

Spacewalk-18: A Benchmark for Multimodal and Long-form Procedural Video Understanding in Novel Domains

Supplementary Material

Appendix

Table of Contents

A Dataset Construction	1
A.1 Labeling Process	2
A.2 Illustrations of Annotated Data	3
A.3 Statistics of Annotations	3
A.4 Question Answering Task	4
A.5 Temporal Certificate	8
B Evaluation Metrics of Step Recognition	8
C Implementation Details	9
C.1 Evaluated Models	9
C.2 VLLM Prompts	10
C.3 Last-layer Fine-tuning Contrastive VLMs	10
C.4 All-layer Fine-tuning Contrastive VLMs	10
C.5 Long-term Feature Bank	11
D Additional Experiments	11
D.1 Accuracy of Each Question Answering Type	11
D.2 Ablation Study on Fine-tuning Data Size	11
D.3 Video Speedup	12
D.4 Ablation on Frame Sampling on QA Task	12
D.5 Temporal Context Incorporation	13
D.6 Qualitative Results	13
E Trends on Ours and Other Benchmarks	14
F Additional Task – Intra-video Retrieval	14
F.1. Task Definition	14
F.2. Evaluation Methods	14
F.3. Evaluation Results	23
F.4. Qualitative Results	23

Expedition Date	Video Duration	Annotated Duration
Training set		
October 6, 2019	8:42:06	5:25:23
November 22, 2019	8:58:45	6:10:02
January 15, 2020	9:03:04	6:26:18
January 25, 2020	7:59:43	5:22:10
January 27, 2021	8:25:35	5:55:27
February 1, 2021	8:39:50	0:48:35
June 16, 2021	9:23:13	6:35:01
December 2, 2021	7:46:53	6:13:30
March 23, 2022	8:29:54	5:14:27
December 3, 2022	7:44:35	2:12:00
Validation set		
March 15, 2022	9:14:30	5:42:27
November 15, 2022	10:38:40	7:02:55
Test set		
November 15, 2019	8:25:23	3:50:47
December 2, 2019	7:49:37	5:33:52
June 26, 2020	8:14:13	4:08:36
February 28, 2021	9:28:07	7:03:46
September 12, 2021	8:44:29	6:50:19
June 9, 2023	8:12:25	5:46:47

Table A1. The Spacewalk recordings used in the Spacewalk-18 dataset, and their video durations. The recordings cover the spacewalk missions from 2019 to 2023. Because each spacewalk recording includes scenes from both before the spacewalk begins and after it ends, only the video segments where the spacewalk activity happens are annotated. Moreover, some spacewalk missions (e.g., February 1, 2021 and November 15, 2019) may deviate from the tasks originally planned in the animation and carry out unplanned tasks during execution. We only annotated the parts that are aligned with the planned tasks.

A. Dataset Construction

Sourcing Live Streams. First, a list of spacewalks at the International Space Station is sourced from [NASA’s website](#). For each spacewalk from 2019 to 2023, we identify whether there is a recording available on YouTube with an animation sequence near the beginning. We find a total of 18 spacewalk recordings that meet this criteria. Table A1 has a list of the NASA summary for each expedition and the duration of each YouTube video.

Building Transcripts. We use Deepgram’s Automatic Speech Recognition (ASR) service to transcribe the spacewalk videos. In order to choose an ASR algorithm, we had a human verifier analyze transcriptions from Deepgram, OpenAI’s Whisper, NVIDIA’s NeMo, YouTube’s Auto Gener-

ated Captions, and Google’s Speech-to-Text on the same spacewalk audio clip. Deepgram proved to be the most accurate. Through a similar process, we found that Deepgram’s ASR algorithm performs the best with approximately 10 minute audio clips. Therefore, we chunk the multi-hour audio files into 10 minute clips and feed them to Deepgram for transcription. This results in a list of sentences with start and end timestamps for each spacewalk video.

A.1. Labeling Process

Building our dataset requires temporally segmenting and labeling very long videos (many hours). To do this, we introduce a new annotation protocol and tool. Existing methods of collecting temporal segment annotations require multiple passes and/or for clips to be pre-labeled [20, 57]. While annotating our dataset would traditionally require three passes (action identification, temporal segmentation, and task labeling), we pre-segment the videos into clips containing a maximum of one step each, allowing us to collect annotations in a single pass. This drastically reduces the number of human worker hours required. The source code for the tool will be publicly released. The protocol is as follows:

Define the Label Space. In our case, the label space comes from the animation videos. We manually segment the animations into clips containing a single step and label each clip with a short description. We show the step lists of three spacewalk missions in Table A2, A3, and A4, one from each of the training, validation, and test sets.

Split the Videos. To reduce the burden of temporal segmentation, we split the long live streams into sub-clips that each contain at most one step. We find that due to the long nature of the steps and tendency of the camera to switch angles often in spacewalk live streams, a given shot-segmented clip (between camera angle changes) will contain at most one action/step. Thus, to split the multi-hour long spacewalk video into clips for the human workers to annotate, we employ a shot detection algorithm. PySceneDetect’s Content-Detector uses changes in color and intensity between individual frames to draw boundaries between shot changes in videos. We find that it errs on the side of over-segmenting for our spacewalk live stream videos, further ensuring that clips will contain at most one step.

Chunk into Segments. It is unreasonable to expect human workers to annotate a multi-hour video in one sitting. Therefore, we chunk the videos into segments of approximately one hour. We design these segments to contain about one hour of recording content and only contain steps from a continuous subset of the animation. Each video has 5.11 segments on average.

Annotation Interface. To collect annotations, we build the Spacewalk Video Annotation Tool, pictured in Figure A2. The human worker inputs their unique user ID and selects

Step ID	Caption
1	EV1 and EV2 exit airlock
2	EV1 heads outward and places safety tether anchors
3	EV2 retrieves foot restraint
4	EV1 goes to carrier with solar array
5	EV1 and EV2 drop of PGT and bags
6	EV2 stows foot restraint and bag
7	EV1 preps iROSA for removal
8	EV2 installs bag and tools on mod kit
9	EV2 sets up cables for future installation
10	EV2 retrieves PGT and goes to EV1
11	EV1 does more prep for iROSA removal
12	EV1 retrieves foot restraint from CETA cart
13	EV1 installs, sets up, and enters foot restraint
14	EV1 removes bolts on iROSA
15	Robotic arm moves EV1 into position
16	EV2 gets into position for iROSA release
17	EV2 prepares iROSA for release
18	EV1 removes iROSA
19	EV2 stows tools and enters foot restraint
20	EV1 carries iROSA on robotic arm to EV2
21	EV1 hands iROSA to EV2
22	EV1 exits foot restraint and goes to EV2
23	EV2 rotates into position
24	EV1 enters foot restraint and receives iROSA
25	EV2 exits foot restraint and gets into position
26	EV1 and EV2 install iROSA
27	EV1 and EV2 swing iROSA into single tube
28	EV1 and EV2 drive mounting bolts
29	EV1 and EV2 clean up and prep for next EVA
30	EV1 and EV2 return to airlock
31	EV1 and EV2 enter airlock

Table A2. Step list of the spacewalk mission on June 16, 2021, which is in the training set.

Step ID	Caption
1	Luca and Drew exit airlock with pump system
2	Luca and Drew take pump system to external support platform 2
3	Luca enter foot restraint on robotic arm
4	Drew hands Luca the pump system
5	Robotic arm takes Luca to AMS
6	Drew move to ELC 2
7	Luca and Drew install pump system
8	Luca and Drew connect power and data cables
9	Robotic arm takes Luca to aft side
10	Luca connect six fluid connections
11	Robotic arm takes Luca to underside of AMS
12	Luca and Drew complete final 2 suages
13	Robotic arm takes Luca to ESP 2
14	Drew bring bags back to airlock

Table A3. Step list of the spacewalk mission on December 2, 2019, which is in the test set.

the video date and segment number that they are assigned. The interface then loads the corresponding pre-labeled animation clips and the recording clips for them to label. The platform saves their progress as they annotate clips so they are able to start the task and return to it at any time. Once they complete the task, a completion URL appears.

Source Human Annotators. We use an online platform to source human workers to annotate the dataset. First, participants are sent the training and screening phases where they learn about the task and demonstrate proficiency in being able to accurately label a small sample of spacewalk live stream clips. In the training phase, as workers select labels for live stream clips, they are given feedback about the correctness of their annotations along with some reasoning. They spend as much time as necessary in the training phase before moving on to the screening phase. In this phase, workers are tasked with labeling a set of 10 clips without feedback, and those who achieve an accuracy of 80% or greater are selected for the annotation phase.

Annotate Clips. In the annotation phase, human workers are presented with animation and annotation clips from a single segment of a spacewalk recording (Figure A2). They watch the annotation clips and select a label for each clip from the set of animation clip labels, “Irrelevant”, or “Unsure”. The “Irrelevant” label is used to categorize any clip that does not contain footage of one of the tasks for the given spacewalk. This includes shots of the mission control center, noisy shots (*e.g.* blue screen), and shots of get-ahead tasks that were not originally planned for the spacewalk. We have three human workers annotate each clip and we choose the most commonly selected label as the true label.

Merge Adjacent Clips. We intentionally over-segment the long spacewalk recordings before collecting annotations, to avoid the same clip spanning across multiple steps. We thus include a final step to merge adjacent clips with the same label. Rather than the traditional method of having human workers provide temporal boundaries, the over-segmentation allows us to break down the challenging tasks of temporal segmentation and action recognition into a series of easier, smaller ones. We then use the collected annotations to obtain true temporal boundaries by concatenating all adjacent clips with identical labels. Figure A1 illustrates how this balances the distribution of clip durations in the dataset.

A.2. Illustrations of Annotated Data

Figure A3, A4, and A5 illustrate three examples of annotations collected from human annotators. We also include a video in the supplementary material showing a few more examples.

Step ID	Caption
1	EV1 exit airlock and receive large bag
2	EV2 exit airlock
3	EV1 move to integrated equipment assembly
4	EV1 stow bag and begin prep work
5	EV2 move to phase 1
6	EV2 stow crew bag and retrieve tools
7	EV1 and EV2 assemble upper triangle
8	EV2 move to and enter foot restraint
9	EV1 hand upper triangle to EV2
10	EV2 dock upper triangle on gimbal assembly
11	EV1 stow pistol grip tool (PGT)
12	EV2 exit foot restraint and tilt it to the left side
13	EV1 pass left mid strut to EV2
14	EV1 hands lower strut to EV2
15	EV1 and EV2 install lower strut
16	EV1 and EV2 install mid strut
17	EV2 exit foot restraint and tilt it to the right side
18	EV1 hand right mid strut to EV2
19	EV1 and EV2 install lower strut
20	EV1 and EV2 install mid strut
21	EV1 finish mid strut install
22	EV2 move to bag and stow tools
23	EV2 return to worksite and stow tools on body restraint tether
24	EV1 take pictures of completed mod kit
25	EV1 move to battery charge/discharge unit and begin prep work
26	EV2 translate to CETA cart and stow tools
27	EV2 retrieve crew lock bag and move to EV1
28	EV1 and EV2 fold and restrain insulation
29	EV1 and EV2 break torque and reinstall
30	EV1 and EV2 clean up and retrieve crew bag
31	EV2 return to airlock
32	EV1 return to airlock

Table A4. Step list of the spacewalk mission on March 15, 2022, which is in the validation set.

A.3. Statistics of Annotations

After merging adjacent clips with the same labels, we obtain in total 3,753 clips with annotated spacewalk steps. We show several distributions about the clips in Figure A6, including the durations of the video clips, total durations of the steps, and numbers of clips per step. On average, each merged clip has a length of 92 seconds, each step spans 9 minutes, and each step is composed of 5 clips.

While each spacewalk video has a list of 25 steps on average, a small portion of the steps do not necessarily occur following their order in the list. After manual examination, we find that around 84% of the adjacent steps are logically non-interchangeable in order. For example, “installing the battery” must happen after “taking the battery to the worksite”. The remaining interchangeable steps are mostly due




Figure A1. Distributions of video segment durations before and after merging adjacent clips.

to the parallel work of the astronauts. For example, if “EV 2 retrieves foot restraint” happens while “EV 1 goes to the carrier”, there is no guarantee about which step will occur first.

To analyze the diversity of the dataset, we filter the nouns and verbs from the step captions to obtain a list of objects and actions. Figure A7 and Figure A8 illustrate the occurrence of each object and action in the step captions. In total, we observe 51 objects and 47 actions across the dataset. For the 455 steps, which is the sum of the number of steps in each video, we emphasize that no steps are exactly the same because they have their own animation videos and narrations as contexts. But we manually cluster semantically similar steps and observe 167 distinct groups of steps.

A.4. Question Answering Task

A.4.1. Overview

As described in Section 3.3, our question answering task includes 376 questions. 349 of them are generated by templates using our step annotations, while the other 27 miscellaneous questions are manually created. Although the manually created questions are fewer, they are more video-specific than those generated by templates. For the template-generated questions, we further categorize them into four types: task before/at/after location, task before/after task, when task, and task order. Examples of the different question types are listed in Table A5. The distribution of the question types can be found in Figure A9. In the following, we describe the construction of each type of question in detail.

A.4.2. Template Generated Questions

For each of the four types of question templates, we first sample some one-hour-long video segments as follows: for all 3,753 clips, we take the start time of each clip as the beginning and the time one hour later as the end to form

a long video segment. We then randomly sample M segments from these. To balance the number of the sub-type questions, we set $M = 300$ for task before/at/after location, $M = 200$ for task before/after task, $M = 100$ for when task, and $M = 100$ for task order. We keep M at a moderate level in order to avoid generating repeated questions as much as possible.

For each sampled long video segment, we follow the descriptions below to construct a question. Note that some segments are invalid to construct certain types of questions. For example, videos in which all events occur at the same location cannot be used to generate task before/at/after location questions. Therefore, the final size of the generated questions (Figure A9) is smaller than the M segments we sampled.

Task before/at/after location. The template of this question type is “Which of the following tasks happens {before the astronaut arrives at}/{while the astronaut is at}/{after the astronaut leaves from} {location}?” We first manually annotate the location where each step is performed. After that, for each sampled one-hour video segment, we uniformly sample one relationship from before/at/after and find all step-location pairs that satisfy this relationship. Specifically, a step is considered before/after a location only when all sub-segments related to the step happen before/after all the sub-segments related to the location. From the set of valid pairs, we sample one to fill in the template. For the negative choices, we preferentially select the other steps from this one-hour video segment that are invalid answers to the generated question. If there are less than three of these negative choices, we randomly sample steps within the same spacewalk mission but outside the video segment.

Task before/after task. The template of this question type is “Which of the following tasks happens {before}/{after} {step}?” Similar to “task before/at/after location”, we uniformly sample one relationship from before/after and then find all step1-step2 pairs that satisfy this relationship. Step 1 is considered before/after step 2 only when all sub-segments of step 1 happen before/after all the sub-segments of step 2. From the set of valid pairs, we sample one to fill in the template. For the negative choices, we preferentially select the other steps from this one-hour video segment that are invalid answers to the generated question. If there are less than three of these negative choices, we randomly sample steps within the same spacewalk mission but outside the video segment.

When task. The template of this question type is “In which part of the video does {step} happen?” The answer can be the first third, the middle third, the last third, or the task does not happen. For a given sampled one-hour video segment, we filter the steps whose related sub-segments all fall within the same third of the video. We then uniformly sample one of them to construct the question. We did not construct

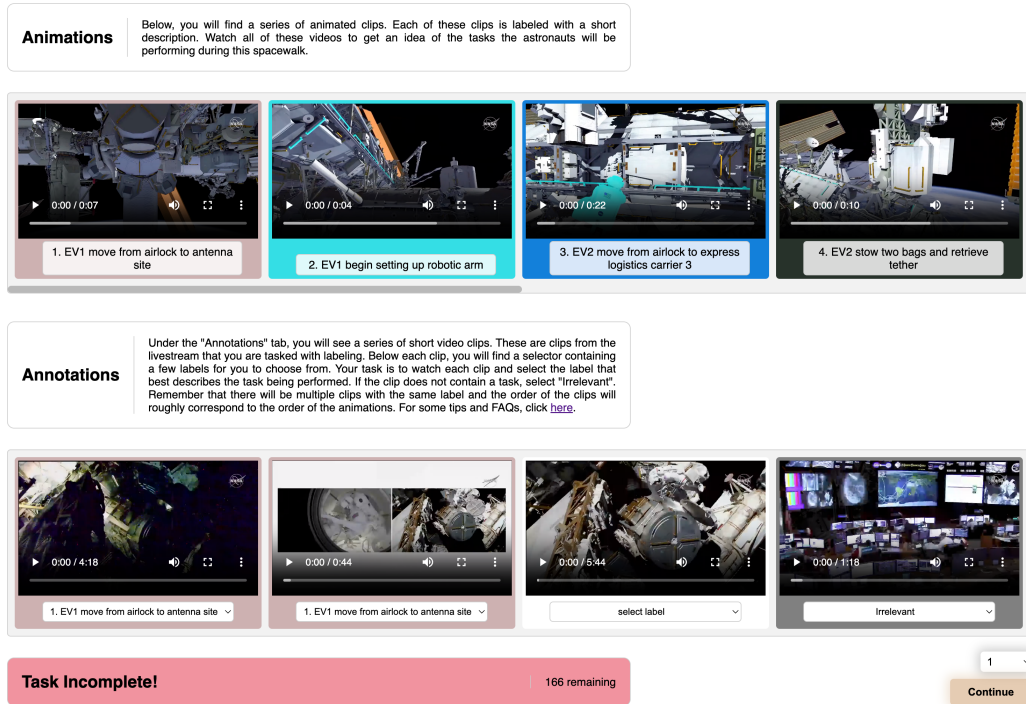


Figure A2. Spacewalk Video Annotation Tool Interface. Row (1) contains a pre-labeled set of steps from the animation video. Row (2) contains a set of live stream clips for the annotator to categorize into steps.

Spacewalk Video Clip

(C-1)_{end} C_{start} C_{middle} C_{end} (C+1)_{start}

Transcript: [1786 words]
 And just heard confirmation from Tina Cartman to the crew that they are building point five psi, and our go for hatch opening. Again, we are looking for that call for it to turn on internal battery power. And then the crew will head out to open the hatch. Okay. Frank, you might be catching on the strip there right now. K. Turkey. You're clear. Sure. That's on. And, Tina, the hatch is open and in the hatch key. Copy hatch open. Okay. Emergency m p have closed. Emergency access code. Copy, Duke. EV one and two on your DCMs. Switch power to bat, stagger your switch throws, ...

Annotator 1	Annotator 2	Annotator 3
Step 1	Step 1	Step 1



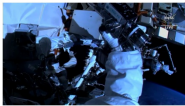


Label: Step 1

Caption:
 EV1 exits airlock and receives bags

Transcript:
 US spacewalk eighty one will be conducted by EV one, Josh Cassada, and EV two, Frank Rubio. EV one exits the airlock and gets handed one of the bags in preparation for the EVA. And the second bag, the strut bag gets handed out and put on EV one's body restraint tip.

Figure A3. Annotation example 1 from the spacewalk recording from Nov. 15, 2022 (training set). All three annotators agree on the label for the video clip.

Spacewalk Video Clip





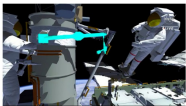






(C-1)_{end} C_{start} C_{middle} C_{end} (C+1)_{start}

Transcript: [300 words]
 Frank, there will be no changes to your PGT settings. You'll be going three to eight and a half additional turns to torque on m one through m four. Josh, for you, you'll be driving m twenty eight to torque, six and a half to nine turns. You can hand start it. Let me know when you're ready for PGT. Copy. There are no changes. You got three to eight terms. Do you want me to wait till Josh, second date? Okay. I think that's a good plan. Don't you worry about a half and a half and a half inch? On the engagement line. Yeah. I'm not sure. The hole. Yeah. Give me a look. We don't buy it again. Right. Far. Just to make it before. Correct? Yeah. Yeah. Could you see it on you can...

Annotator 1 Step 12	Annotator 2 Step 12	Annotator 3 Background
-------------------------------	-------------------------------	----------------------------------

Label: Step 12



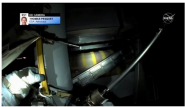

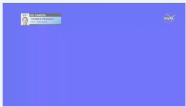






Caption:
 EV1 & EV2 install mid strut on left side

Transcript:
 Frank will apply the final torque on the bolt and get into position for the mid strike. Again, working together, the two crew members will install the mid strike. And drive the bolts to put it in place.

Figure A4. Annotation example 2 from the spacewalk recording from Nov. 15, 2022 (training set). One annotator disagrees with the other two but since the majority of annotators selected step 12, the clip is labeled as step 12. The clips on either side appear to be of the same step, which illustrates the effect of over-segmenting the spacewalk recordings and the necessity of merging adjacent clips after collecting annotations.

Spacewalk Video Clip

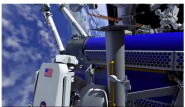
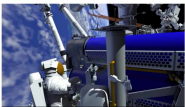
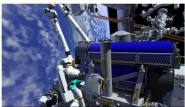

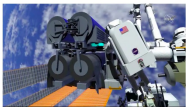






(C-1)_{end} C_{start} C_{middle} C_{end} (C+1)_{start}

Transcript:
 Understand there's not much thing you can do about that at point, you might booms are set, but that's why I'm repeating calls back to you and just correct me if we ever get it wrong. Yeah. That's all good, Jenny. And, I try to be clear and change you. Do you copy me or you hit the same for you? No. I'm copying you. Good. Okay. Copy you. Good. Copy all. You're sure it's not my friend's accident. I've thrown you off.

Annotator 1 Step 7	Annotator 2 Step 7	Annotator 3 Background	Annotator 4 Step 7
------------------------------	------------------------------	----------------------------------	------------------------------

Label: Step 7

Caption:
 EV1 preps iROSA for removal

Transcript:
 Continue translating outboard to the very end of the truss and owes his stowage bag. Toma prepares the Irosa by releasing the first hinge bolt, and stowing his pistol grip tool. Tomma makes his way outboard and releases an anti rotation device, which keeps the primary bolt locked in place during launch loads. Tomma retrieves, a handling aid for Irosa, we refer to this as a scoop.

Figure A5. Annotation example 1 from the spacewalk recording from Jun. 16, 2021 (validation set). One annotator disagrees with the other three but since the majority of annotators selected step 7, the clip is labeled as step 7. The clip immediately following this one demonstrates an example of a blue screen that would be classified as “Background”.




Figure A6. Distributions of (a) Merged video clip duration, (b) Total step duration, and (c) Video clip number per step.

Question Type	Example
Task before/at/after location	<p>Q: Which of the following tasks happens before the astronaut arrives at the external pallet?</p> <p>(A) Chris sets up tools and prepares worksite. (B) Bob moves to IEA. (C) Chris & Bob retrieve battery from slot 1. (D) Chris sets up tether.</p>
Task before/after task	<p>Q: Which of the following tasks happens after EV1 move back inboard?</p> <p>(A) EV1 retrieve portable foot restraint with extension. (B) EV1 & EV2 move to P6. (C) EV1 & EV2 install respective bags on worksites. (D) EV2 move to P1 and install anchor hooks for safety tether.</p>
When task	<p>Q: In which part of the video does the task that EV1 & EV2 install mid strut happen?</p> <p>(A) The first third of the video. (B) The middle third of the video. (C) The task does not happen in the video. (D) The last third of the video.</p>
Task order	<p>Q: In which order do the tasks happen in the video?</p> <p>(A) (1) Robotic arm takes Luca to aft side. (2) Drew move to ELC 2. (3) Robotic arm takes Luca to AMS. (B) (1) Drew move to ELC 2. (2) Robotic arm takes Luca to aft side. (3) Robotic arm takes Luca to AMS. (C) (1) Robotic arm takes Luca to AMS. (2) Drew move to ELC 2. (3) Robotic arm takes Luca to aft side. (D) (1) Drew move to ELC 2. (2) Robotic arm takes Luca to AMS. (3) Robotic arm takes Luca to aft side.</p>
Miscellaneous	<p>Q: What type of equipment does the astronaut retrieve first, and how is it utilized during the mission?</p> <p>(A) Bags containing structure to assemble modification kit. (B) Articulating portable foot restraint to later attach to the Canada arm. (C) Pump to be installed on the Alpha Magnetic Spectrometer. (D) Bags containing structure to support new solar arrays.</p>

Table A5. Question types and examples in the question answering task.

questions whose answer is “the task does not happen”.

Task order. The question of this type is “In which order do the tasks happen in the video?” The four choices are permutations of three given steps. Here, the order of the two steps refers to the order of their first occurrences. If the video contains step 1, step 2, and then step 1 again, we consider step 1 to occur before step 2. To construct the question, we uniformly sample three steps from the given one-hour video. Their ground truth order serves as the correct answer while three permutations of them are used as negative choices.

A.4.3. Manually Annotated (Miscellaneous) Questions

To manually annotate questions for spacewalk videos, we first determine a list of questions for each spacewalk mission – “What type of equipment does the astronaut retrieve

first, and how is it utilized during the mission?”, “What is the goal of this mission?”, and “What did EV1 do while EV2 is doing {step}?”. For the first two types of questions, we watch the early part of the spacewalk video to determine the answer for each video. Negative choices are randomly selected from the correct answers of other spacewalks with similar tasks. For the third type of questions, we analyzed the step annotations to find video segments where the annotation order resembles step 1, step 2, step 1. Each of these was manually analyzed to determine if EV1 and EV2 were working on separate tasks at the same time. If so, the question was formed and the correct answer annotated. Negative choices are randomly selected from steps of the same spacewalk and spacewalks with similar tasks.




Figure A7. Distribution of objects. It counts how many steps contains each object in their captions.

A.5. Temporal Certificate

Table A6 provides the numerical values we use to plot the temporal certificate figure in Figure 3. It is extended from [34] to include Spacewalk-18.

B. Evaluation Metrics of Step Recognition

We use accuracy, mAP, and IoU to evaluate a model’s performance on step recognition. To eliminate the impact of uneven numbers of task steps in each spacewalk video, we first compute the following metrics per video and then average them across videos.




Figure A8. Distribution of actions. It counts how many steps contains each action in their captions.

Accuracy. It is the percentage of timestamps whose corresponding steps are correctly recognized.

Mean Average Precision (mAP). For methods that can give a confidence score for each task step, we employ mAP as an evaluation metric. We first compute the average precision of each task step except “Irrelevant” and then take the mean.

Intersection over Union (IoU). After the models predict the step label for each sampled timestamp, we merge adjacent timestamps with the same predictions into consecutive

Dataset	Temporal Certificate	Average Clip Length	Test Duration
Kinetics	1.931	10	240000
AVA	0.25	1	117900
HVU Concept	0.77	10	650000
HVU Action	1.65375	10	650000
UCF 101	1.81	6.66	9457
Something Something	1.28	3	81471
LVU Relationship	17.12	210	5364
NextQA	2.7	44	47872
EgoSchema	100	180	90000
Youtube8M Segment	0.1	5	2520000000
MSRVTT	0.7	13	38870
IVQA	0.3	18	36000
AGQA	3.7	30	57600
How2QA	1.5	18	16263
ActivityNetQA	2.4	124.58	144000
Spacewalk-18	140	89	119664

Table A6. The numerical values of the temporal certificates. We extend this plot from [34] to include Spacewalk-18. The units for all three columns are seconds. “Test Duration” is the total duration of each test set.




Figure A9. Distribution of the question types of the question answering task.

temporal intervals. They form a segmentation of each step across the entire spacewalk video. For each step, we calculate the IoU with the ground truth segmentation of the step. Finally, we take the average over all the steps as the IoU measurement.

C. Implementation Details

C.1. Evaluated Models

In the following, we briefly introduce the evaluated models, selected checkpoints, and their frame sampling strategies. We by default use the recommended number of sampled frames for each model, rather than unifying them across models.

EgoVLP [31] is an egocentric video-language model trained on EgoClip dataset. In our step recognition experiments, we sample 4 frames from a video for feature extraction.

VideoCLIP [67] is a video-language model trained with video-text contrastive learning. It uses a Transformer to integrate S3D [66] video features and align it with text feature. 150 frames are sampled from a video to extract feature in the step recognition experiments.

InternVideo [61] is a video-language model pre-trained on a large corpus of video-language datasets including HowTo100M [35] and WebVid10M [5]. It achieves state-of-the-art performances across 39 video datasets from extensive tasks. In our step recognition experiments, we use the checkpoint further fine-tuned for video-text retrieval on MSRVTT [68]. 12 frames are sampled from each video clip in the experiments.

InternVideo2 [62] unifies masked video modeling, cross-modal contrastive learning, and next token prediction to develop a family of video foundation models. While their entire training process develops a VLLM, the training stage 2 learns video, audio, and text encoders that share an aligned feature space. In our experiments, we evaluated the stage 2 checkpoint `InternVideo2_s2-1B` and sample 4 frames for each video clip.

LLaVA-Next-Video [73] is a video large language model trained on a large video/image-language corpus. We use the checkpoint `LLaVA-NeXT-Video-34B-DPO` and uniformly sample 32 frames from each video as visual inputs on both tasks.

VideoLLaMA2 [9] is a multimodal large language model

capable of understanding video, audio, and language. It employs Spatial-Temporal Convolution connector to capture the intricate spatial and temporal dynamics in the video input. We use the VideoLLaMA2-7B checkpoint and sample 8 frames from each video on the step recognition task. We use the VideoLLaMA2-7B-16F checkpoint and sample 16 frames on the question answering task.

LongVU [47] proposes a spatiotemporal adaptive compression mechanism to reduce the number of video tokens for long-form video understanding. We use the LongVU_Qwen2_7B checkpoint. This model samples video frames at 1 FPS regardless of the video’s duration. **Qwen2.5VL** [4] employs dynamic resolution processing and absolute time encoding to deal with images and videos with various resolutions and durations. We use the checkpoint Qwen2.5-VL-7B-Instruct. This model by default samples video frames at 2 FPS. However, it has a maximum limit of 768 frames. Consequently, the number of sampled frames is 768 on our question answering task. **InternVL3** [76] develops a VLLM through native multimodal pre-training, which consolidates language pre-training and multi-modal alignment training into a single pre-training stage. We evaluate the checkpoint InternVL3-8B and sample 32 video frames on the step recognition task and 256 frames on the question answering task.

GPT-4o and **GPT-5** are a proprietary multimodal large language model API. As they receive multiple images as input but not a video file, we uniformly sample a few frames from a video and feed them to the API following the temporal order. The used checkpoint is gpt-4o-2024-05-13 for GPT-4o and gpt-5-2025-08-07 for GPT-5. The number of sampled frames 8 on the step recognition task and 32 on the question answering task.

Caption-augmented LLM. Following [72], on both of our two tasks, we use LLaVA-1.5-13B [33] to caption 60 uniformly sampled video frames with prompt “Describe the image in 30 words”. The generated captions are delivered into GPT-4o (gpt-4o-2024-05-13) to answer the given question.

VideoMAE [55] is a video encoder trained with masked video modeling. We use the videomae-base checkpoint and sample 16 frames from each video on the step recognition task.

C.2. VLLM Prompts

To evaluate VLLMs on the step recognition task, we format the task into multi-choice video question answering. Given the video content and transcript, we ask the model to choose a step index from a given list. Specifically, the VLLM prompt for the step recognition task is as following:

```
<video>
You are given a spacewalk video, where the spacewalk mission can be divided into <number of steps> steps. The transcript of the video speech is:
<transcript>.
Please provide a single-number answer (from 0 to <number of steps>) to the following multiple-choice question, and your answer must be one of the numbers from 0 to <number of steps>. You must not provide any other response or explanation. If you are not sure, answer with the most likely answer.
Here is the question: Which step does the frame in the middle of this video belong to?
Here are the choices:
(0) Irrelevant: The mission control center, noisy shots (e.g. blue screen), or tasks not planned for the spacewalk.
(1) <step 1 caption>: <step 1 transcript>.
(2) <step 2 caption>: <step 2 transcript>.
...
(N) <step N caption>: <step N transcript>.
```

For our task 2, the question answering task, we use the following prompt:

```
<video>
Please provide a single-letter answer (from A to D) to the following multiple-choice question, and your answer must be one of the letters from A to D. You must not provide any other response or explanation. If you are not sure, answer with the most likely answer.
Here is the question: <question>
Here are the choices:
(A) <option_1>
(B) <option_2>
(C) <option_3>
(D) <option_4>
```

On both tasks, we set the temperature to 0 to disable sampling during answer generation for all VLLMs except GPT-5. The GPT-5 API only supports a temperature of 1.

C.3. Last-layer Fine-tuning Contrastive VLMs

When last-layer fine-tuning the contrastive VLMs, we use Adam [25] optimizer with a learning rate of $1e-4$ and a batch size of 2048. All models are trained for 20 epochs and the checkpoint with the best validation loss is picked.

C.4. All-layer Fine-tuning Contrastive VLMs

We follow the code base of InternVideo [61] to fine-tune its entire model backbone on our step recognition task. We use Adam [25] optimizer with a learning rate of $4e-6$. The learning rate warms up linearly in the first 10% training

steps, after which cosine annealing is adapted. Besides, a weight decay of 0.2 is used. We fine-tune the model for 1 epoch with a batch size of 16. We find that training for more than 1 epoch always leads to over-fitting.

C.5. Long-term Feature Bank

To employ Long-term Feature Bank (LFB) [64] to solve the step recognition task, we first divide the spacewalk recording clip $V_{t,w}$ into k -second-long segments and extracts their InternVideo features with a frame sampling rate of 1 FPS, respectively. Since InternVideo is pre-trained with 12 frames per video, we set $k = 12$. There are in total $M = 60 \times w/12$ segments for a w -minute-long video and the query timestamp falls in segment $m = M/2$. Denote the feature of the i -th segment as f'_i . LFB aims to incorporate the history feature $f'_H = f'_{1:m-1}$ and the future feature $f'_F = f'_{m+1:M}$ into the query timestamp feature f'_m . Based off [64], we design four context incorporation mechanisms – **LFB Avg**, **LFB Cat**, **LFB NL**, and **LFB TF**. Each of these methods integrate f'_m , f'_H , and f'_F into a feature f for the entire video clip $V_{t,w}$. We substitute f for f_s in Equation (1) to train the linear layer $G_\theta(\cdot)$ and the parameters of NL/TF blocks jointly. The pre-trained video-language encoders are always frozen in this process. Unless otherwise specified, the training hyperparameters are the same as those of last-layer fine-tuning in Appendix C.3.

LFB Avg. We pool the query, history, and future features together to form a single feature, *i.e.*,

$$f = \text{Average}([f'_H, f'_m, f'_F]) = \text{Average}(f'_{1:M}). \quad (3)$$

LFB Cat. We first pool the history and future features, respectively, and then concatenate them with the query feature, *i.e.*,

$$f = [\text{Average}(f'_H), f'_m, \text{Average}(f'_F)]. \quad (4)$$

LFB NL. We stack two non-local blocks [60] to learn the interaction between the query feature and the context features via attention mechanism. The architecture of a non-local block is shown in Figure A10. Before feeding the features into non-local blocks, we first add positional encoding to them and use a linear layer to project them into a 512-dimensional space. The output of the non-local blocks is concatenated with the query feature to form the final clip feature. Formally,

$$f = [f'_m, \text{Non-local}(L(f'_m + P_m), L([f'_H + P_H, f'_F + P_F]))], \quad (5)$$

where $L(\cdot)$ is a linear projection and P_m , P_H , and P_F are positional encoding for the query, history, and future timestamps. We use a learning rate of $1e - 3$ and a batch size of 512 during training.

LFB TF. We employ a two-layer Transformer encoder upon the concatenation of the history, query, and future features,

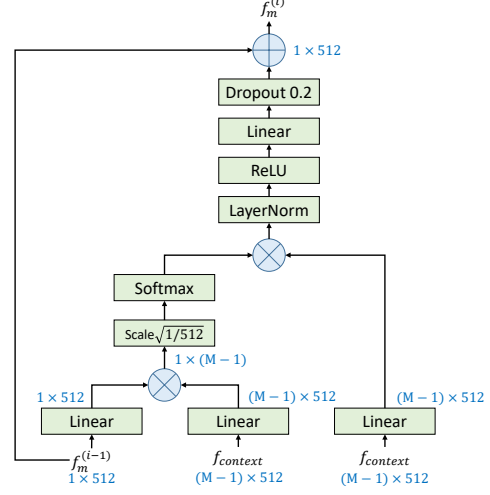


Figure A10. Architecture of a non-local block Non-local($f_m^{(i-1)}, f_{\text{context}}$) borrowed from [64]. It employs attention mechanism between the query feature $f_m^{(i-1)}$ and context feature f_{context} , and produces an updated query feature $f_m^{(i)}$. When stacking multiple non-local blocks, the query feature is iteratively updated while the context feature remains unchanged.

i.e.,

$$f = \text{Transformer}([f'_H, f'_m, f'_F] + P) = \text{Transformer}(f'_{1:M} + P), \quad (6)$$

where P is positional encoding. The hidden size is the same as the feature dimensions, which is 2304. We use a learning rate of $1e - 4$ and a batch size of 512 to train the model.

D. Additional Experiments

D.1. Accuracy of Each Question Answering Type

As discussed in Appendix A.4, our question answering task includes five different types of questions. In Table A7, we decompose the model performance in Table 2 into the five question types. We find that most of models perform the best on the task before/after task (TT) questions, while the miscellaneous (MI) questions are generally difficult to the VLLMs. However, given the transcript of the animation video, the oracle model makes great improvement on the miscellaneous questions, further highlighting the incapability of the VLLMs of abstracting the tasks in spacewalk videos. We notice that the oracle model performs poorly on the when task (WT) questions. This is expected because the time period options of these questions are based on the segmented one-hour video clip, while the oracle information given to the model covers the entire spacewalk mission.

D.2. Ablation Study on Fine-tuning Data Size

We conduct experiments to explore the impact of data size on all-layer fine-tuning InternVideo on the step recognition

Method	#Frames	TL	TT	WT	TO	MI	Acc
Random	-	25.00	25.00	25.00	25.00	25.00	25.00
LLaVA-Next-Video-34B	32	26.67	33.33	25.37	29.82	33.33	29.52
VideoLLaMA2-7B	16	25.71	40.83	20.90	36.84	29.63	31.65
InternVL3-8B	256	35.24	34.17	28.36	19.30	40.74	31.65
Qwen2.5VL-7B	768	40.00	33.33	37.31	21.05	29.63	33.78
LongVU-7B	1 FPS	32.38	38.33	50.75	24.56	33.33	36.44
Caption-enhanced LLM	60	33.33	35.00	26.87	15.79	25.93	30.37
GPT-4o	32	33.33	33.33	32.84	31.58	25.93	32.45
GPT-5	32	39.05	50.83	53.73	47.37	37.04	46.54

Table A7. Accuracy on different question types of the question answering task. TL: task before/at/after location; TT: task before/after task; WT: when task; TO: task order; MI: miscellaneous.



Figure A11. Ablation studies on training data size. (a) Varying number of training samples per video (10 training videos in total). (b) Varying number of training videos (2000 training samples from each video). We full-layer fine-tune InternVideo with a context length of 1 minute. The blue curves are the fine-tuning performances while the orange straight lines are the zero-shot performances on the corresponding metrics.

task. In Figure A11a, we ablate the number of samples drawn from each video in the training set. The model accuracy and IoU first rise as the data size grows. However, when more than 500 samples are drawn from each video (5k samples in total), the performance becomes saturated (less than 0.3% gain from 500 samples to 2k samples). This indicates that the number of training samples per video is not the bottleneck for our fine-tuning approach.

In Figure A11b, we vary the number of training videos, while 2k training samples are drawn from each video. The accuracy and IoU curves show significant upward trends when the videos get more, while mAP is saturated when more than 6 videos are used. This shows that fine-tuning on more spacewalk videos might better adapt the model the novel domain. However, the number of spacewalk videos is naturally limited by the number of real spacewalk missions, which disables large-scale data collection. So we urge the need for domain adaptation method with higher data effi-

Speedup	1×	2×	4×	8×	16×	32×
InternVideo	10.03	10.00	9.88	9.80	9.60	9.42
VideoLLaMA2	5.17	4.00	10.50	15.00	17.33	16.00

Table A8. Step recognition accuracy under different speedups.

Number of frames	32	64	128	256
Accuracy	29.52	29.26	30.85	31.65

Table A9. Ablation on frame sampling on QA task with InternVL3.

Number of frames	8	16	32	60
Accuracy	31.65	30.32	32.45	31.91

Table A10. Ablation on frame sampling on QA task with GPT-4o.

ciency.

All these experiments are conducted under a context window length of 1 minute. For different training data sizes, we keep the number of fine-tuning epochs the same. Therefore, the less training data, the smaller number of fine-tuning steps. However, we find that naively enlarging the number of training steps always leads to over-fitting.

D.3. Video Speedup

Astronauts on spacewalks sometimes have slower motion than earthly activities, which may cause temporal out-of-domain issues to video-language models. To investigate this factor, we speed up the videos at different rates and evaluate InternVideo and VideoLLaMA2 at 1 FPS with 5-minute contexts. In Table A8, the original speed suits InternVideo better, while speedup boosts VideoLLaMA2.

D.4. Ablation on Frame Sampling on QA Task

Videos in our question answering task are hour-long. The number of sampled frames might be essential to represent the video content. Therefore, although each model has a suggested number of frames per query, we may increase the number of frames to achieve better performance on long-form videos.

Method	$w = 1$ min			$w = 2$ min			$w = 3$ min			$w = 5$ min		
	Acc.	mAP	IoU	Acc.	mAP	IoU	Acc.	mAP	IoU	Acc.	mAP	IoU
EgoVLP	6.34	8.92	2.17	7.73	9.83	2.59	9.68	10.66	3.03	10.21	10.86	3.18
VideoCLIP	8.40	9.80	3.21	9.01	10.40	3.28	9.98	11.14	3.71	8.87	10.39	3.46
InternVideo	10.12	12.17	4.02	11.22	12.30	4.22	11.13	12.68	4.04	10.08	12.53	4.04
LLaVA-Next-Video	10.35	-	3.77	12.03	-	4.53	13.71	-	5.07	13.82	-	4.92
VideoLLaMA2	9.34	-	2.88	12.43	-	4.32	14.37	-	5.45	17.32	-	6.28

Method	$w = 10$ min			$w = 15$ min			$w = 20$ min		
	Acc.	mAP	IoU	Acc.	mAP	IoU	Acc.	mAP	IoU
EgoVLP	8.63	10.13	2.65	7.42	9.37	2.12	8.33	9.61	2.48
VideoCLIP	8.08	10.44	3.36	8.47	9.79	3.79	7.23	8.82	3.20
InternVideo	10.26	12.36	4.18	9.66	11.28	3.54	9.94	10.53	3.87
LLaVA-Next-Video	11.85	-	4.58	11.80	-	-	12.05	-	4.00
VideoLLaMA2	20.98	-	8.03	19.72	-	-	17.49	-	6.69

Table A11. Last-layer fine-tuning performances of contrastive VLMs and zero-shot performances of VLLMs on the step recognition task under varying context window lengths. They are used to plot Figure 4.

Method	$w = 1$ min			$w = 2$ min			$w = 3$ min			$w = 5$ min		
	Acc.	mAP	IoU	Acc.	mAP	IoU	Acc.	mAP	IoU	Acc.	mAP	IoU
Sparse Frame Sampling	10.12	12.17	4.02	11.22	12.30	4.22	11.13	12.68	4.04	10.08	12.53	4.04
Dense Frame Sampling	10.28	12.54	4.20	11.98	12.98	4.60	11.75	13.51	4.13	11.14	13.19	4.43
LFB Avg	11.47	13.92	4.39	11.98	15.15	4.49	11.71	16.14	4.29	12.37	17.25	3.73
LFB Cat	12.22	13.43	4.49	12.76	14.25	4.82	12.98	14.83	4.95	13.12	15.46	4.76
LFB NL	9.31	13.23	3.61	10.22	13.96	3.48	9.50	14.98	3.26	10.70	17.28	3.01
LFB TF	11.31	13.08	3.77	11.67	14.83	3.71	12.45	15.68	4.11	11.38	16.10	3.04

Method	$w = 10$ min			$w = 15$ min			$w = 20$ min		
	Acc.	mAP	IoU	Acc.	mAP	IoU	Acc.	mAP	IoU
Sparse Frame Sampling	10.26	12.36	4.18	9.66	11.28	3.54	9.94	10.53	3.87
Dense Frame Sampling	10.72	13.23	4.20	-	-	-	-	-	-
LFB Avg	10.86	18.11	3.35	11.39	18.65	3.03	10.77	18.29	2.68
LFB Cat	13.20	16.69	4.89	11.98	16.43	4.52	12.78	16.41	4.84
LFB NL	12.78	16.81	4.20	10.57	17.55	3.64	10.08	17.41	3.39
LFB TF	13.18	16.96	5.00	12.98	17.06	4.42	11.17	18.79	3.35

Table A12. Performances of different temporal context incorporation methods on the step recognition task. These methods are built upon frozen InternVideo. They are used to plot Figure 5.

To investigate if this is true, we conduct an ablation study with InternVL3 and GPT4o. The results are in Table A9 and Table A10. For InternVL3, increasing the number of frames from 32 to 256 improves the accuracy from 29.52% to 31.65%, indicating the effectiveness of dense frame sampling. While InternVL3 with 256 frames already saturates the GPU memory of a single RTX A6000, we did not conduct experiments with more frames. On the contrary, increasing the number of sampled frames did not improve GPT-4o’s performance. This shows that the effect of dense sampling is model-dependent and cannot generally boost performance.

D.5. Temporal Context Incorporation

In Section 5.5, we test models under varying context window lengths to investigate their capability to incorporate temporal context. The numerical results used in Figure 4 are listed in Table A11. Those used in Figure 5 are in Table A12.

D.6. Qualitative Results

In this section, we demonstrate qualitative results on the step recognition task and the question answering task. While we provide a few key frames for each video in our figures, the corresponding videos can be found in the supplementary material.

We illustrate examples on the step recognition task with a context window length $w = 5$ min from Figure A12 to Figure A17. The model makes predictions mainly based off keywords shared by the transcript of the spacewalk video and the caption/transcript of the step animation. Besides, it is also aware of the objects (*e.g.*, adapter plate in Figure A12) and actions (*e.g.*, moving around in A13) appearing in both the spacewalk and animation videos. While the models correctly recognize some steps, these cues can also lead to failure cases. Sharing keywords in the transcripts does not necessarily mean that the spacewalk clip belongs to the step. The models need to better utilize the semantic details in the spacewalk video and transcript to determine

Models	Spacewalk-SR	Spacewalk-QA	EgoSchema [34]	VideoMME [16]
InternVideo	9.5	-	32.1	-
LLaVA-Next-Video-7B	11.0*	25.53*	44.6	35.6
VideoLLaMA2-7B	17.3	31.7	51.7	47.9
Qwen2.5VL-7B	23.0	33.8	65.0	65.1
GPT4o	26.4	32.5	72.2	71.9

Table A13. Model performance reported on ours and other benchmarks. *: Different from Table 2 because that uses a 34B checkpoint.

the corresponding step. Moreover, all-layer fine-tuning corrects InternVideo’s predictions in some examples. The model learns the visual concept of airlock in Figure A14 and swages in Figure A15. Finally, there are still difficult examples for the evaluated models (Figure A16 and A17), demonstrating the room for model improvement.

In Figures A18 to A20, we show qualitative results on the question answering task. In Figure A18, LLaVA-Next-Video, GPT-4o, and Caption-enhanced LLM recognize that astronauts are connecting the cables after leaving from the robotic arm. However, In Figure A19, only GPT-5 and GPT-4o with animation oracle correctly identify the temporal relations between tasks. Some of the models choose “(D) Luca and Drew remove debris shield”, which indeed happens before they hand off the debris shield. In Figure A20, many of the models have difficulty localizing the spacewalk task of interest.

E. Trends on Ours and Other Benchmarks

In Table A13, we list a few models’ performance on Spacewalk-18 and other video benchmarks. Although our benchmark evaluates novel domain generalization of VLMs, the models exhibit a similar trend on other benchmarks, probably indicating better generalization capability of stronger models. However, all the models perform much worse on Spacewalk-18 than on other benchmarks, highlighting the necessity of models with better domain generalization.

F. Additional Task – Intra-video Retrieval

In our step recognition task, the list of steps in the corresponding spacewalk mission is provided. However, the pre-defined steps are not always available for all the spacewalk mission. In this case, the segmentation of spacewalk steps must be done in a unsupervised manner. Current approaches to unsupervised video segmentation rely on clustering video features [15, 27]. This requires models to form distinguishable embeddings across steps. With this in mind, we explore an additional task – intra-video retrieval. This task evaluates the models capability of retrieving video clips that are from the same step as a query video, without any description about the step.

F.1. Task Definition

In a spacewalk video, given a query timestamp t_q and two candidate timestamps t_{c1} and t_{c2} , the task is to determine which of t_{c1} and t_{c2} is of the same step as t_q . Note that closer timestamps are more likely to be in the same step. Therefore, to avoid this shortcut, we involve only two equidistant candidates satisfying $|t_q - t_{c1}| = |t_q - t_{c2}|$, where one of them to a different step and serves as a hard negative. To make them distinguishable, we ensure that the candidates are at least 30 seconds away from the query. Moreover, both the query and candidates should be at least 15 seconds away from their corresponding step boundaries to avoid ambiguity. Similar to the step recognition task, we set a context window length w and the model can access the video and transcript in the w -long windows centered at t_q , t_{c1} , and t_{c2} .

We construct 2000 samples from each spacewalk video for this task. For each annotated spacewalk clip, we first uniformly sample several query timestamps t_q ’s. Then, for each t_q , we randomly sample a distance d such that only one of $t_q - d$ and $t_q + d$ is in the same step as t_q and they follow the restrictions in the task definition. Finally, we randomly choose one of $t_q - d$ and $t_q + d$ as t_{c1} and the other as t_{c2} . We adjust the number of samples from each spacewalk clip to ensure roughly uniform numbers of queries from different spacewalk steps.

For performance measurement, we calculate the retrieval accuracy as the evaluation metric. Random chance is 50% accuracy on this task.

F.2. Evaluation Methods

F.2.1. Intra-video Retrieval by VLLMs

Zero-shot. To evaluate VLLMs on the intra-video retrieval task, we format it as multi-choice video question answering by asking “Which candidate shows the same step as the query?” after providing the query and candidate videos. Specifically, the prompt is as following:

You are given three spacewalk video clips from a spacewalk mission, one query clip and two candidate clips. The mission can be divided into multiple steps. Please provide a single-number answer (1 or 2) to the following question, and your answer must be either 1 or 2. You must not provide any other response or explanation. If you are not sure, answer with the most likely answer.


Here is the question: Which candidate shows the same step as the query?

Query video clip:

Eight uniformly sampled video frames: <query video frame 1> ... <query video frame 8>

Video speech transcript: <query video transcript>


Spacewalk video clip



Transcript:

The **adapter plate** is soft dock. Copy that, Bob. You can retrieve the PGT with the hex driver. That's the one that's on battery four. And settings will be alpha seven clockwise two. And, Chris, for you, as you get back on structure, you'll wanna prepare the e p for translation. So make sure the ingress aid is stowed towards the boot plate and tethers are clear..... I've got the alpha seven Yep. And counters or clockwise t set. Moving right to drive on each key. Affirm h two. You'll confirm one line flush. Looking for sixteen to seventeen turns. One line flush. K. For sixteen to seventeen turns, here we go..... And Cassidy just completed the work to install the old battery in the slot that was emptied for disposal, and banking is still working to **install the adapter plate** in open slot number two on the truss. Let me confirm the two lines. One of them flush. Verify. And can you again remember the front? It'll be four and a half to five and a half turns.....

Label: Step 23



Caption:

Bob **installs adapter plate** in open slot 2.

Transcript:

Bob translates back to the truss to **install the adapter plate** in open slot two.

Predictions:

Zero-shot InternVideo:	Step 23 ✓	Fine-tuned InternVideo:	Step 23 ✓
LLaVA-Next-Video:	Step 17 ✗	VideoLLaMA2:	Step 23 ✓
LongVU:	Step 23 ✓	Qwen2.5-VL:	Step 23 ✓
InternVL3:	Step 23 ✓	Caption-enhanced LLM:	Step 23 ✓
GPT-4o:	Step 23 ✓	GPT-5:	Step 23 ✓

Step 17



Caption:

Chris & Bob retrieve battery from slot 2.

Transcript:

Next, the battery in slot two is removed and translated over to the pallet.

Figure A12. In this example of step recognition, the spacewalk video transcript mentions that the astronaut is installing an adapter plate, and the video also shows them tightening the screws on the adapter plate. All models except LLaVA-Next-Video match this video with the step of installing adapter plate.

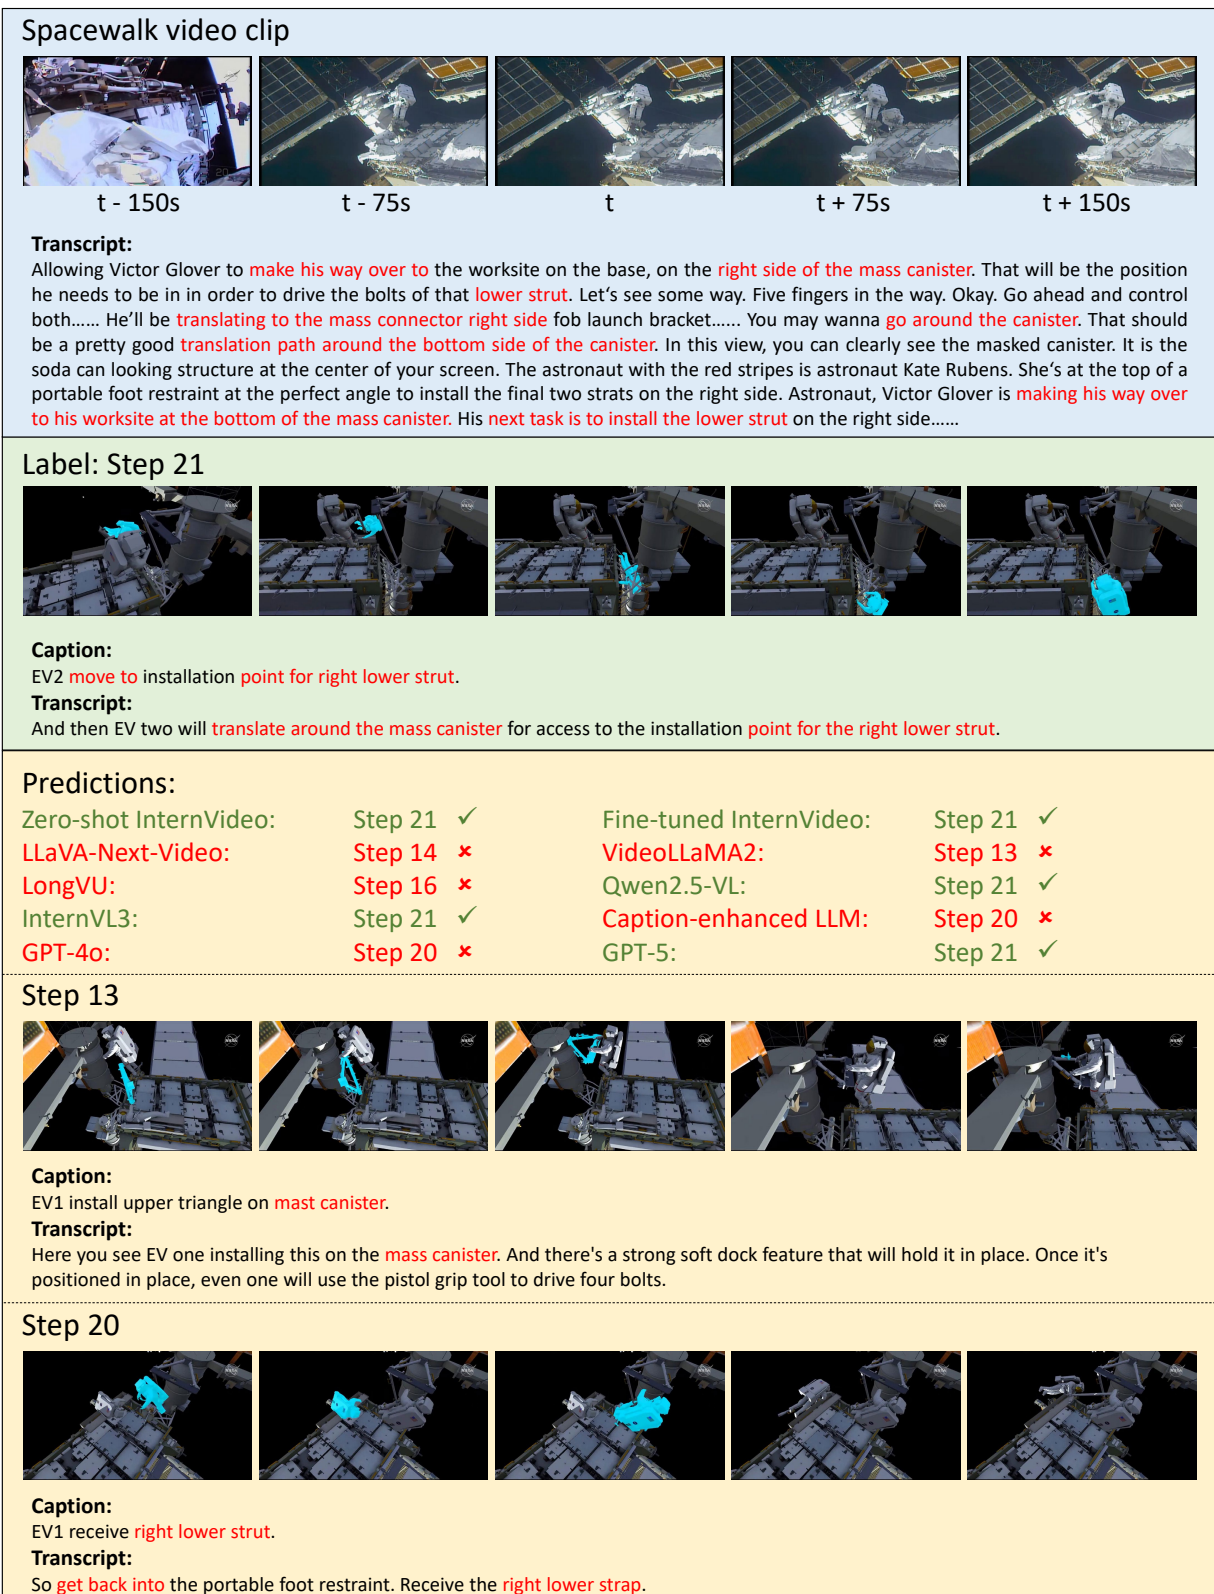


Figure A13. In this example of step recognition, both the spacewalk video and transcript indicate that the astronaut is translating around the mass canister, which aligns with the transcript of step 21. However, five of the VLLMs mismatch it, probably based on keywords “mast canister” (step 13) and “lower strut” (step 20).

Spacewalk video clip



Transcript:

His poles have been blocked. I'm going for the log. Copy. EV one anchor..... Floting in the **airlock**, you can see the **strut bag**. That is a specially designed bag that will that is carrying the pieces for the modification kit that will be built today. E v one, then I see a green light. Okay. To my I am in position, so we The the the camera as it's exposed camera is on the forward camera side. K. Eight. Are you ready for the **bag**? Yeah..... That **strap bag** carries the parts that will be used to assemble the modification kit this morning. Okay. I think you can look right. Straight back to your lock. You have a light on it. And here it comes the back..... I have control. Okay. Right now. Uh-huh. Go ahead. I'm gonna stay on the outside. What are you with that? I think I can disconnect my waist tether from the dial up during extender issue. Yes. You have a go for that. Yeah. We concur as well. Alright. And, Akie, while you're, waiting for Samantha **egress**.....

Label: Step 1



Caption:

EV1 & EV2 **exit airlock**.

Transcript:

US EVA seventy seven will begin at the Quest **Airlock**. Jackson astronaut, Aki Hochde, EV one, noted by the red stripes, will **egress** first and receive and hold the eight foot mod kid's **strut bag**. European astronaut to Ma Pes Gay will **egress** second with the full white suit.

Predictions:

Zero-shot InternVideo:	Step 14 ✗	Fine-tuned InternVideo:	Step 1 ✓
LLaVA-Next-Video:	Step 1 ✓	VideoLLaMA2:	Step 1 ✓
LongVU:	Step 2 ✗	Qwen2.5-VL:	Step 11 ✗
InternVL3:	Step 4 ✗	Caption-enhanced LLM:	Step 4 ✗
GPT-4o:	Step 1 ✓	GPT-5:	Step 1 ✓

Step 4



Caption:

EV2 stow **strut bag**.

Transcript:

While tomorrow is doing the very large strip bag at p four, Aki will complete his translation out to the p four beta Gimbal assembly where he will install and set up his quick fix.

Step 14



Caption:

EV2 collect tools and prepare for right side.

Transcript:

Small then gather tools and reset for the right hand side. Rubbing your tools back to the **bag** while arching **egresses** and reposition the foot restraint to bias it to the right side.

Figure A14. In this example of step recognition, the spacewalk video shows the astronauts exiting the airlock. While zero-shot InternVideo makes an incorrect prediction, fine-tuned InternVideo learns the concept of airlock and gives the true answer. The mistakes made by zero-shot InternVideo, InternVL3, and caption-enhanced LLM are probably due to keywords “strut bag” (step 4) and “egress” (step 14).

Spacewalk video clip



t - 150s

t - 75s

t

t + 75s

t + 150s

Transcript:

Hey, Drew. Next, I can get the straightener. Okay. In hand. It looks like this card is unlocked. K. There we go. Back to me, administrator. The right length's been on the straightener..... Luca. Oh, okay. K. I have the **tape installed**. Check me down the gauge. Okay. Copy that. No. I have to retake it. It's moved to auto doing it. Okay. Copy that. Able to get it off completely. That it's not it's not completely, but I won't be able to reboot it. Certainly not representative of Major's work, amazing work on this project.... For a good **swage**, the tube has to be put into the alternate fitting, a very specific length marked off by the Kapton **tape**, measured by the tube straightener, both in the hands of Luca Parmitano right now. Those precise measurements, required for a good **swage** to make sure it's pinched at just the right point. This one's on. Primatano just working to make sure that that, measurement is precise.....

Label: Step 10



Caption:

Luca **connect** six fluid **connections**.


Transcript:

There'll be six fluid **connections** completed here at this VSP work site using, commercial off the shelf **swages**. Those **swages** have any encased in a, custom fitting that the engineering team to develop to allow the astronauts pressurized gloves to operate the small ferrals of the **swage**.

Predictions:

Zero-shot InternVideo:	Step 7 ✗	Fine-tuned InternVideo:	Step 10 ✓
LLaVA-Next-Video:	Step 14 ✗	VideoLLaMA2:	Step 1 ✗
LongVU:	Step 12 ✗	Qwen2.5-VL:	Step 11 ✗
InternVL3:	Step 12 ✗	Caption-enhanced LLM:	Step 10 ✓
GPT-4o:	Step 11 ✗	GPT-5:	Step 12 ✗

Step 7



Caption:

Luca & Drew install pump system.

Transcript:

Drew will provide eyes on to assist Luca in aligning system and install it onto AMS into the mechanical attachment device.

Step 11



Caption:

Robotic arm takes Luca to underside of AMS.

Transcript:

Once those six **swages** are completed, the SSRMS will take Luca to the underside of AMS for two final **swages**.

Figure A15. In this example of step recognition, the spacewalk videos shows the astronaut connecting swages, while its transcript also mentions them. However, only fine-tuned InternVideo and caption-enhanced LLM make the correct predictions. The mismatch to step 11 by GPT-4o is possibly because of the keyword “swage”.

Spacewalk video clip



t - 150s

t - 75s

t


t + 75s

t + 150s

Transcript:

Am I having something? Oh, okay. Good..... I'll just do the five pound. Good words. Hey, Reddit. Kindly missed your ret. Nope. You have three left. Okay. Oh, where's the refi? I can do it. Yeah. Yeah. Okay. I'm already turned out for sure. And, as in some IPs, I have the car. Okay. Toma, you'll be retrieving the Okay. Okay. Next tomorrow, you'll be **retrieving the left lower stretch** under straps five and six. And **handing that off**. That's all the slots five and six coming up. I mean, I got it's red on my BRT with the rest. It's rattles out. Copy. Good config, Okey. Pesk now working to treat the **left lower strut** from the strap bag. Once he has that, he'll pass it along to Hoshide, who is in the AP far or articulating portable foot restraint? Once Hoshide has the lower strut, he will install it to the mounting bracket using the PGT or the pistol grip tool, which has been used quite so far this morning. And tomorrow, you'll be handing the clevis bolt side.....

Label: Step 11



Caption:

EV1 enter foot restraint and **receive strut**.

Transcript:

Once complete with the upper triangle, Aki will egress the foot restraint, bias it to the left hand side while Tamal prepares the **hand a left mid strut** for BRT stow. Aki will ingress the APFR and **hold the lower strut** while Tama reposition to the solar array blanket box for saab bearing to install the left lower struts.

Predictions:

Zero-shot InternVideo:	Step 28 ✗	Fine-tuned InternVideo:	Step 28 ✗
LLaVA-Next-Video:	Step 1 ✗	VideoLLaMA2:	Step 11 ✓
LongVU:	Step 12 ✗	Qwen2.5-VL:	Step 12 ✗
InternVL3:	Step 12 ✗	Caption-enhanced LLM:	Step 12 ✗
GPT-4o:	Step 11 ✓	GPT-5:	Step 12 ✗

Step 12




Caption:

EV1 & EV2 install **lower left strut**.

Transcript:

Tomah will begin driving this bolt by hand four turns while Aki aligns his end and drives his to the mounting bracket two turns. Tomaw will finish his bolt by driving it with a pistol grip tool, followed by high torque with the EVA torque wrench. Once the bolt is deemed good, off you will be given a go to drive his bolt of torque using the crystal grip tool. This will complete our minimum config for the mod kit.

Step 28



Caption:

EV1 stow spare FPMU.

Transcript:

And meet Tamah at the FPMU at p one. He will stow the spare in a location optimal for hand off.

Figure A16. In this example of step recognition, the spacewalk video is dark but the transcript includes a command that one astronaut would hand off a strut to the other. As this example is complicated to understand, only VideoLLaMA2 and GPT-4o predict the step label correctly.

Spacewalk video clip



Transcript:

The alpha seven clockwise two. You're going to h two. Confirm socket tape line flush. Looking for sixteen to seventeen turns. Alpha seven, a price two on h two. And if they find this flush okay. Matures moving. Are you appears to be going ahead? I agree. This one seems snugger up against the wall than the other guy..... You'll confirm the tape line flush looking for four to five and a half turns. Copy. You can remove it? It's twelve thousand indicator is not quite to a block nine point two And it looks like we had about that four turns. Copy. Four turns. And torque was nine point two, but the status indicator doesn't quite look locked. Status indicator on **battery** is maybe a quarter inch for a block. Okay. We're good with that. You can hand that PGT back to Chris to Stone a swing arm and release that RET. That far. Plus with the JSON patterns, everything looks straight.....

Label: Step 31



Caption:

Chris & Bob **install new battery** in empty slot 3.

Transcript:

They repeat the steps to release the bolts, translate back to the truss, And together, install an empty slot number three.

Predictions:

Zero-shot InternVideo:	Step 23 ✗	Fine-tuned InternVideo:	Step 25 ✗
LLaVA-Next-Video:	Step 14 ✗	VideoLLaMA2:	Step 18 ✗
LongVU:	Step 26 ✗	Qwen2.5-VL:	Step 23 ✗
InternVL3:	Step 2 ✗	Caption-enhanced LLM:	Step 6 ✗
GPT-4o:	Step 16 ✗	GPT-5:	Step 23 ✗

Step 18



Caption:

Chris & Bob move **battery** to EP.

Transcript:

Next, the **battery** in slot two is removed and translated over to the pallet.

Step 23



Caption:



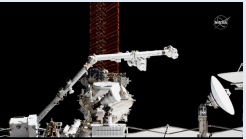


Bob **installs** adapter plate in open slot 2.

Transcript:

Bob translates back to the truss to **install** the adapter plate in open slot two.

Figure A17. In this example of step recognition, the video clearly shows an astronaut tightening the screws on a white battery. However, all the models fail to recognize it.

Spacewalk video clip

0 min 6 min 18 min 42 min 60 min

Question: Which of the following tasks happens after the astronaut leaves from the robotic arm?






Choices: (A) Robotic arm takes Luca to AMS
(B) Luca & Drew connect power and data cables
(C) Drew move to ELC 2
(D) Drew hands Luca the pump system

Predictions:

LLaVA-Next-Video:	(B) ✓	VideoLLaMA2:	(A) ✗	LongVU:	(A) ✗
Qwen2.5-VL:	(A) ✗	InternVL3:	(D) ✗	GPT-4o:	(B) ✓
GPT-5:	(C) ✗	Caption-enhanced LLM:	(B) ✓	GPT-4o-Oracle:	(B) ✓

Figure A18. In this example of spacewalk question answering, the astronaut is on the robotic arm in the first 18 minutes. After that, as shown in the 42nd minute, they are connecting cables to an equipment. So the answer to the question is “(B) Luca & Drew connect power and data cables”.

Spacewalk video clip

0 min 6 min 12 min 42 min 60 min

Question: Which of the following tasks happens after Luca hands off the debris shield to Drew?

Choices: (A) Luca and Drew install handrails
(B) Luca picks up WIF extender
(C) Arm wraps around the backside of AMS and brings Luca to worksite
(D) Luca and Drew remove debris shield

Predictions:

LLaVA-Next-Video:	(C) ✗	VideoLLaMA2:	(D) ✗	LongVU:	(D) ✗
Qwen2.5-VL:	(D) ✗	InternVL3:	(D) ✗	GPT-4o:	(C) ✗
GPT-5:	(A) ✓	Caption-enhanced LLM:	(D) ✗	GPT-4o-Oracle:	(A) ✓

Figure A19. In this example of spacewalk question answering, the astronaut first remove and jettison the debris shield and then install handrails. They remove the debris shield in the 0th minute, hand it off in the 6th minute, and jettison it in the 12th minute. So the only option happened after they hand off the debris shield is “(A) Luca and Drew install handrails”. Only GPT-5 and the GPT-4o oracle model with spacewalk mission animation correctly answer the question.

Spacewalk video clip					
					
0 min	18 min	36 min	48 min	60 min	
Question:	In which part of the video does the task that EV1 & EV2 install respective bags on worksites happen?				
Choices:	(A) The task does not happen in the video (B) The first third of the video (C) The last third of the video (D) The middle third of the video				
Predictions:					
LLaVA-Next-Video:	(B) ✗	VideoLLaMA2:	(D) ✗	LongVU:	(C) ✓
Qwen2.5-VL:	(C) ✓	InternVL3:	(D) ✗	GPT-4o:	(D) ✗
GPT-5:	(D) ✗	Caption-enhanced LLM:	(C) ✓	GPT-4o-Oracle:	(D) ✗

Figure A20. In this example of spacewalk question answering, the astronauts exit the airlock in the first third of the video, move along the handrails in the middle third, and finally install the respective bags in the last third. Therefore, the correct answer is “(C) The last third of the video”. Only LongVU, Qwen2.5-VL, and the caption-enhanced LLM give the correct answer to the question.

```

Candidate 1 video clip:
Eight uniformly sampled video frames: <candidate 1
video frame 1> ... <candidate 1 video frame 8>
Video speech transcript: <candidate 1 video
transcript>
Candidate 2 video clip:
Eight uniformly sampled video frames: <candidate 2
video frame 1> ... <candidate 2 video frame 8>
Video speech transcript: <candidate 2 video
transcript>

```

Because this task has three videos as input, we only evaluate VLLMs that can process interleaved videos and texts.

F.2.2. Intra-video Retrieval by Contrastive VLMs

Zero-shot. In this task, we match a query spacewalk clip centered at timestamp t_q with one of two candidate clips centered at t_{c1} and t_{c2} respectively. In contrast to step recognition, we **only use video features** and ignore the transcripts. This proved the best in Section F.3. After calculating the similarities between the query and the candidates, we pick the candidate with higher score as the prediction.

Last-layer Fine-tuning. We freeze the models and train two linear layers to align the query and candidate features. One is $G_\theta^q(\cdot)$ for queries and the other is $G_\theta^c(\cdot)$ for candidates. We optimize cross entropy loss during training. Denoting the query feature as f^q and candidate features as f_1^c and f_2^c , it is

$$\mathcal{L} = -\log \frac{\exp(G_\theta^q(f^q) \cdot G_\theta^c(f_y^c))}{\sum_{1 \leq i \leq 2} \exp(G_\theta^q(f^q) \cdot G_\theta^c(f_i^c))}. \quad (7)$$

After training, the candidate with higher similarity to the query, $G_\theta^q(f^q) \cdot G_\theta^c(f_i^c)$, is pick as the prediction.

We use the same hyperparameters as the last-layer fine-tuning for the step recognition task in Appendix C.3.

All-layer Fine-tuning. We fine-tune the backbone of the pre-trained models using the same training objective as that of the last-layer fine-tuning. The models learn to extract features that are critical to distinguish spacewalk steps. Specifically, we fine-tune EgoVLP on the intra-video retrieval task using AdamW [22] optimizer with a learning rate of $1e-6$. The model is trained for 1 epoch with a batch size of 32.

F.2.3. Unsupervised Segmentation Method

TW-FINCH [45] is a clustering-based unsupervised action segmentation method. We employ it to solve the intra-video retrieval task as a baseline. For a spacewalk recording in the test set, we first extract InternVideo features with a sliding window. The window length is 12 seconds and the step is 1 second. We sample video frames at 1 FPS to feed the video encoder, *i.e.*, 12 frames per video clip. Based on these clip features, TW-FINCH clusters the entire spacewalk recording into N_{seg} segments, each of which is regarded as a step. For a retrieval query timestamp t_q with two candidate timestamps t_{c1} and t_{c2} , if exactly one of the candidates falls in the same segment as t_q , we regard it as our prediction. Otherwise, we randomly pick one of them as the prediction. We find that the best configuration is $N_{\text{seg}} = 100$ with only video features.

Method	Accuracy					
	(w=)10s	20s	30s	1min	2min	3min
Random	50.00	50.00	50.00	50.00	50.00	50.00
TW-FINCH*	-	-	-	-	-	64.10
Zero-shot						
EgoVLP	69.04	69.95	70.92	65.76	60.20	56.43
VideoCLIP	66.91	67.48	67.88	63.52	50.72	52.07
InternVideo	67.60	69.12	69.82	66.76	56.46	56.45
GPT-4o	71.00	71.00	70.83	58.50	53.08	56.67
Last-layer Fine-tuning						
EgoVLP	68.63	69.91	70.12	66.51	57.46	52.36
VideoCLIP	61.15	60.53	61.29	59.00	54.08	52.23
InternVideo	65.11	65.33	64.97	62.62	56.02	52.24
All-layer Fine-tuning						
EgoVLP	68.78	70.18	70.93	69.79	62.58	58.80

Table A14. Model performances on the intra-video retrieval task. GPT-4o achieves the highest accuracy while EgoVLP is comparative. *: TW-FINCH have unlimited context.

F.3. Evaluation Results

Table A14 shows the model performances on the intra-video retrieval task. GPT-4o achieves the best performance of 71%, while the 70.93% accuracy of EgoVLP is competitive. Different from the step recognition task, last-layer fine-tuning cannot improve the model performance in most cases. All-layer fine-tuning only increases the model accuracy very slightly. Besides, while the context window length expands, the model performance first increases and then degrades after a peak. This observation is in line with the step recognition results, which reveals their incapability of digesting long video contexts.

Table A15 ablates the modalities on intra-video retrieval in zero-shot setting. The video-only approach achieves the highest accuracy across all models, while the text-only one performs poorly. This is different from step recognition where videos contribute less than texts. Because the query and candidates here are the same modality with the same visual distribution, whereas we need to match the real spacewalk videos with their animations in the step recognition task. Moreover, combining video and text does not improve the performance, indicating that transcript is effectively noise for the current method on this task. This motivates the future development of stronger models that can utilize the text signal on this task to facilitate unsupervised segmentation of spacewalk videos.

F.4. Qualitative Results

On the intra-video retrieval task, we show two qualitative examples with context window length $w = 30$ s in Figures A21 and A22. We find that the models tend to retrieve candidate videos naively based off the visual scene appearance. This helps them do well in simple scenarios (Figure A21). However, there are more difficult test samples.

Method	Accuracy					
	$w = 20$ s			$w = 30$ s		
	V	T	V+T	V	T	V+T
EgoVLP	69.95	48.73	63.71	70.92	47.85	64.95
VideoCLIP	67.48	51.65	56.56	67.88	50.04	57.14
InternVideo	69.12	51.55	63.15	69.82	50.23	63.08
GPT-4o	71.00	52.42	69.25	70.83	53.67	69.92

Table A15. Ablation about input modality on intra-video retrieval. The models are evaluated in zero-shot scenario. V: video; T: text (captions and transcripts).

In Figure A22, the query and candidate 1 show the same step but from different views. All the models fail in this example. Better spacewalk video understanding capability is necessary for future models to solve it.

Query



t - 15s

t - 7.5s

t

t + 7.5s

t + 15s


Step label:

Step 13 - EV1 install upper triangle on mast canister

Transcript:

Now it's up. And you'll engage the left side first and then pivot the right side to engage the soft dock. Mike, you'll be working on stowing your short PGT on handrail five three seven two. That's near the APFR wrenches. Handbrake embedded in the soft dock.

Candidate 1




Step label:

Step 13 - EV1 install upper triangle on mast canister

Transcript:

Great views from the high definition camera view of Kate Rubin's helmet camera. Alpha two, clockwise two. She's confirming the pistol grip settings before she goes ahead and starts driving some of the bolt that you see on the center pad there that's currently soft dock to the mass canister. Driving those bolts will secure it in place. Through mic eight to torque. We anticipate eleven and a half turns. Copy. Eleven and a half torque.

Candidate 2



Step label:

Step 12 - EV2 hand upper triangle to EV1

Transcript:

Yeah. We want the RET back in the bag. Sorry. You you don't have to release it from the mounting bracket. Okay. Either you can relocate or you can hook it back this way, and I'll be No. It'll text me. Okay. I'm sending it back to a queue. It's still of control. Again, Rubens has control of that upper bracket. They're just making sure that the tether configuration is okay before she begins the work of installing it to the mass canister.

Predictions:


Zero-shot EgoVLP: Candidate 1 ✓

Fine-tuned EgoVLP: Candidate 1 ✓

GPT-4o: Candidate 1 ✓

Figure A21. In this example, both the query video and candidate 1 show an astronaut installing an upper triangle. They are in the same scene while candidate 2 is different. All the models make the correct predictions.

Figure 10 displays a sequence of five satellite images showing the movement of a white object (likely a space station module) relative to a larger structure (likely a space station) over time. The images are labeled $t - 15s$, $t - 7.5s$, t , $t + 7.5s$, and $t + 15s$, showing the object moving from left to right.



Step label:

Step 12 - EV2 hand upper triangle to EV1

Transcript:

And I have my RET on the right. K. Let me know when you already take control. I already take control. And you guys I I got control. Alright. Glover just handed off the upper triangle to Kate Rubins. Okay. And, like, on case go, you can release the bad RET. And I have control. I have it ready to go. The back. RET. He wants the RET to go back to the back or to the mounting bracket.

Candidate 1



Step label:

Step 12 - EV2 hand upper triangle to EV1

Transcript:

Like, looks like you're doing a great job on MOI. I think they're happy with that config, and you can pass the triangle up to Kate when you're ready. Works.

Candidate 2



Step label:

Step 13 - EV1 install upper triangle on mast canister

Transcript:

And you'll engage the left side first and then pivot the right side to engage the soft dock. Mike, you'll be working on stowing your short PGT on handrail five three seven two. That's near the APFR wrenches. Handbrake embedded in the soft dock. Copy. Soft dock. Thank you. K. On your long PGT, your settings will be alpha two clockwise two.

Predictions:

Zero-shot EgoVLP: Candidate 2 ✖
GPT-4o: Candidate 2 ✖

Fine-tuned EgoVLP: Candidate 2 ✖

Figure A22. In this example, all three videos are visually dissimilar. The query video shows an astronaut handing an upper triangle to the other from a third-person view, while candidate 1 shows it from a first-person view. In candidate 2, they are installing the upper triangle. However, all the models fail to recognize candidate 1 as the same step as the query.