. Carefully read the question and all answer choices.
4. Select the correct answer based on the video and labeled parts.
  - Respond **\*\*only\*\*** with a JSON object containing a single key, "answer", whose value is the letter of your chosen option (e.g., "A", "B", "C").

**\*\*Do not include any explanations or additional text -- reply with only the JSON string.\*\***

Now answer the following question:

Figure S4. **Task Instructions for Collage Prompts.** *Assumptions* remain the same (denoted by . . .), but the *Inputs* and *Instructions* are modified to describe the prompt format. In this case, the prompt image(s) are attached to the left of the video frame. For TRACK questions, the jumbled visual prompt is the second (middle) image in each frame.

Table S3. **Contact-Reasoning Results.** We show the performance of Qwen2.5-VL (32B & 72B) across two question templates. The overall performance is quite poor across all settings.

Model	Template	Micro Avg.	Yes	No
Qwen2.5VL-32B	1	62.60	41.20	84.00
Qwen2.5VL-32B	2	64.33	52.93	75.73
Qwen2.5VL-72B	1	58.80	33.33	84.26
Qwen2.5VL-72B	2	63.00	46.40	79.60

**Tracking Issues.** We used SAM2 [23] (sam2.1\_hiera\_l) for our tracking experiments. We used trimmed videos for this experiment. SAM2 cannot handle very long videos, so we implemented a chunk-wise propagation algorithm. First, we split the video into 256-frame clips with a single frame

overlap between adjacent clips. Then, we first ran SAM2 on the clip containing the prompt frame. After this, we propagate the masks forward in the video by using the masks predicted in the last frame of the initial clip to prompt the first frame in the next clip. Similarly, we propagate the masks backward from the initial clip by using the masks predicted in the first frame of the initial clip to prompt the last frame in the preceding clip. We continue this process until the entire video has been covered. We subsampled frames by a factor of 4 to maintain a trade-off between compute constraints and temporal continuity.

### D.3. Linguistic Prompt Engineering

This section provides additional details about the linguistic prompt engineering experiments discussed in Sec. 5.1. Ini-

## Task Instructions for Concat Prompts

You are a furniture-assembly expert. You are given a video input of the assembly process of a furniture item. The first frame of the video is a frame taken from the assembly process, highlighted with each visible furniture part, and displaying numeric IDs on each furniture part.

You will be asked multiple-choice questions referring to the assembly video and the labeled parts shown on the left.

### For TRACK Questions

You are a furniture-assembly expert. You are given a video input of the assembly process of a furniture item. The first two frames of the video are frames taken from the assembly process, highlighted with each visible furniture part, and displaying numeric IDs on each furniture part. We will refer to them as Image A (the first frame) and Image B (the second frame).

Note: The numeric IDs in Image A and Image B are not necessarily the same; the same ID may refer to different parts in each image. You will be asked multiple-choice questions referring to the assembly video and the labeled parts shown in the first frame.

• • •

Your responsibilities:

1. Use the video frames after the first frame to observe the assembly steps.
2. Use the first labeled frame to identify and relate the furniture parts mentioned in the question.

### For TRACK Questions

1. Use the video frames after the first two frames to observe the assembly steps.
2. Use the first two labeled frames, i.e., Image A and Image B, to identify and relate the furniture parts mentioned in the question.

3. Carefully read the question and all answer choices.
4. Select the correct answer based on the video and labeled parts.
  - Respond **\*\*only\*\*** with a JSON object containing a single key, "answer", whose value is the letter of your chosen option (e.g., "A", "B", "C").

**\*\*Do not include any explanations or additional text -- reply with only the JSON string.\*\***

Now answer the following question:

Figure S5. **Task Instructions for Concat Prompts.** *Assumptions* remain the same (denoted by . . .), but the *Inputs* and *Instructions* are modified to describe the prompt format. In this case, the prompt image(s) are concatenated as the initial frames of the video. For TRACK questions, the jumbled visual prompt is the second frame in the video.

1. Each question consists of:
  - a. A furniture assembly video.
  - b. 1-2 visual prompts or frames extracted from the video, with certain parts shaded and labeled.
  - c. An MCQ question with at most 4 options.
2. Examine the labeled frame(s) and understand the relationships between the highlighted parts using the video.
3. Read the question and then determine the correct option.

1. Each question consists of:
  - a. 1-2 images showing different stages in the assembly of a furniture item. Certain parts of the furniture will be shaded and labeled.
  - b. An MCQ question about the furniture assembly process with at most 4 options.
2. First, examine the labeled images(s) and understand the relationships between the highlighted parts.
3. Read the question and then determine the correct option.
4. Even if you cannot infer the answer from the question and the image alone, we request that you use your best judgment and select an option.

Figure S6. **Human Evaluation Instructions.** **Left:** Instructions for the standard task, where participants were provided the assembly video, visual prompt, and question text. **Right:** Instructions for the image-only task, they were provided only the visual prompts and question text.

tially, we added “Please explain this answer step-by-step” to the task instruction. However, we found that the model was just paraphrasing its selected option. Hence, we expanded this instruction to explicitly push the model to explain its rea-

soning. The results in Sec. 5.1, and the qualitative results discussed in App. D.1 are from this modified prompt. We also amended the *Instructions* portion of the task instruction to add an `explanation` key to the output JSON. Figure S10

#### Question 4

Identify the correct set of matches from Image A to Image B.

#### Options

- A. Part 0 in Image A matches Part 0 in Image B, Part 6 in Image A matches Part 6 in Image B, Part 5 in Image A matches Part 2 in Image B, Part 4 in Image A matches Part 5 in Image B, and Part 2 in Image A matches Part 4 in Image B
- B. Part 0 in Image A matches Part 6 in Image B, Part 6 in Image A matches Part 0 in Image B, Part 5 in Image A matches Part 2 in Image B, Part 4 in Image A matches Part 5 in Image B, and Part 2 in Image A matches Part 4 in Image B
- C. Part 0 in Image A matches Part 6 in Image B, Part 6 in Image A matches Part 0 in Image B, Part 5 in Image A matches Part 4 in Image B, Part 4 in Image A matches Part 2 in Image B, and Part 2 in Image A matches Part 5 in Image B
- D. Part 0 in Image A matches Part 0 in Image B, Part 6 in Image A matches Part 6 in Image B, Part 5 in Image A matches Part 5 in Image B, Part 4 in Image A matches Part 4 in Image B, and Part 2 in Image A matches Part 2 in Image B

0 1 2 3 4 Help

Not sure

Comment:

Reason you're not sure, or any other comment

Submit Task




Figure S7. **Prolific Human Annotation Portal**. Users can scroll through the questions using the numbered buttons. Once they attempt all options, a button appears prompting them to submit their responses. They can access the instructions at any time using the “Help” button.

shows the complete task instructions. In SC-CoT [31], for temperature sampling, we set the temperature to 0.7, and  $\text{top}_k$  and  $\text{top}_p$  set to 1 and 40, respectively [37].

## D.4. Image-only Prompts

In this section, we discuss the experimental details for the image-only prompts discussed in Sec. 5.3. Figure S11 shows the task instruction used for our image-only experiments. The remaining settings were the same as our main zero-shot evaluation in Sec. 4.

## E. Analysis on InternVL3

Here, we replicate the experiments discussed in Sec. 5 for the best model on our benchmark (Tab. 2), InternVL3-78B [48].

**Linguistic Prompt Engineering.** We follow the same prompts and generation settings as Qwen2.5-VL-72B (App. D). We use Concat prompts with Key-frame videos as they are the best choice for this model (See Tab. S6). Tab. S4 shows our results. ZS-CoT does lead to a minor improvement on TRACK but it does not improve overall performance, while SC-CoT leads to a significant decline in accuracy across all tasks. Thus, similar to Qwen2.5-VL-72B, linguistic prompting strategies do not aid the spatio-temporal reasoning abilities of InternVL3-78B.

**Visual Prompt Ablation.** Figure S12 shows the results of visual prompt ablation on InternVL3-78B. Firstly, we look at the different choices video and prompt types (Fig. S12(a)). For InternVL3-78B the choice of visual prompt type does have a big impact as Concat prompt perform significantly

Table S4. **Results of Linguistic Prompting Strategies for InternVL3-78B**. Similar to Qwen2.5-VL-72B, both ZS-CoT and SC-CoT fail to improve performance on FLAT-PACK BENCH.

	Micro Avg.	TORD	TLOC	TRACK	MATE
–	41.03	43.87	39.80	42.02	34.48
ZS-CoT	41.03	41.29	37.86	44.75	33.33
SC-CoT	31.56	33.55	28.15	32.68	28.73

Table S5. **Image-only Prompt for InternVL3**. Performance of InternVL3-78B using image-only prompts, along with the change ( $\Delta$ ) in performance from when the video is included in the prompt.

	Micro Avg.	TORD	TLOC	TRACK	MATE
<i>Video+Image</i>	41.03	43.87	39.81	42.02	34.48
Image-only	27.57	32.19	36.89	20.23	26.44
$\Delta$	-13.46	-9.68	-2.91	-21.79	-8.05
Shuffled Parts	29.35 $\pm$ 1.25	35.48 $\pm$ 1.94	39.16 $\pm$ 2.02	19.20 $\pm$ 0.98	36.78 $\pm$ 1.15
$\Delta$	-11.68	-8.39	-0.65	-22.83	+2.30
<b>Human Perf.</b>	42.69	30.32	45.63	48.24	44.82

better, unlike the mixed-media setting for Qwen2.5-VL-72B. This is intuitive as InternVL3 was primarily trained on image-text and video-text sequence but not all three together (image-video-text sequences) [48]. As before, the video type does not impact performance significantly. With concat prompts and key-frame videos, we observe similar trends on color scheme, mark type and marker size to Qwen2.5-VL-72B.

**Image-only Results.** On image-only prompts, InternVL3-78B shows clear degradation in performance across all tasks (See Tab. S5). This suggests that it uses the video context

## Video Difficulty Annotation Instructions

You will be given a set of furniture assembly videos. Your task is to watch the video in its entirety, and then answer some questions about the presence of certain artifacts or characteristics in the video and how they affected your ability to understand the assembly process.

Please judge difficulty with respect to answering fine-grained questions about part interactions, ordering, and attachments (not merely recognizing the final assembled object).

- Zoomed-in/out cameras:** Are there frames where the camera is too zoomed-in/zoomed-out such that it becomes difficult to understand the assembly video? Answer based on overall impact on your understanding, not peak moments where this occurs. Please specify this difficulty on a scale of 1-3 where:
  - 1 means zoom-in/zoom-out effects did not affect your understanding of the video at all
  - 2 means you could not make out when 1-2 of the parts were attached due to the camera being too close or too far
  - 3 means you could not understand  $> 50\%$  of the assembly process
- Cut-shots:** Please assign a score to the video based on the number of cuts
  - 0-1 cuts: 1
  - 2-3 cuts: 2
  - $> 3$  cuts: 3
- Out-of-frame rotations:** How difficult is it to understand out-of-frame (outside the camera's view) rotations/attachments of parts in the video? Rate this on a scale of 1-3 where:
  - 1 means that the parts go out-of-frame for  $< 10$  frames and even when they do, they remain partially visible or their motion is very unambiguous
  - 2 means that a noticeable number of parts ( $\geq 20\%$ ) go completely out of frame for a few frames (10-50 frames), but you are still able to make out how they moved based on the other in-frame parts and the rest of the video
  - 3 means that a noticeable number of parts ( $\geq 20\%$ ) go out of frame for a significant duration of time ( $> 50$  frames), and you need to use guess work or your common-sense to understand how they have been manipulated when outside of the camera's view.
- Tracking or tracing the motion of moving parts:** We are trying to understand the amount of motion and interactions between the parts in the video and whether it makes it difficult to understand which part was connected where. Rate the video from 1 to 3 where:
  - 1 means that the parts are interacted with minimally, typically only just before they are about to be attached.
  - 2 means that some of the parts move around significantly, undergoing rotations even before they are about to be attached, but by 1-2 repeated viewings of the video, it is possible to understand the assembly process
  - 3 means that the parts undergo complex motions and rotations during the video, and for some parts it is even difficult to be sure about its position in every frame. You had to use some guess work and elimination to track the parts through the full video.
- Visual effects:** Does the video have visual effects like picture-in-picture, side-by-side visuals, etc.?
  - 0 if no, 1 if yes
- Text/Visual overlay:** Some of the videos you see might have text or pictures from the furniture assembly guide overlaid on the frames. Depending on how difficult did the overlaid content make it for you to understand the video either due to obscuring the visuals or providing ambiguous instructions, please assign a score from 1 to 3 where:
  - 1 means that there was no overlaid context or even if it was there, it does not obscure the visuals or make it difficult to understand the assembly process
  - 2 means that even though overlaid content is present in some frames, it is translucent or does not obscure the visual content in the scene or adds to your understanding of the assembly process.
  - 3 means that the overlaid content in some frames is opaque and completely obscures the visual content in the scene, making it difficult for you to understand the assembly and track the parts in the video
- Initial orientation mismatch:** Does the assembly start in a flipped, upside-down, or non-canonical orientation relative to standard furniture assembly guides?
  - 0: Canonical/upright orientation
  - 1: Non-canonical (e.g., upside-down or rotated)

Figure S8. **Instructions for annotation difficulty of videos.** We asked annotators to rate the videos in our benchmark on different characteristics. The scores are then summed across characteristics and averaged across annotators to obtain a single score for the video.

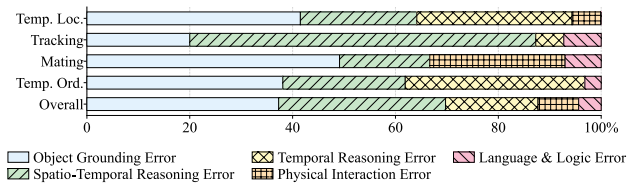


Figure S9. **Error Types for Rationales.** We provide the full quantitative distribution of error types for the rationales generated by Gemini 2.5 Pro, as discussed in Sec. 5.4.

more effectively, compared to Qwen2.5-VL-72B. This is also consistent with our observation that InternVL3-78B outperforms Qwen2.5-VL-72 on the main benchmark. Despite this, the earlier trends of majority of the decline in performance stemming from TRACK persist, suggesting that even InternVL3-78B does not video context very effectively.

These results show that InternVL3-78B displays largely similar trends to Qwen2.5-VL-72B across all the questions we discussed in Sec. 5. This suggests that our conclusions are not specific to a single model and these findings may extend beyond any one model family or architecture.

## Task Instructions for CoT Prompting

You are a furniture-assembly expert. You are given:

1. A video of a furniture assembly in progress.
2. A single labeled frame from the same video, displaying numeric IDs on each visible furniture part. Call this a visual prompt.

### For TRACK Questions

2. Two labeled frames from the same video, called Image A and Image B, each displaying numeric IDs on visible furniture parts. Note: The numeric IDs in Image A and Image B are not necessarily the same; the same ID may refer to different parts in each image.

3. A multiple-choice question (with its list of answer options) that refers to both the video and the image.

Some additional assumptions to keep in mind:

- Two furniture parts are "connected" if they are directly attached (physically in contact), like they would be in the final assembly.
- Parts simply touching each other physically, but not in their final assembly position, are not considered "connected".
- While deciding the order of "connection" events between two parts, consider the order in which they were connected last as their actual connection, i.e., if two parts are connected, then disconnected and then connected again, the last connection is the one that matters.
- In the course of the video, some parts may be attached physically to each other before they are fixed in their final assembly position by screwing, gluing, or otherwise securing them. In such cases, the parts are considered "connected" from the moment they are physically attached, not when they are fixed. Hence, when answering questions about the order of connections, consider the physical attachment as the connection event, even if the parts are not yet secured to each other.

Your responsibilities:

- Watch and interpret the assembly steps in the video.
- Examine the labeled image frame to understand part relationships.
- Read the question and all answer choices.
- Determine the correct option.
- Please explain your answer step-by-step. This explanation should contain your reasoning process, including how you analyzed the video and image to arrive at your conclusion. Do not simply restate the question or the answer choices.
- Respond **\*\*only\*\*** with a JSON object containing two keys -
  - A) "answer", whose value is the letter of your chosen option (for example "A", "B", "C", etc.).
  - B) "explanation", whose value is a step-by-step explanation of your answer. This must not be a restatement of the question or the answer choices, but rather a detailed reasoning process that leads to your answer. For example, describe how you analyzed the video and image to arrive at your conclusion or why certain options were eliminated.

**\*\*Do not include any explanations or additional text - reply with only the JSON string.\*\***

Now answer the following question:

Figure S10. **Task Instructions for CoT.** We ask the model to explain its response, explicitly asking it not to paraphrase the correct option.

## Task Instructions for Image-only Experiments

You are a furniture-assembly expert. You are given:

1. A still image showing a step in the assembly process of a furniture item, with the furniture parts shaded and labeled, displaying numeric IDs on each visible furniture part. Call this a visual prompt.

### For TRACK Questions

2. Two still images, Image A and Image B showing different steps in the assembly process of a furniture item, with the furniture parts shaded and labeled, displaying numeric IDs on each visible furniture part.

2. A multiple-choice question (with its list of answer options) that refers to the visual prompt.

Some additional assumptions to keep in mind:

- Two furniture parts are "connected" if they are directly attached (physically in contact), like they would be in the final assembly.
- Parts simply touching each other physically, but not in their final assembly position, are not considered "connected".

Your responsibilities:

- Observe and interpret the assembly step shown in the image.

### For TRACK Questions

- Observe and interpret the assembly steps shown in the images.

- Examine the labeled image frames to understand part relationships.
- Read the question and all answer choices.
- Determine the correct option.
- Respond **\*\*only\*\*** with a JSON object containing a single key, `"answer"`, whose value is the letter of your chosen option (e.g., `"A"`, `"B"`, `"C"`).

**\*\*Do not include any explanations or additional text - reply with only the JSON string.\*\***

Now answer the following question:

Figure S11. **Image-only Task Instructions.** For image-only questions, the *Inputs* only consist of 1-2 prompt images. We remove any mention of the video from the *Assumptions* and the *Instructions*.

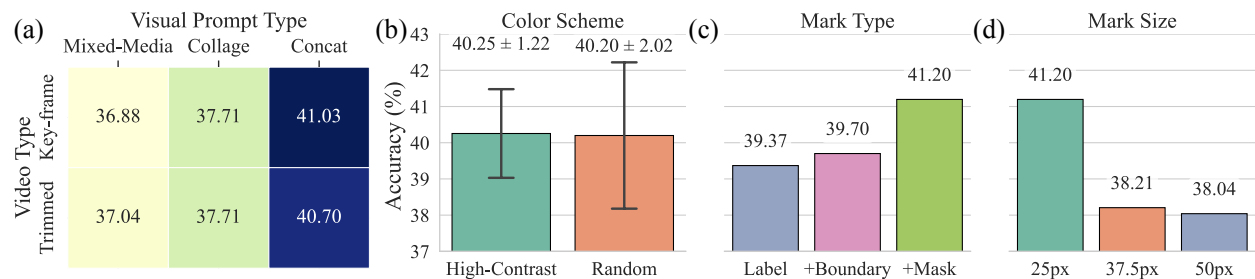


Figure S12. **Visual Data Ablation for InternVL3-78B.** InternVL3 performs better on Concat prompts with Key-frame videos, unlike Qwen2.5-VL 72B. However, when it comes to rendering the visual prompt images, both seem to follow similar trends

Table S6. **Full Zero-shot Results on FLAT-PACK BENCH.** In this table, we show the performance of all the models that we evaluated on each setting of video type and visual prompt type.

Model	Prompt Type	Video Type	Micro Avg.	95% CI	TORD	TLOC	TRACK	MATE
Video-LLaVA-7B	Collage	Key-frame	21.26	[15.65, 26.88]	21.94	30.10	10.51	41.38
Video-LLaVA-7B	Collage	Trimmed	20.76	[15.37, 26.10]	22.58	26.21	10.89	40.23
Video-LLaVA-7B	Concat	Key-frame	22.59	[18.02, 27.11]	23.87	32.04	12.45	39.08
Video-LLaVA-7B	Concat	Trimmed	21.59	[16.80, 26.37]	23.23	29.13	12.06	37.93
Video-LLaVA-7B	Mixed-Media	Key-frame	23.75	[18.12, 29.40]	21.29	35.92	10.89	51.72
Video-LLaVA-7B	Mixed-Media	Trimmed	23.59	[18.00, 29.23]	21.94	34.95	10.51	51.72
InternVL3-14B	Collage	Key-frame	37.71	[32.02, 43.41]	42.58	21.36	37.74	48.28
InternVL3-14B	Collage	Trimmed	37.04	[32.58, 41.29]	40.00	28.16	35.41	47.13
InternVL3-14B	Concat	Key-frame	34.88	[30.46, 38.87]	37.42	31.07	33.46	39.08
InternVL3-14B	Concat	Trimmed	33.55	[27.98, 38.47]	38.06	30.10	31.13	36.78
InternVL3-14B	Mixed-Media	Key-frame	35.05	[31.30, 38.59]	39.35	37.86	31.13	35.63
InternVL3-14B	Mixed-Media	Trimmed	35.88	[31.88, 40.19]	38.71	39.81	31.52	39.08
InternVL3-38B	Collage	Key-frame	34.72	[29.82, 39.26]	37.42	33.98	26.85	54.02
InternVL3-38B	Collage	Trimmed	35.05	[30.20, 39.41]	38.06	33.98	26.85	55.17
InternVL3-38B	Concat	Key-frame	36.05	[31.38, 40.06]	42.58	37.86	25.68	52.87
InternVL3-38B	Concat	Trimmed	35.22	[30.60, 39.23]	37.42	41.75	26.07	50.57
InternVL3-38B	Mixed-Media	Key-frame	33.06	[28.49, 37.91]	45.81	47.57	15.95	43.68
InternVL3-38B	Mixed-Media	Trimmed	31.73	[26.77, 36.44]	45.16	41.75	15.56	43.68
InternVL3-78B	Collage	Key-frame	37.71	[32.31, 42.72]	32.90	31.07	43.97	35.63
InternVL3-78B	Collage	Trimmed	37.71	[32.22, 42.73]	33.55	33.01	42.80	35.63
InternVL3-78B	Concat	Key-frame	41.03	[36.21, 45.64]	43.87	39.81	42.02	34.48
InternVL3-78B	Concat	Trimmed	40.70	[35.81, 45.00]	44.52	37.86	42.02	33.33
InternVL3-78B	Mixed-Media	Key-frame	36.88	[32.66, 41.27]	33.55	44.66	36.96	33.33
InternVL3-78B	Mixed-Media	Trimmed	37.04	[33.33, 40.92]	36.77	44.66	36.96	28.74
Qwen2.5-VL-32B	Collage	Key-frame	35.88	[31.18, 40.38]	34.84	29.13	33.07	54.02
Qwen2.5-VL-32B	Collage	Trimmed	32.72	[28.40, 36.92]	34.19	24.27	28.02	54.02
Qwen2.5-VL-32B	Concat	Key-frame	28.57	[23.86, 33.06]	32.90	30.10	18.68	48.28
Qwen2.5-VL-32B	Concat	Trimmed	27.57	[22.95, 32.23]	32.26	28.16	16.73	50.57
Qwen2.5-VL-32B	Mixed-Media	Key-frame	32.72	[28.27, 37.16]	38.71	33.98	24.51	44.83
Qwen2.5-VL-32B	Mixed-Media	Trimmed	34.72	[30.40, 38.77]	38.71	42.72	26.07	43.68
Qwen2.5-VL-72B	Collage	Key-frame	37.87	[31.79, 43.22]	33.55	23.30	42.41	49.43
Qwen2.5-VL-72B	Collage	Trimmed	34.88	[29.23, 40.03]	38.06	25.24	32.68	47.13
Qwen2.5-VL-72B	Concat	Key-frame	33.06	[26.84, 38.73]	35.48	26.21	28.40	50.57
Qwen2.5-VL-72B	Concat	Trimmed	32.72	[26.72, 38.12]	34.19	24.27	29.18	50.57
Qwen2.5-VL-72B	Mixed-Media	Key-frame	40.20	[35.22, 44.60]	40.65	30.10	45.53	35.63
Qwen2.5-VL-72B	Mixed-Media	Trimmed	40.37	[34.81, 45.29]	41.29	30.10	45.14	36.78
Qwen2.5-VL-7B	Collage	Key-frame	25.58	[22.24, 28.95]	30.32	23.30	18.68	40.23
Qwen2.5-VL-7B	Collage	Trimmed	25.08	[21.44, 28.79]	30.32	24.27	17.90	37.93
Qwen2.5-VL-7B	Concat	Key-frame	28.74	[25.07, 32.39]	30.97	21.36	25.68	42.53
Qwen2.5-VL-7B	Concat	Trimmed	29.24	[25.13, 33.10]	29.68	22.33	26.07	45.98
Qwen2.5-VL-7B	Mixed-Media	Key-frame	30.23	[24.83, 35.38]	27.10	18.45	33.07	41.38
Qwen2.5-VL-7B	Mixed-Media	Trimmed	29.57	[25.27, 33.67]	30.97	21.36	28.02	41.38
Qwen3-VL-235B-A22B	Collage	Key-frame	37.21	[32.77, 41.59]	37.42	25.24	39.69	43.68
Qwen3-VL-235B-A22B	Collage	Trimmed	36.05	[29.68, 42.02]	36.77	23.30	35.41	51.72
Qwen3-VL-235B-A22B	Concat	Key-frame	34.88	[30.31, 39.31]	38.06	33.01	32.30	39.08
Qwen3-VL-235B-A22B	Concat	Trimmed	34.72	[28.65, 40.31]	36.13	24.27	32.30	51.72

(table continues on next page)

(continued from previous page)

Model	Prompt Type	Video Type	Micro Avg.	95% CI	TORD	TLOC	TRACK	MATE
Qwen3-VL-235B-A22B	Mixed-Media	Key-frame	33.55	[28.60, 39.05]	37.42	38.83	28.79	34.48
Qwen3-VL-235B-A22B	Mixed-Media	Trimmed	32.23	[27.85, 36.27]	35.48	40.78	26.07	34.48
Qwen3-VL-30B-A3B	Collage	Key-frame	35.38	[30.74, 40.11]	33.55	24.27	38.52	42.53
Qwen3-VL-30B-A3B	Collage	Trimmed	35.22	[29.94, 40.36]	31.61	22.33	36.96	51.72
Qwen3-VL-30B-A3B	Concat	Key-frame	36.05	[30.65, 41.31]	34.84	20.39	42.02	39.08
Qwen3-VL-30B-A3B	Concat	Trimmed	36.71	[30.96, 42.17]	30.32	22.33	42.02	49.43
Qwen3-VL-30B-A3B	Mixed-Media	Key-frame	36.54	[32.25, 40.82]	35.48	35.92	36.19	40.23
Qwen3-VL-30B-A3B	Mixed-Media	Trimmed	36.38	[32.22, 40.76]	36.13	32.04	38.91	34.48
Qwen3-VL-32B	Collage	Key-frame	35.05	[30.37, 39.16]	36.13	32.04	33.07	42.53
Qwen3-VL-32B	Collage	Trimmed	34.05	[28.57, 39.24]	40.65	25.24	31.13	41.38
Qwen3-VL-32B	Concat	Key-frame	32.06	[26.96, 36.86]	36.13	28.16	28.02	41.38
Qwen3-VL-32B	Concat	Trimmed	33.06	[27.63, 38.35]	35.48	24.27	31.13	44.83
Qwen3-VL-32B	Mixed-Media	Key-frame	37.71	[33.09, 42.11]	38.71	46.60	31.91	42.53
Qwen3-VL-32B	Mixed-Media	Trimmed	35.71	[31.51, 39.88]	38.71	44.66	28.40	41.38
Qwen3-VL-32B-Think	Collage	Key-frame	31.56	[26.42, 36.20]	36.13	26.21	27.63	41.38
Qwen3-VL-32B-Think	Collage	Trimmed	40.03	[33.82, 45.50]	38.71	22.33	45.53	47.13
Qwen3-VL-32B-Think	Concat	Key-frame	24.25	[20.55, 28.13]	34.19	21.36	18.29	27.59
Qwen3-VL-32B-Think	Concat	Trimmed	34.88	[30.09, 39.29]	35.48	25.24	33.85	48.28
Qwen3-VL-32B-Think	Mixed-Media	Key-frame	34.72	[29.98, 39.07]	41.29	33.01	30.35	37.93
Qwen3-VL-32B-Think	Mixed-Media	Trimmed	35.38	[30.30, 39.91]	40.00	33.98	32.30	37.93
Qwen3-VL-4B	Collage	Key-frame	35.55	[30.69, 40.10]	32.26	33.98	31.13	56.32
Qwen3-VL-4B	Collage	Trimmed	35.38	[31.03, 39.52]	31.61	31.07	33.85	51.72
Qwen3-VL-4B	Concat	Key-frame	32.56	[27.21, 37.92]	32.90	25.24	29.96	48.28
Qwen3-VL-4B	Concat	Trimmed	36.54	[31.72, 41.20]	34.19	33.01	32.68	56.32
Qwen3-VL-4B	Mixed-Media	Key-frame	34.39	[29.45, 39.21]	32.90	30.10	32.68	47.13
Qwen3-VL-4B	Mixed-Media	Trimmed	34.55	[29.46, 39.43]	30.97	29.13	34.24	48.28
Qwen3-VL-4B-Think	Collage	Key-frame	27.57	[21.96, 32.45]	30.32	25.24	19.46	49.43
Qwen3-VL-4B-Think	Collage	Trimmed	31.73	[26.13, 37.16]	34.84	25.24	26.85	48.28
Qwen3-VL-4B-Think	Concat	Key-frame	32.23	[27.05, 36.64]	32.90	29.13	29.96	41.38
Qwen3-VL-4B-Think	Concat	Trimmed	37.21	[32.11, 41.78]	31.61	25.24	37.74	59.77
Qwen3-VL-4B-Think	Mixed-Media	Key-frame	27.41	[21.65, 32.31]	29.03	25.24	23.35	39.08
Qwen3-VL-4B-Think	Mixed-Media	Trimmed	28.90	[23.72, 33.42]	28.39	31.07	23.74	42.53
Qwen3-VL-8B	Collage	Key-frame	28.57	[24.44, 32.87]	36.13	32.04	20.62	34.48
Qwen3-VL-8B	Collage	Trimmed	25.58	[21.02, 31.00]	31.61	23.30	17.90	40.23
Qwen3-VL-8B	Concat	Key-frame	28.07	[23.31, 33.39]	32.26	27.18	22.57	37.93
Qwen3-VL-8B	Concat	Trimmed	29.40	[24.41, 34.18]	33.55	23.30	26.46	37.93
Qwen3-VL-8B	Mixed-Media	Key-frame	33.72	[29.31, 38.21]	36.13	30.10	33.85	33.33
Qwen3-VL-8B	Mixed-Media	Trimmed	31.73	[26.70, 37.28]	34.19	31.07	31.91	27.59
Qwen3-VL-8B-Think	Collage	Key-frame	28.41	[23.60, 32.55]	30.97	27.18	21.40	45.98
Qwen3-VL-8B-Think	Collage	Trimmed	29.07	[24.14, 33.70]	35.48	26.21	20.23	47.13
Qwen3-VL-8B-Think	Concat	Key-frame	25.25	[20.37, 29.74]	32.90	31.07	14.79	35.63
Qwen3-VL-8B-Think	Concat	Trimmed	27.41	[22.83, 31.52]	34.19	21.36	19.07	47.13
Qwen3-VL-8B-Think	Mixed-Media	Key-frame	26.58	[22.00, 30.82]	34.84	26.21	19.46	33.33
Qwen3-VL-8B-Think	Mixed-Media	Trimmed	31.73	[27.12, 35.80]	34.19	33.01	25.29	44.83
Perception-LM-1B	Collage	Key-frame	27.74	[22.85, 31.96]	28.39	26.21	25.29	35.63
Perception-LM-1B	Collage	Trimmed	27.41	[22.92, 31.37]	27.10	25.24	26.07	34.48
Perception-LM-1B	Concat	Key-frame	27.57	[23.48, 31.23]	24.52	28.16	26.46	35.63
Perception-LM-1B	Concat	Trimmed	27.41	[23.44, 30.88]	23.87	27.18	26.85	35.63
Perception-LM-3B	Collage	Key-frame	29.40	[24.95, 33.63]	27.74	32.04	26.46	37.93

(table continues on next page)

(continued from previous page)

Model	Prompt Type	Video Type	Micro Avg.	95% CI	TORD	TLOC	TRACK	MATE
Perception-LM-3B	Collage	Trimmed	31.40	[27.15, 35.26]	28.39	32.04	29.96	40.23
Perception-LM-3B	Concat	Key-frame	29.40	[24.38, 34.12]	29.03	33.98	26.46	33.33
Perception-LM-3B	Concat	Trimmed	28.90	[24.13, 33.13]	28.39	31.07	27.24	32.18
Perception-LM-8B	Collage	Key-frame	35.22	[29.53, 40.28]	25.16	26.21	44.75	35.63
Perception-LM-8B	Collage	Trimmed	35.38	[29.29, 41.22]	26.45	26.21	44.75	34.48
Perception-LM-8B	Concat	Key-frame	29.90	[24.55, 34.66]	20.65	19.42	38.13	34.48
Perception-LM-8B	Concat	Trimmed	30.90	[25.77, 35.68]	25.16	22.33	37.35	32.18
Gemini 2.5 Flash	Collage	Key-frame	23.26	[19.93, 26.29]	25.81	33.98	12.06	39.08
Gemini 2.5 Flash	Collage	Trimmed	18.77	[15.23, 21.91]	19.35	30.10	10.89	27.59
Gemini 2.5 Flash	Concat	Key-frame	27.41	[24.01, 31.05]	32.90	42.72	13.23	41.38
Gemini 2.5 Flash	Concat	Trimmed	23.09	[18.94, 27.76]	27.10	33.01	11.28	39.08
Gemini 2.5 Flash	Mixed-Media	Key-frame	31.06	[27.12, 35.45]	31.61	41.75	23.35	40.23
Gemini 2.5 Flash	Mixed-Media	Trimmed	26.25	[21.55, 30.63]	27.74	33.01	20.23	33.33
Gemini 2.5 Pro	Concat	Key-frame	32.22	[28.92, 35.85]	39.35	51.45	18.67	36.78
Gemini 2.5 Pro	Mixed-Media	Key-frame	33.72	[30.16, 37.11]	40.65	44.66	23.35	39.08
Gemini 3.1 Pro	Concat	Key-frame	32.89	[28.53, 37.13]	34.84	43.69	21.79	49.42
GPT-5	Mixed-Media	Key-frame	37.71	[32.79, 42.89]	40.65	53.40	25.68	49.43
LLaVA-Next-Vid-34B	Collage	Key-frame	27.91	[23.24, 32.71]	29.68	28.16	28.02	24.14
LLaVA-Next-Vid-34B	Collage	Trimmed	26.91	[22.58, 31.38]	30.97	26.21	26.07	22.99
LLaVA-Next-Vid-34B	Concat	Key-frame	28.07	[22.98, 33.44]	31.61	24.27	29.57	21.84
LLaVA-Next-Vid-34B	Concat	Trimmed	29.07	[24.53, 33.67]	32.26	24.27	30.74	24.14
LLaVA-Next-Vid-34B	Mixed-Media	Key-frame	28.90	[24.44, 33.08]	27.74	26.21	31.52	26.44
LLaVA-Next-Vid-34B	Mixed-Media	Trimmed	30.40	[25.87, 34.69]	30.32	24.27	32.68	31.03
LLaVA-Next-Vid-7B	Collage	Key-frame	23.26	[20.14, 26.46]	29.03	26.21	15.56	32.18
LLaVA-Next-Vid-7B	Collage	Trimmed	25.08	[22.12, 28.39]	33.55	24.27	16.73	35.63
LLaVA-Next-Vid-7B	Concat	Key-frame	23.75	[20.46, 26.86]	29.68	29.13	10.89	44.83
LLaVA-Next-Vid-7B	Concat	Trimmed	24.25	[21.23, 27.14]	30.97	29.13	11.67	43.68
LLaVA-Next-Vid-7B	Mixed-Media	Key-frame	24.58	[21.50, 27.37]	26.45	26.21	17.12	41.38
LLaVA-Next-Vid-7B	Mixed-Media	Trimmed	24.42	[21.66, 27.24]	27.10	26.21	15.18	44.83
LlaVA-OneVision-72B	Collage	Key-frame	37.87	[32.87, 42.55]	34.84	26.21	37.74	57.47
LlaVA-OneVision-72B	Collage	Trimmed	38.37	[32.97, 43.26]	35.48	25.24	38.91	57.47
LlaVA-OneVision-72B	Concat	Key-frame	36.05	[31.45, 40.09]	35.48	27.18	33.46	55.17
LlaVA-OneVision-72B	Concat	Trimmed	36.05	[31.32, 40.33]	35.48	25.24	34.24	55.17
LlaVA-OneVision-7B	Collage	Key-frame	29.07	[25.60, 32.33]	23.87	25.24	30.35	39.08
LlaVA-OneVision-7B	Collage	Trimmed	28.24	[24.46, 31.39]	23.87	20.39	29.18	42.53
LlaVA-OneVision-7B	Concat	Key-frame	32.39	[28.09, 36.16]	25.16	31.07	33.07	44.83
LlaVA-OneVision-7B	Concat	Trimmed	32.89	[28.76, 36.48]	26.45	30.10	34.24	43.68
LlaVA-OneVision-7B	Mixed-Media	Key-frame	31.23	[26.83, 35.14]	27.74	29.13	33.46	33.33
LlaVA-OneVision-7B	Mixed-Media	Trimmed	31.56	[27.64, 35.13]	28.39	26.21	34.63	34.48
LLaVA-Video-72B	Collage	Key-frame	37.54	[31.93, 42.53]	36.77	27.18	35.80	56.32
LLaVA-Video-72B	Collage	Trimmed	34.39	[29.64, 38.63]	33.55	21.36	35.02	49.43
LLaVA-Video-72B	Concat	Key-frame	34.05	[28.85, 38.97]	38.71	27.18	27.24	54.02
LLaVA-Video-72B	Concat	Trimmed	33.39	[28.82, 37.39]	40.00	25.24	25.68	54.02
LLaVA-Video-7B	Collage	Key-frame	28.57	[23.85, 33.06]	23.23	28.16	23.74	52.87
LLaVA-Video-7B	Collage	Trimmed	28.57	[24.88, 32.70]	26.45	28.16	24.12	45.98
LLaVA-Video-7B	Concat	Key-frame	30.73	[25.33, 36.06]	30.97	24.27	25.68	52.87
LLaVA-Video-7B	Concat	Trimmed	27.74	[22.11, 33.33]	24.52	21.36	26.85	43.68
ArrowRL-7B	Collage	Key-frame	28.90	[24.40, 33.02]	36.77	26.21	20.62	42.53
ArrowRL-7B	Collage	Trimmed	24.92	[20.51, 29.29]	27.74	21.36	20.23	37.93

(table continues on next page)

(continued from previous page)

<b>Model</b>	<b>Prompt Type</b>	<b>Video Type</b>	<b>Micro Avg.</b>	<b>95% CI</b>	<b>TORD</b>	<b>TLOC</b>	<b>TRACK</b>	<b>MATE</b>
ArrowRL-7B	Concat	Key-frame	30.07	[26.68, 33.14]	34.84	24.27	25.68	41.38
ArrowRL-7B	Concat	Trimmed	30.56	[26.71, 34.19]	30.97	24.27	29.18	41.38
ArrowRL-7B	Mixed-Media	Key-frame	29.90	[25.36, 34.35]	33.55	23.30	28.40	35.63
ArrowRL-7B	Mixed-Media	Trimmed	30.40	[25.63, 35.18]	30.97	28.16	28.02	39.08
Gemini 2.5 Pro + GenS	Collage	Key-frame	25.58	[21.70, 29.67]	33.55	32.04	13.23	40.23
VideoRefer	—	Key-frame	28.57	[28.50, 33.33]	32.90	30.10	17.51	51.72

## References

- [1] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025. [arXiv:2502.13923](https://arxiv.org/abs/2502.13923) [cs]. 1, 4, 5, 6, 2
- [2] Yuwei Bao, Keunwoo Yu, Yichi Zhang, Shane Storcks, Itamar Bar-Yossef, Alex de la Iglesia, Megan Su, Xiao Zheng, and Joyce Chai. Can foundation models watch, talk and guide you step by step to make a cake? In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12325–12341, Singapore, 2023. Association for Computational Linguistics. 2
- [3] Apratim Bhattacharyya, Bicheng Xu, Sanjay Haresh, Reza Pourreza, Litian Liu, Sunny Panchal, Leonid Sigal, and Roland Memisevic. Can multi-modal LLMs provide live step-by-step task guidance? In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. 2
- [4] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. 1, 2
- [5] Jang Hyun Cho, Andrea Madotto, Effrosyni Mavroudi, Triantafyllos Afouras, Tushar Nagarajan, Muhammad Maaz, Yale Song, Tengyu Ma, Shuming Hu, Hanoona Rasheed, Peize Sun, Po-Yao Huang, Daniel Bolya, Suyog Jain, Miguel Martin, Huiyu Wang, Nikhila Ravi, Shashank Jain, Temmy Stark, Shane Moon, Babak Damavandi, Vivian Lee, Andrew Westbury, Salman Khan, Philipp Krähenbühl, Piotr Dollár, Lorenzo Torresani, Kristen Grauman, and Christoph Feichtenhofer. Perceptionlm: Open-access data and models for detailed visual understanding. [arXiv:2504.13180](https://arxiv.org/abs/2504.13180), 2025. 1, 2, 4, 5
- [6] Google AI for Developers. Gemini Thinking — Gemini API Documentation. <https://ai.google.dev/gemini-api/docs/thinking#summaries>, 2025. Accessed: 2025-11-13. 7
- [7] Google DeepMind. Gemini 2.5 pro model card. <https://modelcards.withgoogle.com/assets/documents/gemini-2.5-pro.pdf>, 2025. Accessed: 2025-11-10. 4, 7
- [8] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The” something something” video database for learning and evaluating visual common sense. In *Proceedings of the IEEE international conference on computer vision*, pages 5842–5850, 2017. 2
- [9] Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large language models are zero-shot reasoners. In *Advances in Neural Information Processing Systems*, pages 22199–22213. Curran Associates, Inc., 2022. 5
- [10] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. [arXiv preprint arXiv:2408.03326](https://arxiv.org/abs/2408.03326), 2024. 1, 4
- [11] Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. MVBench: A Comprehensive Multi-modal Video Understanding Benchmark, 2024. [arXiv:2311.17005](https://arxiv.org/abs/2311.17005) [cs]. 1, 2
- [12] Yun Li, Yiming Zhang, Tao Lin, Xiangrui Liu, Wenxiao Cai, Zheng Liu, and Bo Zhao. STI-Bench: Are MLLMs Ready for Precise Spatial-Temporal World Understanding?, 2025. [arXiv:2503.23765](https://arxiv.org/abs/2503.23765) [cs]. 2
- [13] Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. [arXiv preprint arXiv:2311.10122](https://arxiv.org/abs/2311.10122), 2023. 4, 2
- [14] Benlin Liu, Yuhao Dong, Yiqin Wang, Zixian Ma, Yansong Tang, Luming Tang, Yongming Rao, Wei-Chiu Ma, and Ranjay Krishna. Coarse correspondences boost spatial-temporal reasoning in multimodal language model. In *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*, pages 3783–3792, 2025. 6
- [15] Yunong Liu, Cristobal Eyzaguirre, Manling Li, Shubh Khanna, Juan Carlos Niebles, Vineeth Ravi, Saumitra Mishra, Weiyu Liu, and Jiajun Wu. IKEA manuals at work: 4d grounding of assembly instructions on internet videos. In *NeurIPS Datasets and Benchmarks Track*, 2024. 2, 3, 8, 1
- [16] Zuyan Liu, Yuhao Dong, Ziwei Liu, Winston Hu, Jiwen Lu, and Yongming Rao. Oryx MLLM: On-Demand Spatial-Temporal Understanding at Arbitrary Resolution, 2025. [arXiv:2409.12961](https://arxiv.org/abs/2409.12961) [cs]. 2
- [17] Karttikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. Egoschema: A diagnostic benchmark for very long-form video language understanding. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 1, 2
- [18] Shehan Munasinghe, Hanan Gani, Wenqi Zhu, Jiale Cao, Eric Xing, Fahad Shahbaz Khan, and Salman Khan. VideoGLaMM : A Large Multimodal Model for Pixel-Level Visual Grounding in Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19036–19046, 2025. 2
- [19] OpenAI. Gpt-5 system card. <https://cdn.openai.com/gpt-5-system-card.pdf>, 2025. Accessed: 2025-10-19. 1, 2, 4, 3
- [20] Sunny Panchal, Apratim Bhattacharyya, Guillaume Berger, Antoine Mercier, Cornelius Böhm, Florian Dietrichkeit, Reza Pourreza, Xuanlin Li, Pulkit Madan, Mingu Lee, Mark Todorovich, Ingo Bax, and Roland Memisevic. What to say and when to say it: Live fitness coaching as a testbed for situated interaction. In *Advances in Neural Information Processing Systems*, pages 75853–75882. Curran Associates, Inc., 2024. 2
- [21] Prolific. Prolific. <https://www.prolific.com>, 2025. London, UK. Version used: March 2026. Accessed: 2026-03-25. 3

- [22] Qwen Team. Qwen3-vl github repository. <https://github.com/QwenLM/Qwen3-VL>, 2025. Accessed: 2025-10-20. 4, 5, 2
- [23] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollar, and Christoph Feichtenhofer. SAM 2: Segment anything in images and videos. In *The Thirtieth International Conference on Learning Representations*, 2025. 2, 7, 8, 5
- [24] Ye Sun, Hao Zhang, Henghui Ding, Tiehua Zhang, Xingjun Ma, and Yu-Gang Jiang. SAMA: Towards Multi-Turn Referential Grounded Video Chat with Large Language Models. *arXiv preprint arXiv:2505.18812*, 2025. 2
- [25] Dídac Surís, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for reasoning. *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, 2023. 2, 5, 7
- [26] Kexian Tang, Junyao Gao, Yanhong Zeng, Haodong Duan, Yanan Sun, Zhening Xing, Wenran Liu, Kaifeng Lyu, and Kai Chen. Lego-puzzles: How good are mlms at multi-step spatial reasoning? *arXiv preprint arXiv:2503.19990*, 2025. 2
- [27] Zongheng Tang, Yue Liao, Si Liu, Guanbin Li, Xiaojie Jin, Hongxu Jiang, Qian Yu, and Dong Xu. Human-centric spatio-temporal video grounding with visual transformers. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(12):8238–8249, 2021. 2
- [28] Gemini Team et al. Gemini: A Family of Highly Capable Multimodal Models, 2025. *arXiv:2312.11805 [cs]*. 1, 4, 2
- [29] Huaxiaoyue Wang, Gonzalo Gonzalez-Pumariiega, Yash Sharma, and Sanjiban Choudhury. Demo2code: From summarizing demonstrations to synthesizing code via extended chain-of-thought. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. 5
- [30] Xin Wang, Taein Kwon, Mahdi Rad, Bowen Pan, Ishani Chakraborty, Sean Andrist, Dan Bohus, Ashley Feniello, Bugra Tekin, Felipe Vieira Frujeri, et al. Holoassist: an egocentric human interaction dataset for interactive ai assistants in the real world. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 20270–20281, 2023. 2
- [31] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*, 2023. 5, 7
- [32] Xingrui Wang, Wufei Ma, Angtian Wang, Shuo Chen, Adam Kortylewski, and Alan Yuille. Compositional 4D Dynamic Scenes Understanding with Physics Priors for Video Question Answering, 2025. *arXiv:2406.00622 [cs]*. 1, 2
- [33] Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. Longvideobench: A benchmark for long-context interleaved video-language understanding. In *NeurIPS*, 2024. 2
- [34] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. NExT-QA: Next Phase of Question-Answering to Explaining Temporal Actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9777–9786, 2021. 1
- [35] Zihui Xue, Mi Luo, and Kristen Grauman. Seeing the arrow of time in large multimodal models. *arXiv preprint arXiv:2506.03340*, 2025. 2, 4, 5
- [36] Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. *arXiv preprint arXiv:2310.11441*, 2023. 3
- [37] Jihan Yang, Shusheng Yang, Anjali W. Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in Space: How Multimodal Large Language Models See, Remember and Recall Spaces. *arXiv preprint arXiv:2412.14171*, 2024. 2, 4, 7
- [38] Linli Yao, Haoning Wu, Kun Ouyang, Yuanxing Zhang, Caiming Xiong, Bei Chen, Xu Sun, and Junnan Li. Generative frame sampler for long video understanding. In *Findings of the Association for Computational Linguistics: ACL 2025*, Vienna, Austria, 2025. Association for Computational Linguistics. 4, 2
- [39] Yuqian Yuan, Hang Zhang, Wentong Li, Zesen Cheng, Boqiang Zhang, Long Li, Xin Li, Deli Zhao, Wenqiao Zhang, Yueting Zhuang, Jianke Zhu, and Lidong Bing. VideoRefer Suite: Advancing Spatial-Temporal Object Understanding with Video LLM, 2025. *arXiv:2501.00599 [cs]*. 1, 2, 4, 5
- [40] Weichen Zhang, Ruiying Peng, Chen Gao, Jianjie Fang, Xin Zeng, Kaiyuan Li, Ziyong Wang, Jinqiang Cui, Xin Wang, Xinlei Chen, and Yong Li. The Point, the Vision and the Text: Does Point Cloud Boost Spatial Reasoning of Large Language Models?, 2025. *arXiv:2504.04540 [cs]*. 2
- [41] Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. Llava-next: A strong zero-shot video understanding model. <https://llava-vl.github.io/blog/2024-04-30-llava-next-video/>, 2024. Accessed March 15, 2026. 1, 4
- [42] Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. LLaVA-Video: Video Instruction Tuning With Synthetic Data, 2025. *arXiv:2410.02713 [cs]*. 1, 4, 3
- [43] Zhu Zhang, Zhou Zhao, Yang Zhao, Qi Wang, Huasheng Liu, and Lianli Gao. Where does it exist: Spatio-temporal video grounding for multi-form sentences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10668–10677, 2020. 2
- [44] Zhicheng Zheng, Xin Yan, Zhenfang Chen, Jingzhou Wang, Qin Zhi Eddie Lim, Joshua B Tenenbaum, and Chuang Gan. Contphy: Continuum physical concept learning and reasoning from videos. In *International Conference on Machine Learning*. PMLR, 2024. 1, 2
- [45] Baichuan Zhou, Haote Yang, Dairong Chen, Junyan Ye, Tianyi Bai, Jinhua Yu, Songyang Zhang, Dahua Lin, Conghui He, and Weijia Li. UrBench: A Comprehensive Benchmark for Evaluating Large Multimodal Models in Multi-View Urban Scenarios, 2025. *arXiv:2408.17267 [cs]*. 2
- [46] Honglu Zhou, Xiangyu Peng, Shrikant Kendre, Michael S. Ryoo, Silvio Savarese, Caiming Xiong, and Juan Carlos Nibbles. Strefer: Empowering Video LLMs with Space-Time

- Referring and Reasoning via Synthetic Instruction Data, 2025. arXiv:2509.03501 [cs]. [1](#), [2](#)
- [47] Shijie Zhou, Alexander Vilesov, Xuehai He, Ziyu Wan, Shuwang Zhang, Aditya Nagachandra, Di Chang, Dongdong Chen, Eric Xin Wang, and Achuta Kadambi. Vlm4d: Towards spatiotemporal awareness in vision language models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2025. [2](#)
- [48] Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen Duan, Weijie Su, Jie Shao, Zhangwei Gao, Erfei Cui, Xuehui Wang, Yue Cao, Yangzhou Liu, Xingguang Wei, Hongjie Zhang, Haomin Wang, Weiye Xu, Hao Li, Jiahao Wang, Nianchen Deng, Songze Li, Yanan He, Tan Jiang, Jiapeng Luo, Yi Wang, Conghui He, Botian Shi, Xingcheng Zhang, Wenqi Shao, Junjun He, Yingtong Xiong, Wenwen Qu, Peng Sun, Penglong Jiao, Han Lv, Lijun Wu, Kaipeng Zhang, Huipeng Deng, Jiaye Ge, Kai Chen, Limin Wang, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhai Wang. InternVL3: Exploring Advanced Training and Test-Time Recipes for Open-Source Multimodal Models, 2025. arXiv:2504.10479 [cs]. [1](#), [4](#), [5](#), [7](#)
- [49] Orr Zohar, Xiaohan Wang, Yann Dubois, Nikhil Mehta, Tong Xiao, Philippe Hansen-Estruch, Licheng Yu, Xiaofang Wang, Felix Juefei-Xu, Ning Zhang, Serena Yeung-Levy, and Xide Xia. Apollo: An Exploration of Video Understanding in Large Multimodal Models, 2024. arXiv:2412.10360 [cs]. [2](#)