

These supplemental materials encompass details of investigation of user search behavior feedback (§A), more related works (§B), the details of the dataset (§C), the details of the baselines (§E), more visualization quality (§D), response time of video retrieval (§F), limitations (§G), and quantitative results (§H).

A DETAILS OF INVESTIGATION OF USER SEARCH BEHAVIOR FEEDBACK

To verify users’ search needs for interactive retrieval, we design a user search behavior feedback survey. It aims to gain in-depth insights into users’ real experiences during the search process, identifying potential issues and areas for improvement. The survey covers several aspects: user demographics (such as age, gender, occupation), video search habits, video search usage (whether users prefer to directly search for video moment), evaluation of search effectiveness (including the accuracy and usability of current video search systems, user satisfaction, and expectations for conversational search tools), and issues and improvement suggestions (open-ended questions for users to share problems and suggestions when using interactive search). To better accommodate the habits and preferences of users from different countries, we implemented a diversified survey approach. On the Amazon Mechanical Turk³ platform, we conducted an online survey specifically designed to target an international audience, successfully engaging 500 participants from various regions. In China, we utilized the Wenjuanxing⁴ platform to accurately reach and collect feedback from 500 Chinese users. Additionally, to gather more comprehensive and in-depth data, we organized offline paper surveys. Over the weekends, we collected insights from another 500 participants in high-traffic areas such as university campuses, shopping malls, parks, popular tourist destinations, and subway stations.

B MORE RELATED WORK

B.1 CROSS-MODAL VIDEO RETRIEVAL

Video Retrieval. Since TRECVID (Smeaton et al., 2006), the task of cross-modal video retrieval based on text queries has undergone a long development process. The goal of this task is to retrieve relevant videos from a set of video candidates given a text query. In this section, we focus on methods using the deep learning-based paradigm (Chen et al., 2020; Dong et al., 2019; Song et al., 2021; Luo et al., 2022). The basic idea is to encode texts and videos into embeddings, and then learn a common embedding space to do the matching between them. For text encoding, BERT-based models (Devlin et al., 2018) have become the mainstream. While for the video encoding in T2VR, a typical solution has recently evolved from Convolutional Neural Networks (CNN) (e.g., I3D (Carreira & Zisserman, 2017) and 3D ResNets (Hara et al., 2018)) to visual transformers (ViT) (e.g., ViT of CLIP (Radford et al., 2021) and TimeSformer (Bertasius et al., 2021)). We can divide related methods into three categories. 1) *Pre-extracted multi-modality fusion* methods (Mithun et al., 2018; Liu et al., 2019; Gabeur et al., 2020; Wang et al., 2021) integrate rich multi-modality information (eg., motion, audio, and face) to improve the performance of T2VR. 2) *Joint text-video pre-training* methods (Wang et al., 2022a; Yan et al., 2023; Li et al., 2022; Ge et al., 2022) train the model with raw video and paired text in an end-to-end manner. 3) *Pre-trained CLIP-based* methods (Luo et al., 2022; Fang et al., 2021; Gorti et al., 2022; Liu et al., 2022b) use CLIP as a text-video backbone and adapt it to T2VR tasks. However, the pre-extracted multi-modality fusion methods are limited by the pre-extracted single-modal features, since these features are not particularly learnt for the target downstream tasks. The joint text-video pre-training methods achieve marginal improvements using a joint text-video pre-training paradigm due to the lack of large-scale text-video datasets. The pre-trained CLIP-based methods largely benefit from the pre-learned vision-text associations inherited from CLIP and rely on naive mean-pooling or text-conditioned pooling to aggregate visual features.

Video Moment Retrieval. As an extension of video retrieval, video moment retrieval task aims to identify specific clips or moments within a video based on a given textual query (Gao et al., 2017; He et al., 2019). Pioneering works have explored various technical avenues, including attention-based retrieval, reinforcement learning, visual-language pretraining, among others. Liu et al. maintain

³<https://www.mturk.com/>

⁴<https://www.wjx.cn/>

a focus on retrieval paradigm, incorporating window segmentation as a preprocessing step, then retrieving moments. (He et al., 2019) devised an agent to pinpoint the start and end timestamps of moments based on reinforcement learning. Following this, some researchers expanded RL-based methods into the spatiotemporal or semantic domains (Cao et al., 2020). Moreover, leveraging pretraining techniques (Zeng et al., 2021; Pan & Zeng, 2023), several prompt-based models have emerged (Zeng, 2022; Liu et al., 2022a), facilitating timestamp prediction through regression.

B.2 INTERACTIVE RETRIEVAL

The concept of interactive retrieval has long been proposed in the context of combining human-machine learning techniques for multimedia content search. Some studies (Thomee & Lew, 2012; Snoek et al., 2008) have demonstrated that interactive retrieval can significantly improve search performance by enabling users to review search results and refine queries. With the significant progress of deep learning technology in the field of cross-modal video retrieval, interactive video retrieval has re-attracted the attention of researchers. Currently, only a few works (Madasu et al., 2022; Maeoki et al., 2020; Ma & Ngo, 2022; Liang & Albanie, 2023) have explored this task. For example, Madasu et al. (Madasu et al., 2022) and Maeoki et al. (Maeoki et al., 2020) adopt a dialogue-based approach, utilizing a series of video-related questions and answers generated by different models as retrieval queries. Meanwhile, addressing the issue that the above methods did not directly involve video question answering, Liang et al. (Liang & Albanie, 2023) employed a video question answering model to generate question-and-answer information in order to improve the retrieval accuracy. Furthermore, Ma et al. (Ma & Ngo, 2022) develop a user simulation for intelligent multimedia applications, leveraging advanced techniques in multimedia content analysis, including concept detection and cross-modal embedding, to enable precise video segment search through human-computer interaction. The aforementioned works are limited to achieving single-task interactive retrieval through methods such as reconstructing retrieval text, using visual question answering models, and simulating users. The technical challenges in modeling multi-turn dialogue retrieval have contributed to the slow development in this direction.

C DETAILS OF DATASET

C.1 DATASET INSTANCE

The instances in our dataset consist of four fields: (1) id: a unique identifier generated using the video name and a random number; (2) type: retrieval intent categorized into seven types—0: chat intent, 1: video retrieval intent, 2: video moment retrieval intent, 3: video to video moment retrieval intent, 4: video moment to video retrieval intent, 5: abstract search intent, 6: analogous search intent, 7: context-independent intent; (3) split: the dataset is divided into training, testing, and validation sets; (4) conversations: multi-turn retrieval formats where "from: human" indicates a query from a human, with the corresponding "value" and "Chinese_value" representing the query content in English and Chinese, respectively. "From: gpt" indicates feedback from GPT, with "gt" representing the retrieved video or video moment—note that "gt_se: [-1,1]" indicates video retrieval, while other values indicate video moment retrieval. Additionally, "video_source" indicates the dataset from which the video is sourced, with the corresponding value providing an interpretable description. Figure 8 shows an example in a unified format. These clearly defined fields allow benchmark users to flexibly construct the necessary training instances and easily evaluate the model.

C.2 DIVERSITY QUALITY

We conducted an analysis of our video sources, the different types of videos (Figure 9), and performed a frequency analysis of annotated sentences (Figure 10 and 11) to ensure a comprehensive diversity. In addition, as shown in Figure 12, we present the statistics for video retrieval cases.

Table 5: Comparing InterLLaVA performance across different data distributions

N training samples	Category	R@1 \uparrow	R@1 IoU=0.5 \uparrow	R@1 IoU=0.7 \uparrow
2K	Movies	21.36	6.21	2.24
20K	TV shows	26.73	8.24	3.98
1K	ALL	36.6	8.49	4.94
6K	ALL	45.7	11.31	5.8
20K	ALL	47.96	11.52	6.45
60K	ALL	54.86	12.28	7.13

C.3 TRAINING DATA DISTRIBUTION ANALYSIS.

Our InterLLaVA model is trained on four different data modes, with the number of training samples ranging from 1K to 60K. Table 5 summarizes the performance evaluation results for all training samples. We observed the following points: 1) In low-sample scenarios, particularly when the sample size is less than 6K, InterLLaVA’s accuracy is significantly limited (36.6 vs 54.86 for R@1), showing much lower performance compared to conditions with larger sample sizes. 2) We further explored training the model using samples from a single category (e.g., movies). The experimental results indicate that compared to training on data of the same scale but with more diverse categories, InterLLaVA’s video and moment retrieval performance decreased by 21.23% and 3.28% in R@1 and R@1 IoU=0.5, respectively. This result aligns with expectations, as training on more diverse categories allows the model to capture richer features and enhance its generalization ability.

D MORE VISUALIZATION QUALITY

We present more examples from our IVCR-200K dataset in Figures 13-21.

E DETAILS OF BASELINES

For video retrieval, we selected the following five state-of-the-art models as benchmarks. We adopt their original setup, using both video and text as model inputs for the video retrieval task. CLIP4Clip(Luo et al., 2022) uses CLIP to extract the frame features and the text features, and then uses the mean pooling to aggregate the feature of all frames for video representation. X-Pool(Gorti et al., 2022) adopts text-conditioned pooling to aggregate visual features. TS2-Net(Liu et al., 2022b) proposes different token shift operations in ViT to learn short-term temporal dependencies across locally adjacent frames. T-MASS(Wang et al., 2024) proposes a stochastic modeling approach to achieve expressive and flexible text embeddings, enhancing the alignment of text and video semantics in the joint space. Furthermore, we reimplement a video retrieval model named BLIP-2, utilizing the video and text features encoded by BLIP-2(Li et al., 2023a), with X-Pool(Gorti et al., 2022) serving as the base model.

For video moment retrieval, we selected six methods as benchmarks. We utilize BLIP-2(Li et al., 2023a) as the encoder to extract video and text features, which are then used as inputs for all models in video moment retrieval. 2D-TAN(Zhang et al., 2020) proposes a novel two-dimensional temporal matrix for moment localization. MMN(Wang et al., 2022b) introduces a mutual matching network that directly models the similarity between language queries and video moments within a joint embedding space. UMT(Liu et al., 2022a) proposes a unified framework for solving joint moment retrieval and highlight detection. CG-DETR(Moon et al., 2023) explores the provision of cues for query-associated video clips within cross-modal attention. MomentDiff(Li et al., 2024a) utilizes diffusion models to diffuse real span to random noise, and then learns to denoise the random noise back to the original span under the guidance of text and video similarity. Moreover, we chose a model based on multi-modal large language models as additional benchmarks. TimeChat(Ren et al., 2023)

Table 6: Comparing average response times of the model across different test data scales

N Testing Samples	Average Response Time of Video Retrieval(s)	Average Response Time of MLLM (s)	Total Average Response Time (s)
1008	0.06	0.67	0.73
1534	0.06	0.59	0.65
1878	0.06	0.73	0.79
2436	0.06	0.8	0.86

proposes a time-sensitive multimodal large language model for long video understanding and precise temporal localization.

F RESPONSE TIME OF VIDEO RETRIEVAL

We compared the average response times of video retrieval and Multi-Modal Large Language Model (MLLM) inference, and the results are summarized in the Table 6. Our observations are as follows: 1) The average response time for video retrieval is notably lower compared to the MLLM inference. However, the average response time of the MLLM remains below 1 second, which is acceptable. 2) Testing with different numbers of video samples revealed a slight increase in the average response time for video retrieval. Meanwhile, the average response time of the MLLM remains almost constant, demonstrating that our model does not introduce significant delays even when handling larger volumes of video data.

G LIMITATIONS

The IVCR-200K dataset is constrained by the depth of manual annotation and the diversity of real-world data types. It needs to be expanded to cover a wider array of interactive retrieval scenarios, including complex analogy searches, diverse contextual searches, and fine-grained interactive search requirements. Additionally, the current model does not achieve seamless integration of video retrieval and moment retrieval into a unified, efficient end-to-end system. There is considerable potential for improvement in areas such as temporal video modeling, accurate capture of user retrieval intent, and the natural and fluid execution of multi-round dialogues.

H QUANTITATIVE RESULTS

Figures 22-25 presents a qualitative comparison between InterLLaVA and other video large language models. Our observations are as follows: 1) Video-LLaVA (Lin et al., 2023) has limitations in handling video retrieval, as it is limited to describing the direct content relationship between video and text. It lacks intent analysis for text retrieval and cannot provide interpretable feedback for identifying relevant videos and moment. 2) In contrast, TimeChat (Ren et al., 2023) merely offers the start and end timestamps of videos, lacking any form of feedback on the retrieval outcomes, thus insufficient to cater to users' personalized search demands. 3) InterLLaVA excels at accurately matching the desired retrieval video, precisely locating specific moments within videos, and providing interpretable feedback. However, there is still potential to further enhance the richness and detail of the feedback generated by our model.


```

972 {
973   "id": "v_n9eNF1WqFu",
974   "type": [1,2,4], #video-moment-similar
975   "split": "train",
976   "conversations": [
977     {
978       "from": "human",
979       "value": "Person in white clothes are dancing.",
980       "Chinese_value": "穿白色衣服的人在跳舞。",
981       "text_id": "247775f9"
982     },
983     {
984       "from": "gpt",
985       "gt": ["s234d9"],
986       "gt_se": [-1, -1],
987       "video_source": "ActivityNet",
988       "value": "The intent of this query is to find a complete video.",
989       "Chinese_value": "这个查询的意图是找到一个完整的视频。"
990     },
991   ],
992   {
993     "from": "human",
994     "value": "Locate a video moment of the person spinning in circles.",
995     "Chinese_value": "定位视频中该人转圈的瞬间。",
996     "text_id": "0c1dce9c"
997   },
998   {
999     "from": "gpt",
1000     "gt": ["s234d9"],
1001     "gt_se": [1.3, 5.6],
1002     "video_source": "ActivityNet",
1003     "value": "The intent of this query is to further locate a spinning
1004             moment.",
1005     "Chinese_value": "这个查询的意图是进一步定位转圈的瞬间。"
1006   },
1007   {
1008     "from": "human",
1009     "value": "Is there any similar moment?",
1010     "Chinese_value": "有没有类似的瞬间?",
1011     "text_id": "5553fa8a"
1012   },
1013   {
1014     "from": "gpt",
1015     "gt": ["45866g"],
1016     "gt_se": [-1, -1],
1017     "video_source": "MSR-VTT",
1018     "value": "The 'similar' in the query may refer to videos of dancing
1019             in circles.",
1020     "Chinese_value": "查询中的类似可能指的是在圆圈中跳舞的视频。"
1021   }
1022 ]
1023 }
1024
1025

```

Figure 8: Example JSON of our IVCr-200K dataset.

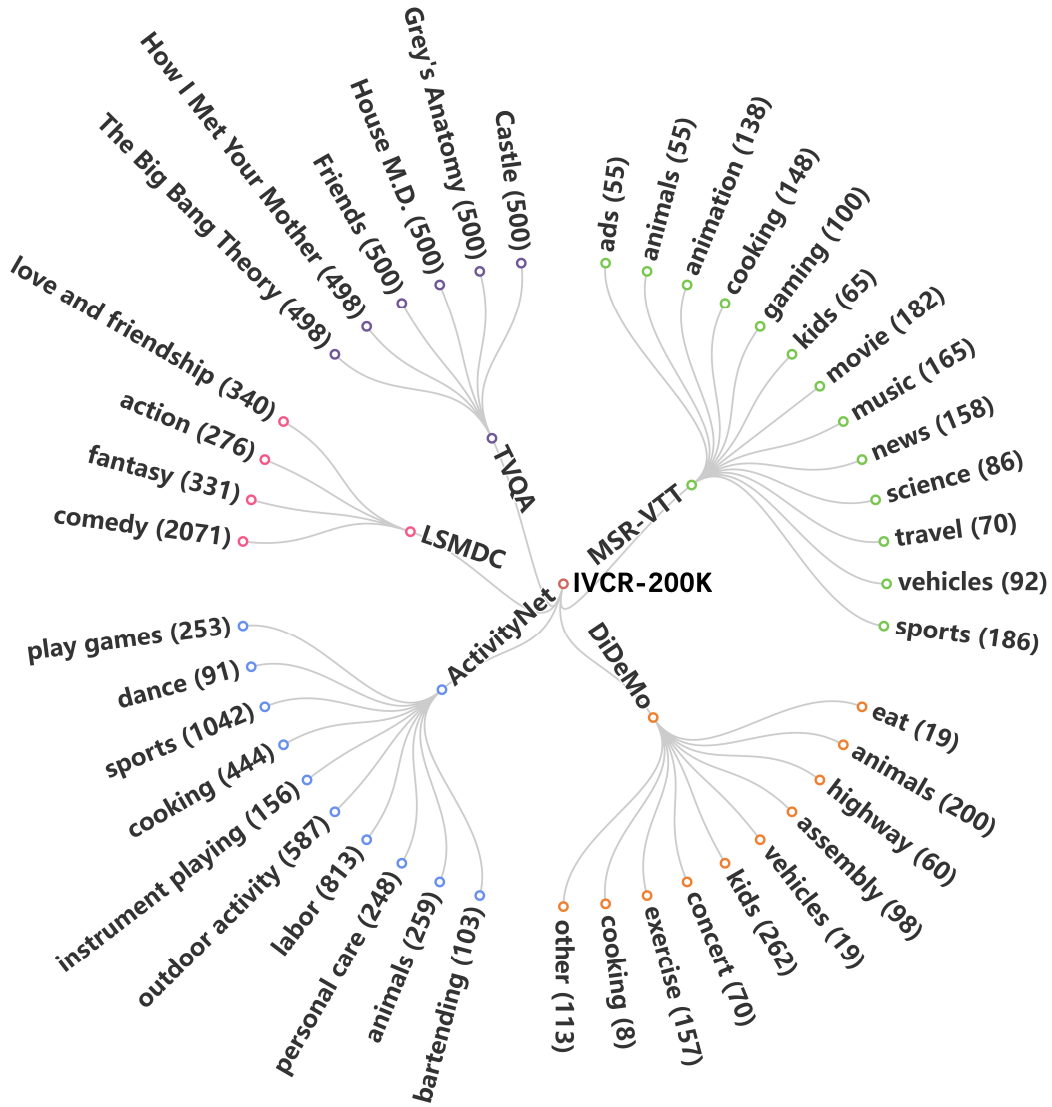


Figure 9: The hierarchical structure of the sources and categories of IVCR-200K.

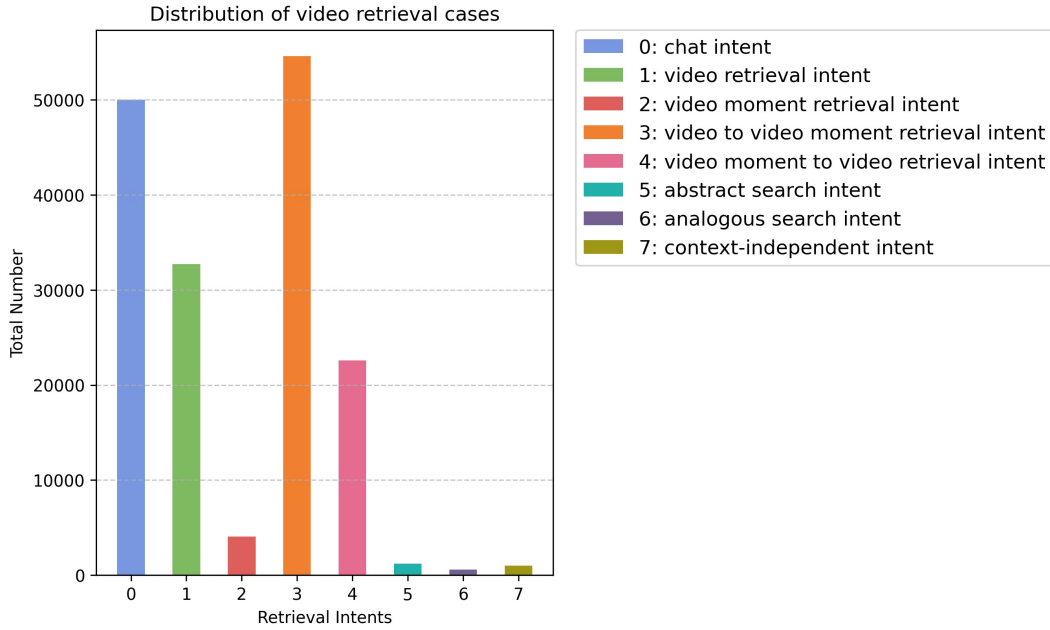


Figure 12: Distribution of video retrieval cases.

User: can you suggest some videos to understand Monte Carlo simulation of golf shots

InterLLaVA: Sure, here are some videos that may be helpful in understanding Monte Carlo simulation of golf shots: 1. "Simulating a Golf Shot" by Aaron Parsons: This video provides a simple introduction to Monte Carlo simulation of golf shots using Excel. 2. "Golf Simulation in Excel with Monte Carlo Method" by Peter Rasmussen: This video provides a more advanced tutorial on simulating golf shots using Monte Carlo methods in Excel. The video covers topics such as random number generation, probability distributions, and sensitivity analysis. These videos should provide a good starting point for understanding Monte Carlo simulation of golf shots. I hope you find them helpful!

User: In a humorous sitcom, a couple is mocked by the airport staff for being late. They explain their reasons and then sit down helplessly to wait.

InterLLaVA: The intent is abstract search. The retrieval question is related to a humorous sitcom. The retrieved video is about Robin and Ted, who are also late for their flight and are sarcastically commented on by an airline attendant. They explain their reasons for the delay, and the attendant makes a phone call while they take a seat to wait at the airport. The retrieval question matches the retrieved video because it is a common scenario in airport settings, and the video captures a similar situation.

User: can you suggest videos to understand Optimization of golf shot trajectory

InterLLaVA: Sure! Here are some videos that can help you understand the optimization of golf shot trajectory: 1. "Optimization of Golf Shots using MATLAB" by MathWorks: This video provides an introduction to using optimization algorithms in MATLAB to optimize golf shot trajectory. It covers topics such as setting up the optimization problem, defining the objective function, and specifying the constraints. 2. "Optimization of Golf Shots with Toptracer" by Toptracer: This video provides an overview of how the Toptracer technology is used to optimize golf shot trajectory. It covers topics such as club selection, ball spin rate, and launch angle optimization. These videos cover different aspects of the optimization of golf shot trajectory.. I hope you find them helpful!

User: In a house, a group of people are fighting fiercely

InterLLaVA: The intent is abstract search. The logical relationship between the retrieval question and the retrieved video is that the question is a general description of the content, while the video provides a specific instance of that content. The video captures a moment of intense conflict between people in a domestic setting, which aligns with the abstract search intent of finding content related to violence or conflict.

Figure 13: An example of abstract concept retrieval and dialogue from the IVCR-200K dataset.

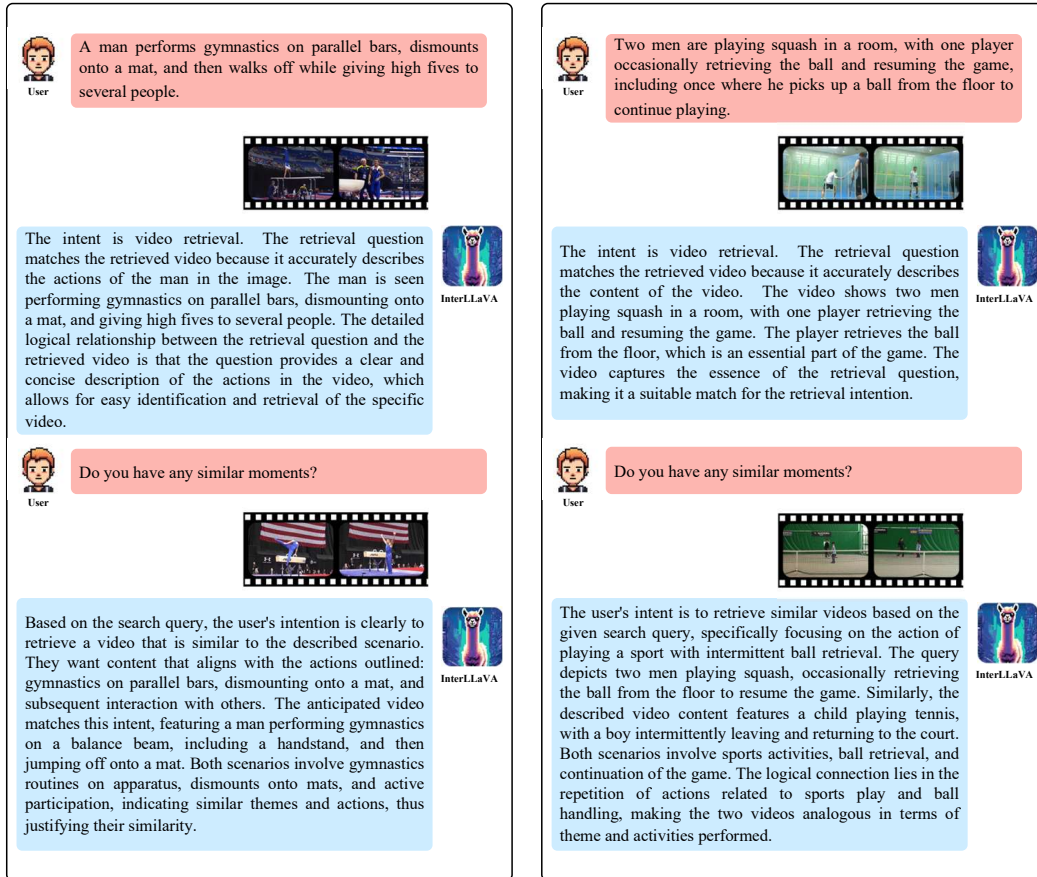


Figure 14: An example of analogy retrieval from the IVCr-200K dataset.

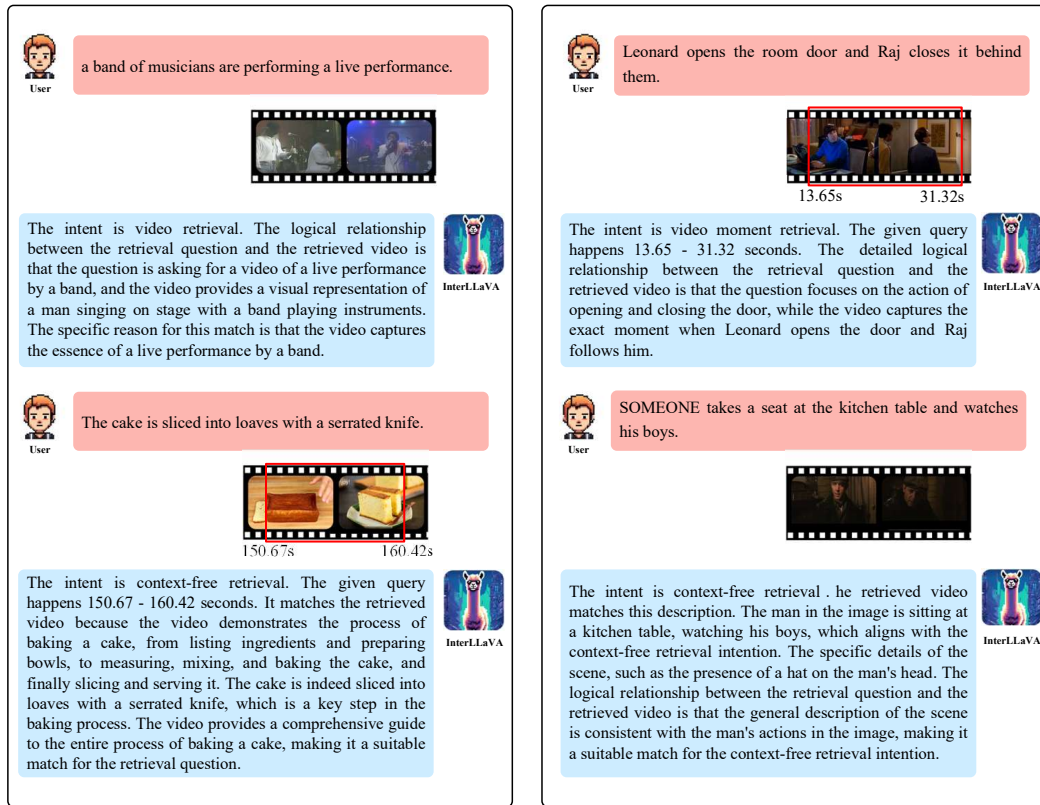


Figure 15: An example of context-independent retrieval from the IVCR-200K dataset.

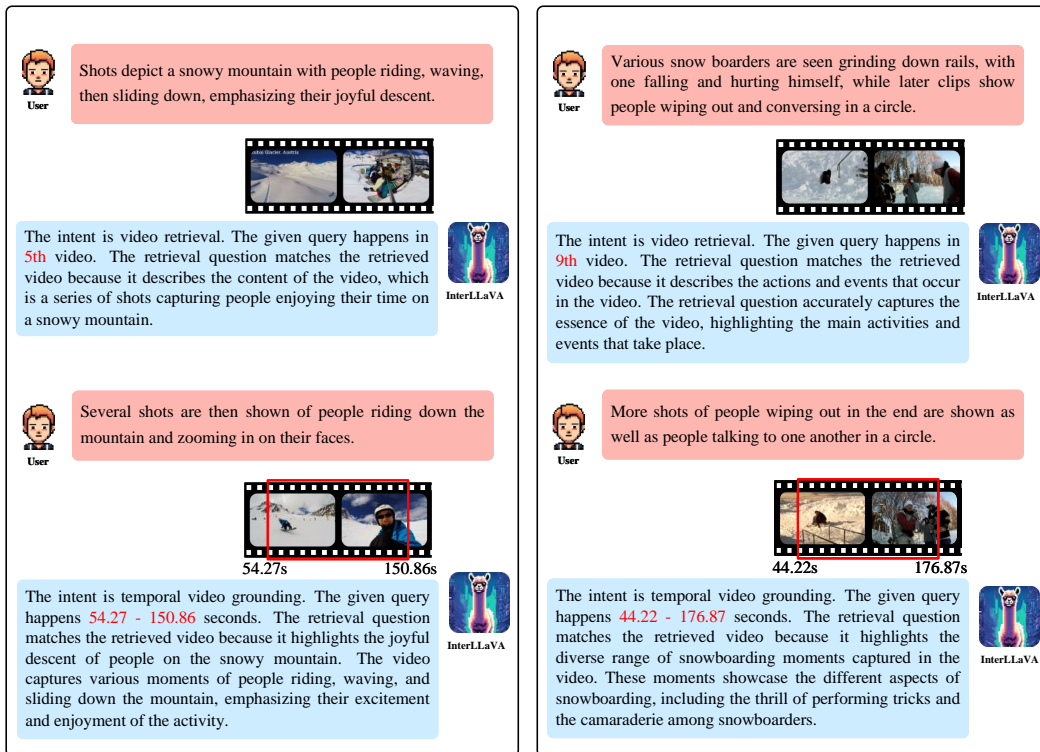


Figure 16: An example of video-to-video moment retrieval from the IVCR-200K dataset.

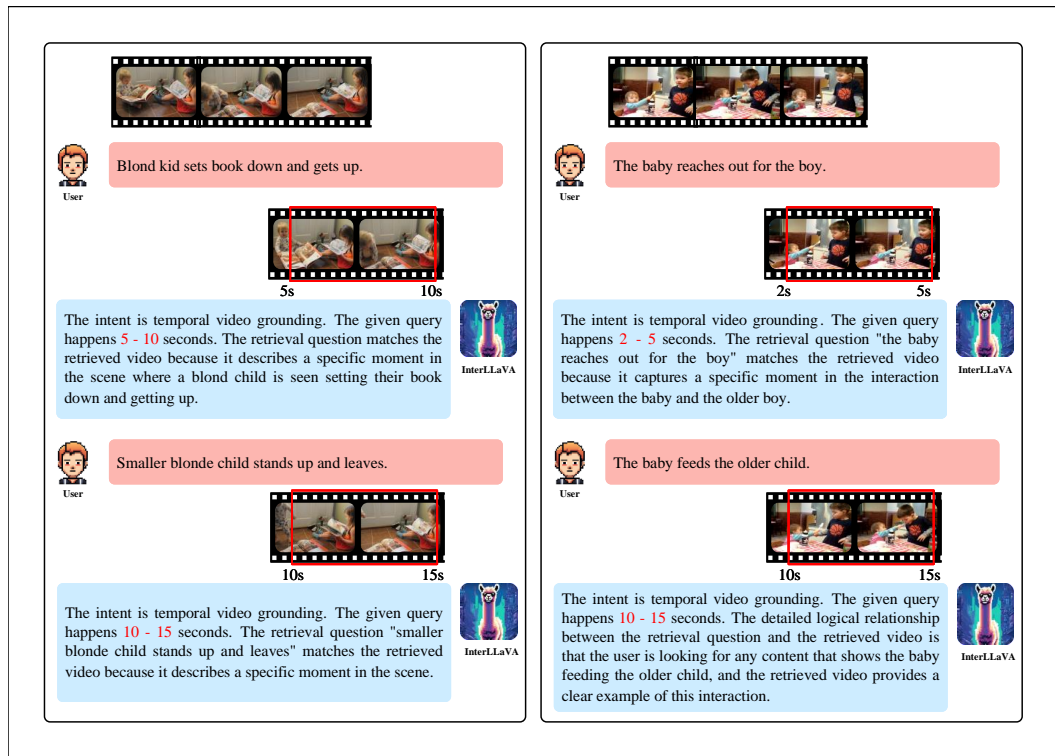





Figure 17: An example of video-to-video moment retrieval from the IVCR-200K dataset.




Phoebe blew on her coffee, sipped it, and adjusted her hair while Ross entered the apartment, walked through, and turned toward Chandler and Joey who were talking.




The intent is video retrieval. The given query happens in 4th video. The retrieval question matches the retrieved video because it describes a scene from a TV show Friends. In the video, Phoebe is blowing on her coffee, sipping it, and adjusting her hair while Ross enters the apartment, walks through, and turns towards Chandler and Joey who are talking.







Phoebe looks up from her coffee and uses her fingers to moved hair out of her face.




14.83s 16.54s

The intent is temporal video grounding. The given query happens 14.83 - 16.54 seconds. The retrieval question matches the retrieved video because it captures a specific moment in the scene where Phoebe is adjusting her hair while sipping her coffee. The video shows Phoebe in the process of blowing on her coffee, sipping it, and adjusting her hair.







Ross walks through the door and closes it.




35.35s 39.92s

The intent is temporal video grounding. The given query happens 35.35 - 39.92 seconds. The retrieval question matches the retrieved video because it describes a specific moment in the scene where Ross enters the apartment and walks through the door. This description accurately captures the moment when Ross enters the room and interacts with the other characters, which aligns with the retrieval question.






Phoebe blows on her coffee and then starts to take a sip.



12.26s 21.1s

The intent is temporal video grounding. The given query happens 12.26 - 21.1 seconds. The retrieval question matches the retrieved video because it captures a specific moment in the scene where Phoebe is blowing on her coffee and preparing to take a sip.





Beckett interacts with Shay, Castle, and Gates in different settings, involving conversations, walking, and exchanging glances.



The intent is video retrieval. The given query happens in 5th video. The retrieval question matches the retrieved video because it describes the content of the video, which shows Beckett interacting with Shay, Castle, and Gates in various settings. The video captures the conversations, walking, and exchanging glances between the characters, which are all elements present in the retrieval question.





Beckett discusses with Gates while opening and closing her folder, then immediately walks out of her boss's office, as Gates stares and Castle converses from Gates' office.



46.92s 90.62s

The intent is temporal video grounding. The given query happens 46.92 - 90.62 seconds. The retrieval question matches the retrieved video because it captures a specific moment in the scene where Beckett is interacting with multiple people, including Gates and Castle. The video moment is relevant to the retrieval question as it provides context and details about the interactions between the characters in the scene.





Beckett walks out of her boss's office and Gates stares at Beckett while Castle converses from her office.



73.6s 90.62s

The intent is temporal video grounding. The given query happens 73.6 - 90.62 seconds. The retrieval question matches the retrieved video because it captures a specific moment in time where Beckett is walking out of her boss's office, Gates is staring at her, and Castle is conversing from his office.





Beckett walks out from the office of her boss immediately.



73.6s 78.2s

The intent is temporal video grounding. The given query happens 73.6 - 78.2 seconds. The retrieval question matches the retrieved video because it captures a specific moment in the scene where Beckett is walking out of her boss's office. The video also involves conversations, walking, and exchanging glances with other characters, such as Shay, Castle, and Gates, which are relevant to the retrieval question.



Figure 18: An example of video-to-video moment retrieval from the IVCR-200K dataset.

Two men in white shirts and hats freestyle wrestle indoors while spectators watch and a choir cheers them on. The fighters alternately attack, perform flips, and speed up their moves in the final round.

The intent is video retrieval. The given query happens in 10th video. The retrieval question is about two men in white shirts and hats freestyle wrestling indoors while spectators watch and a choir cheers them on. The retrieved video matches the retrieval question because it shows two men in white shirts and hats engaging in a wrestling match indoors, with spectators watching and a choir cheering them on.

In the final round, they speed up the fight with faster kicks and motion.

169.25s 179.1s

The intent is temporal video grounding. The given query happens 169.25 - 179.1 seconds. The retrieval question matches the retrieved video because it describes the action of the two men in the image. The men are seen fighting in the middle of a circle, with one man eventually getting up. This action is consistent with the final round of a fight, where the intensity increases and the participants become more aggressive.

There are several spectators watching them, some sitting on chairs and some sitting on the ground.

38.51s 84.18s

The intent is video to video moment retrieval. The given query happens 38.51 - 84.18 seconds. The scene depicts two men in white shirts and hats freestyle wrestling indoors while spectators watch and a choir cheers them on. The fighters are performing flips and speeding up their moves in the final round, which is a common sight in such events. The spectators are engaged and attentive, sitting on chairs and on the ground, which is consistent with the image.

The two men continue wrestling in a non aggressive manner by taking turns to attack the opponent.

124.47s 135.22s

The intent is video to video moment retrieval. The given query happens 124.47 - 135.22 seconds. The retrieval question matches the retrieved video because it describes the actions of the two men in the image. The spectators and choir cheering them on add to the overall atmosphere of the event. The retrieval question captures the essence of the video, highlighting the key moments of the wrestling match and the interactions between the two men.

A group of cheerleaders performed an energetic and synchronized routine, including mid-air flips, runs, and formations, culminating in a stunning triangle formation held by three rows of women.

The intent is video retrieval. The given query happens in 2th video. The retrieval question matches the retrieved video because it accurately describes the cheerleaders' performance in the video. The video showcases a group of cheerleaders executing a highly coordinated routine, with various flips, runs, and formations, ultimately resulting in a captivating triangle formation.

The girls start with mid-air flips, then cheer and flip while staying synchronized and graceful. They finish with a triangle formation held by three rows of women.

40.6s 198.07s

The intent is temporal video grounding. The given query happens 40.6 - 198.07 seconds. The retrieval question matches the retrieved video because it describes the cheerleaders' performance in the video. The video shows a group of cheerleaders performing an energetic and synchronized routine, including mid-air flips, runs, and formations, culminating in a stunning triangle formation held by three rows of women.

They perform an impressive cheer, staying perfectly synchronized and executing all moves gracefully, culminating in a stunning triangle formation with a girl held aloft by three rows of women.

96.06s 198.07s

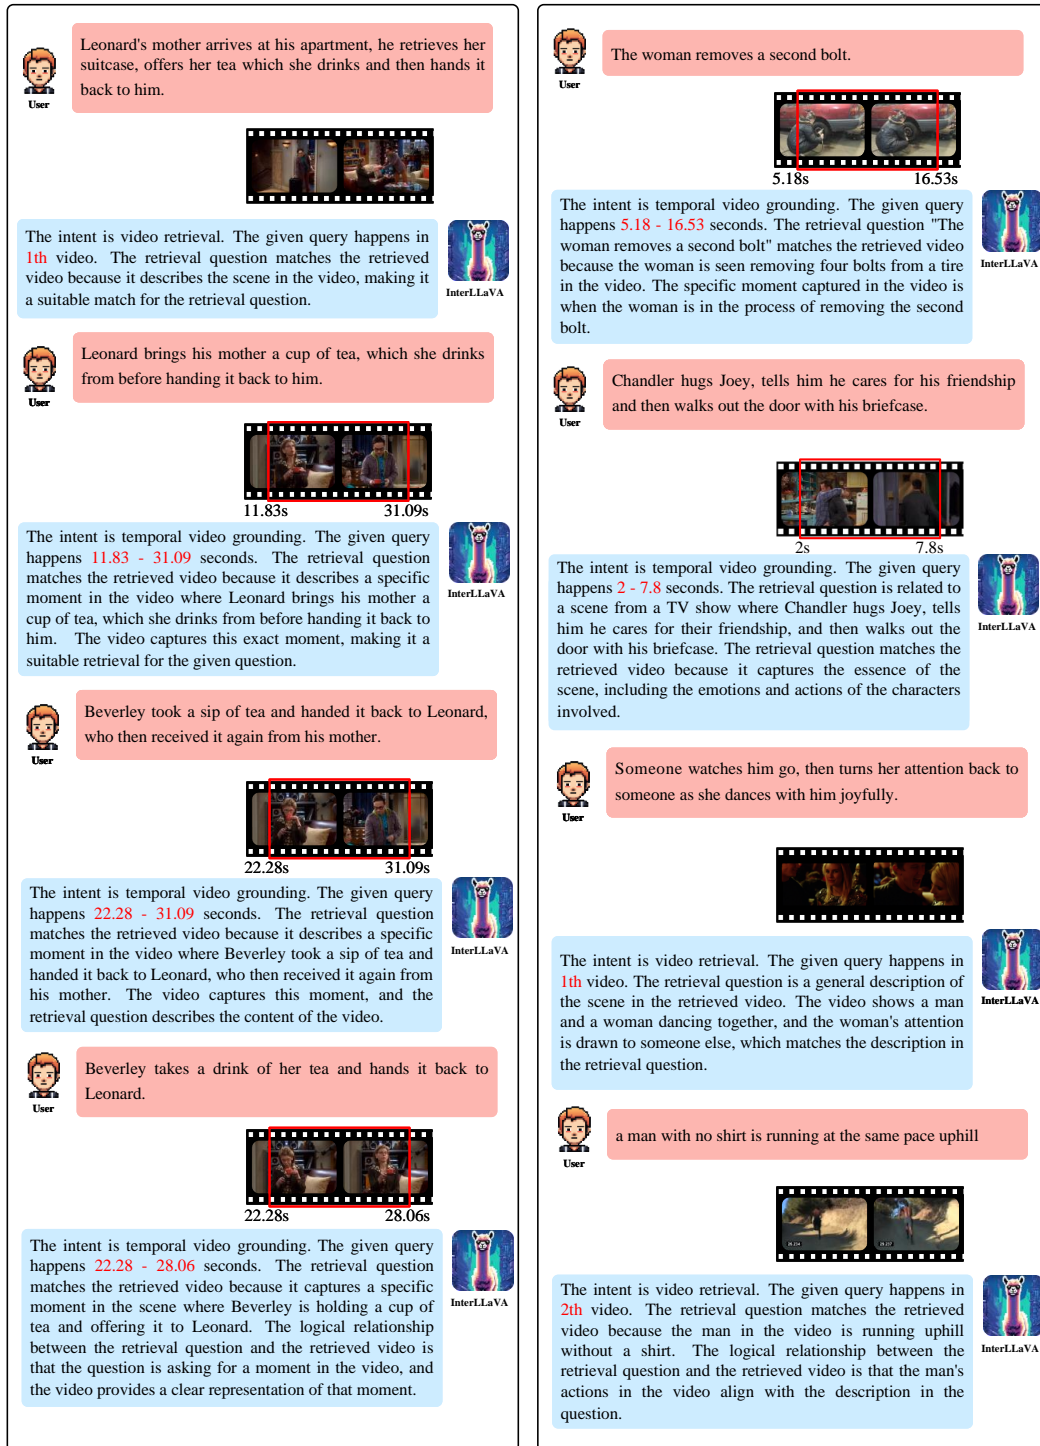
The intent is temporal video grounding. The given query happens 96.06 - 198.07 seconds. The retrieval question matches the retrieved video because it describes the cheerleading performance in detail. The video captures the cheerleaders performing an energetic and synchronized routine, including mid-air flips, runs, and formations, culminating in a stunning triangle formation held by three rows of women.

They stay in sync throughout the entire cheer, looking great and hitting all the moves with grace.

96.06s 154.49s

The intent is temporal video grounding. The given query happens 96.06 - 154.49 seconds. The retrieval question matches the retrieved video because the cheerleaders in the image are performing an energetic and synchronized routine, including mid-air flips, runs, and formations, culminating in a stunning triangle formation held by three rows of women.

Figure 19: An example of video-to-video moment retrieval from the IVCR-200K dataset.



The woman removes a second bolt.

The intent is temporal video grounding. The given query happens 5.18 - 16.53 seconds. The retrieval question "The woman removes a second bolt" matches the retrieved video because the woman is seen removing four bolts from a tire in the video. The specific moment captured in the video is when the woman is in the process of removing the second bolt.

Chandler hugs Joey, tells him he cares for his friendship and then walks out the door with his briefcase.

The intent is temporal video grounding. The given query happens 2 - 7.8 seconds. The retrieval question is related to a scene from a TV show where Chandler hugs Joey, tells him he cares for their friendship, and then walks out the door with his briefcase. The retrieval question matches the retrieved video because it captures the essence of the scene, including the emotions and actions of the characters involved.

Someone watches him go, then turns her attention back to someone as she dances with him joyfully.

The intent is video retrieval. The given query happens in 1th video. The retrieval question is a general description of the scene in the retrieved video. The video shows a man and a woman dancing together, and the woman's attention is drawn to someone else, which matches the description in the retrieval question.

a man with no shirt is running at the same pace uphill

The intent is video retrieval. The given query happens in 2th video. The retrieval question matches the retrieved video because the man in the video is running uphill without a shirt. The logical relationship between the retrieval question and the retrieved video is that the man's actions in the video align with the description in the question.

Figure 20: An example of video-to-video moment retrieval and moment-to-video retrieval from the IVCR-200K dataset.

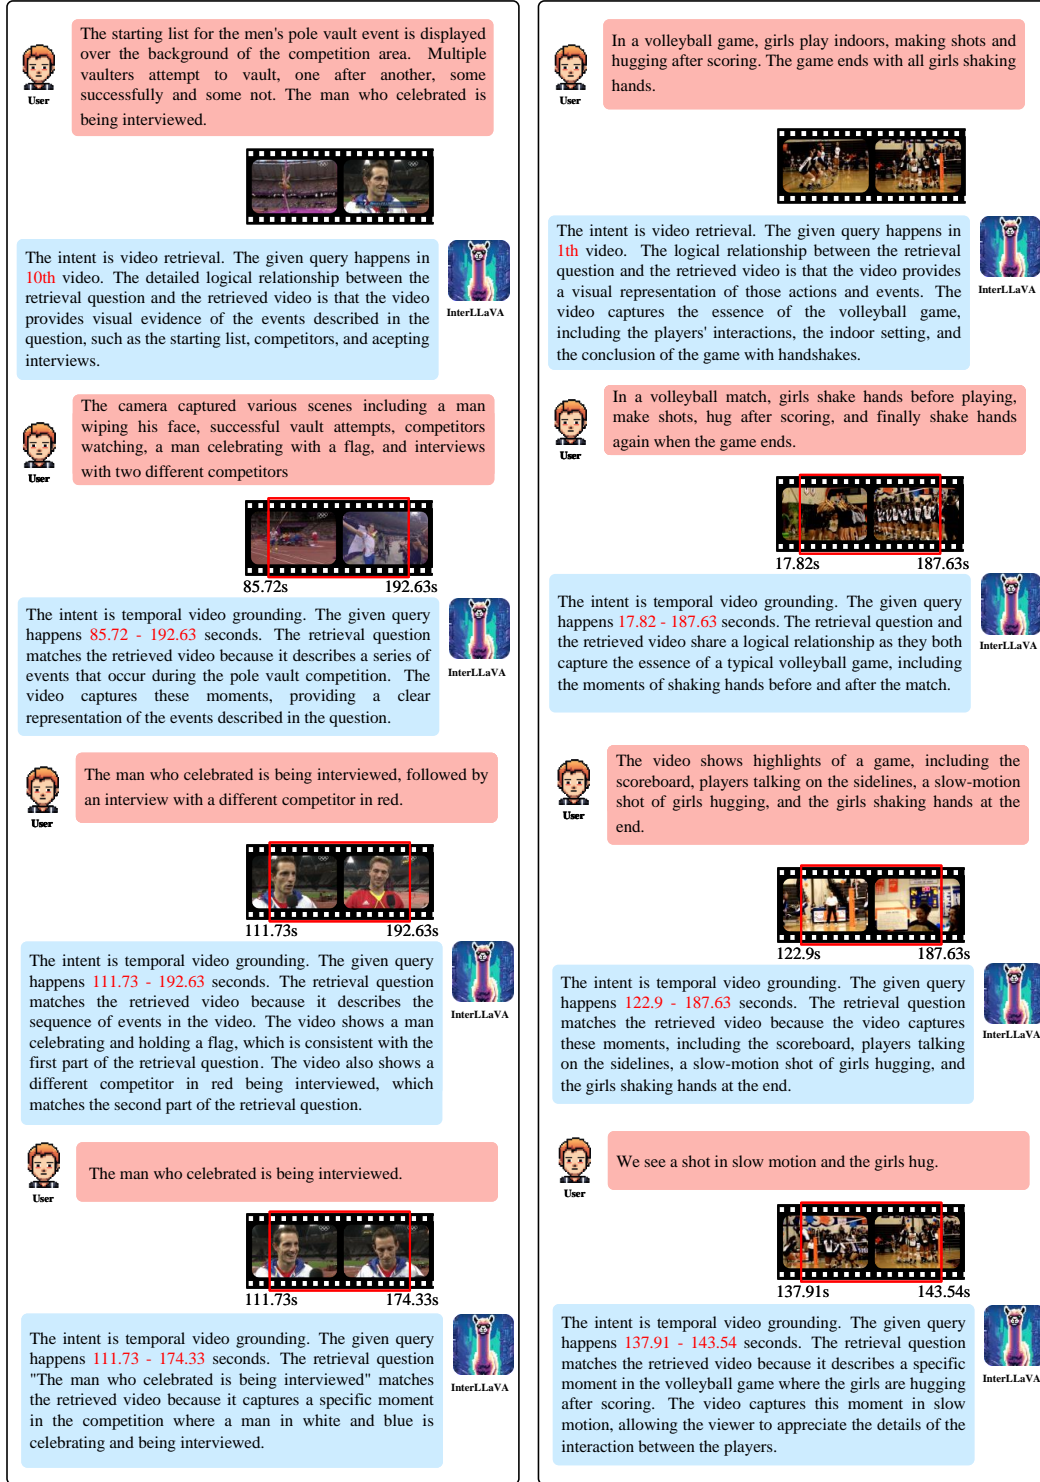


Figure 21: An example of video-to-video moment retrieval and moment-to-video retrieval from the IVCr-200K dataset.

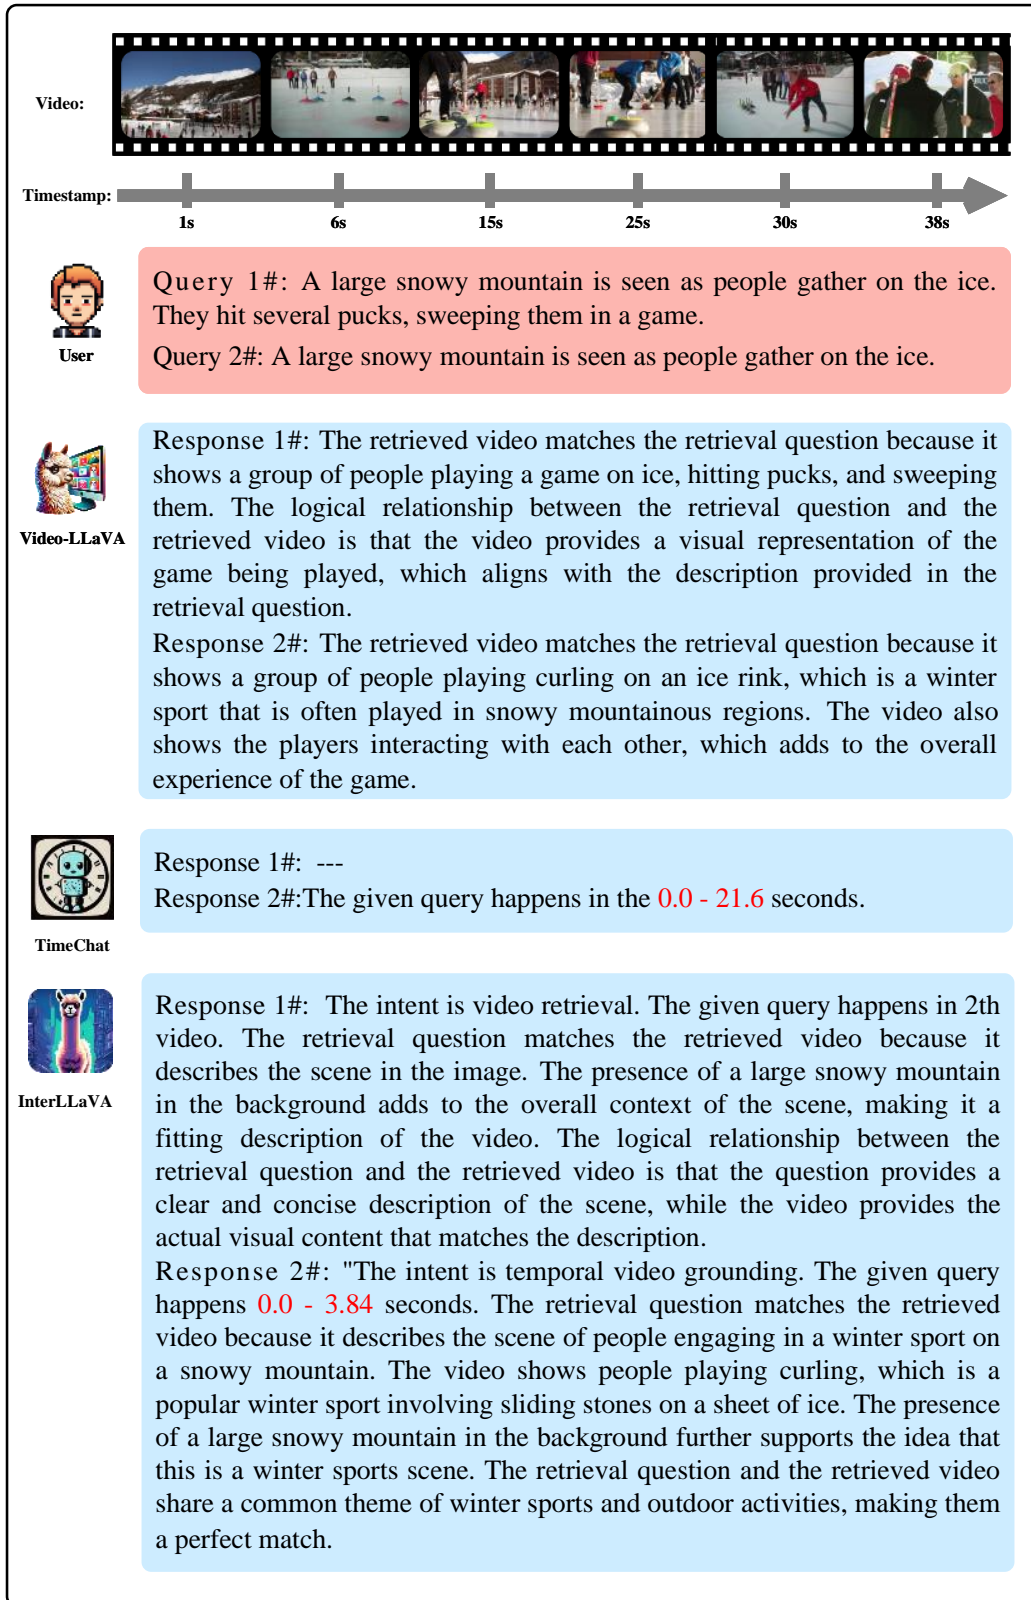


Figure 22: Comparing InterLLaVA with other video large language models.

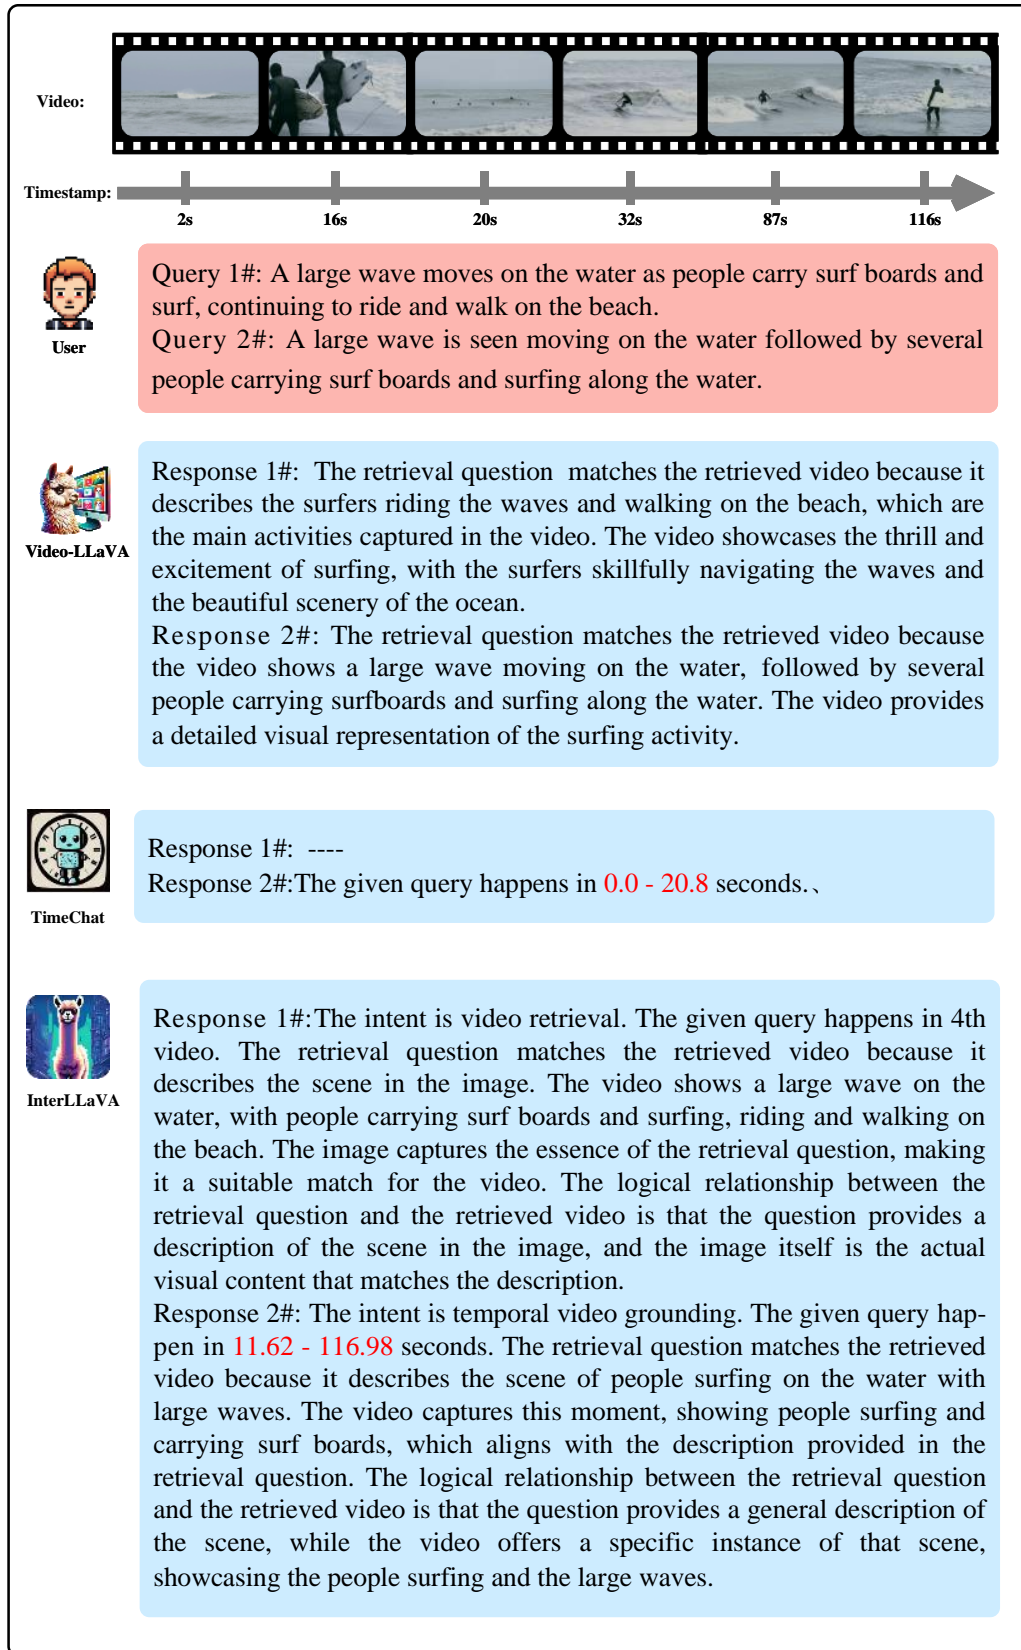


Figure 23: Comparing InterLLaVA with other video large language models.

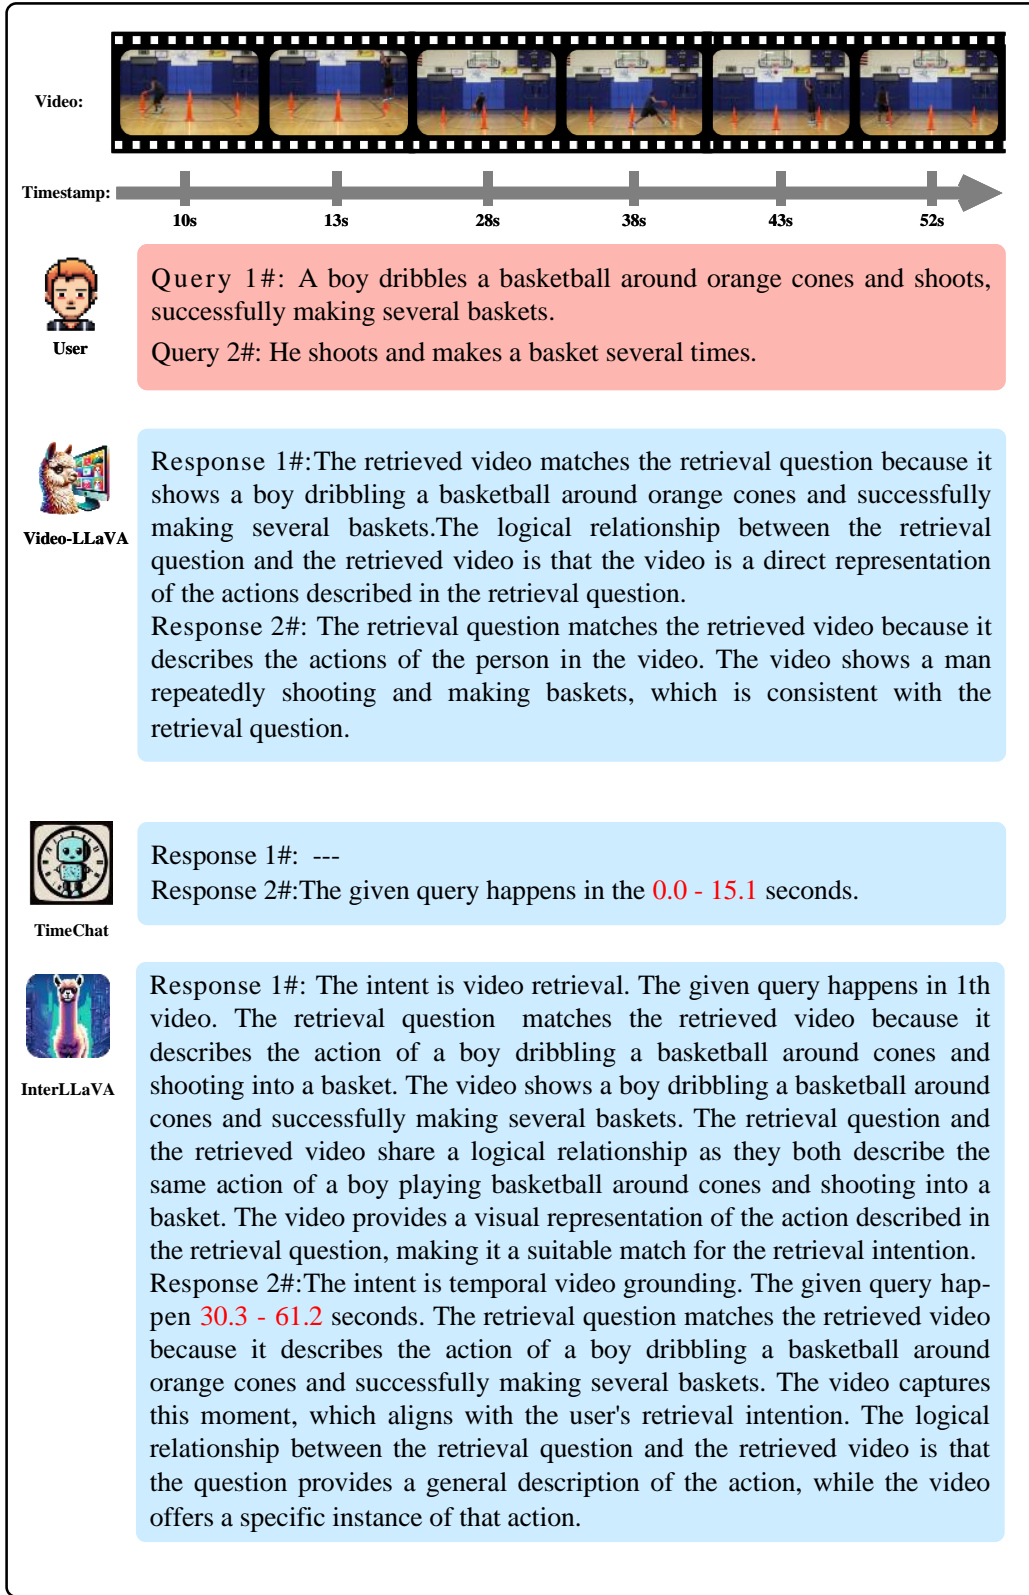


Figure 24: Comparing InterLLaVA with other video large language models.

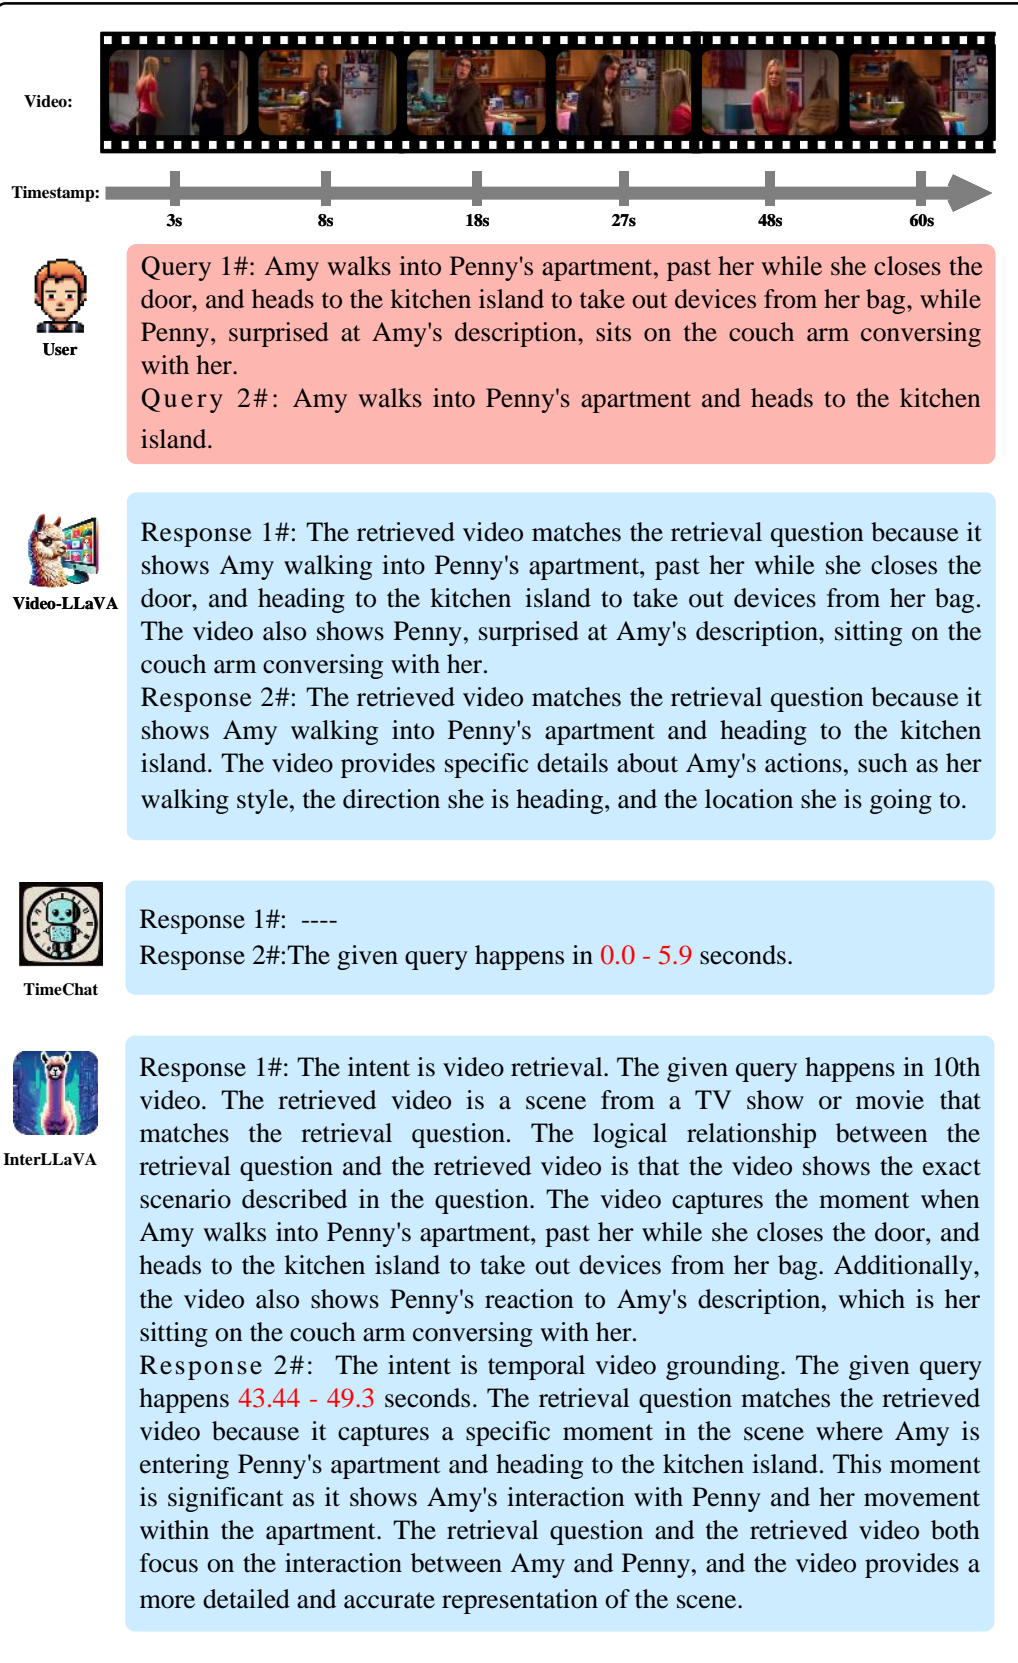


Figure 25: Comparing InterLLaVA with other video large language models.