

Tuning Large Multimodal Models for Videos using Reinforcement Learning from AI Feedback

Anonymous ACL submission

Abstract

Recent advancements in large language models have influenced the development of video large multimodal models (VLMs). Previous approaches for VLMs involve Supervised Fine-Tuning (SFT) with instruction-tuned datasets, integrating LLM with visual encoders, and additional learnable parameters. Here, aligning video with text, and vice versa, remains a challenge, primarily due to the insufficient quality and quantity of multimodal instruction-tune data compared to that of text-only. This discrepancy often results in alignments that poorly ground the video content. To address this, we present a novel alignment strategy that employs a multimodal AI system equipped with Reinforcement Learning from AI Feedback (RLAIF), providing self-preference feedback to refine itself and facilitating the alignment of video and text modalities. Our approach uniquely integrates detailed video descriptions as context into a multimodal AI system during preference feedback generation to enrich the understanding of video content, a process we call *context-aware reward modeling*. Empirical evaluations on various video benchmarks demonstrate that our VLM-RLAIF outperforms existing approaches, including the SFT model. Public release of our code, models, and datasets is forthcoming.

1 Introduction

Large language models (LLMs) are advancing many language and multimodal AI tasks, including those involved with video large multimodal models (VLMs) (Li et al., 2023b; Muhammad Maaz and Khan, 2023; Lin et al., 2023). Extending the logical reasoning and advanced cognitive capabilities of LLMs to the visual domain, VLMs are now remarkably proficient in tasks such as video understanding (Li et al., 2023b), video question answering (Ko et al., 2023) and instruction-following tasks (Muhammad Maaz and Khan, 2023; Luo




Figure 1: **Quantitative comparison of VLMs on various video benchmarks.** The video question answering (VQA) task is marked in purple, video-based generative task in blue, the text-to-video (T2V) retrieval task in violet and the action recognition (AR) task in orange color. VLM-RLAIF achieves superior performances on a broad range of video benchmarks compared to previous approaches, including VLM-SFT. Comprehensive comparisons are provided in Tables 1, 2 and 3.

et al., 2023). These models include publicly available LLMs (Touvron et al., 2023; Chiang et al., 2023; Taori et al., 2023) with visual encoders and additional learnable parameters (Hu et al., 2022; Liu et al., 2023b; Li et al., 2023a). To adapt LLMs to the video modality, thus improving their ability to interpret visual content, they all undergo a supervised fine-tuning (SFT) stage using multimodal instruction-tune data (Luo et al., 2023; Muhammad Maaz and Khan, 2023; Li et al., 2023b).

However, multimodal alignment between video and text faces a significant challenge of deficiency in volume and quality of multimodal instruction-tune data compared to text-only data; text-only data are typically abundant and diverse, while multimodal data are often limited in both quantity and comprehensiveness (Wei et al., 2021; Liu et al.,

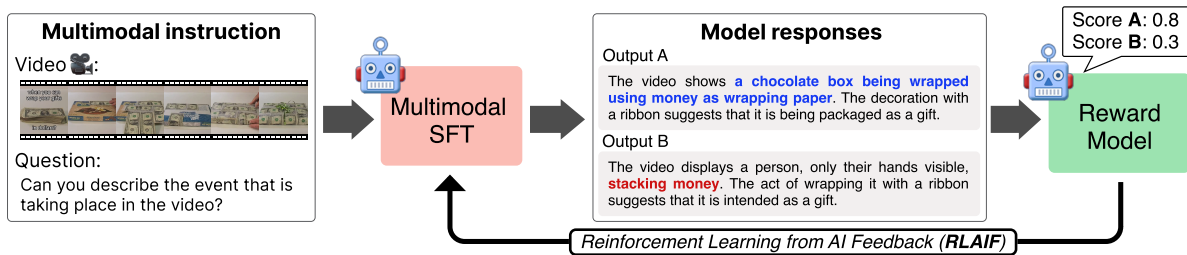


Figure 2: **Illustration of the proposed VLM-RLAIF.** An LLM tuned with video-text instruction-following data, *i.e.*, multimodal SFT model, often produces responses that are not temporally and visually grounded to the video input, as depicted in red color. We propose a method that involves using the VLMM to supervise itself by providing self-preference feedback of generated responses using reward model, refining itself and facilitating the alignment of video and text modalities.

2023a). This often leads the VLMMs to generate responses that are not properly grounded in the visual content, as illustrated by the SFT model responses in Fig. 2.

To address the issue of VLMMs producing inadequately grounded response, we propose a novel method to align video with text that involves using the VLMM to supervise itself by providing preference feedback of generated responses, as shown in Fig. 2. Specifically, we propose to use Reinforcement Learning from AI Feedback (RLAIF) (Bai et al., 2022; Lee et al., 2023) for multimodal alignment. Unlike the Reinforcement Learning with Human Feedback (RLHF), which has been successful in aligning text-only or image-text based AI agents with human intentions (Ouyang et al., 2022; Sun et al., 2023a; Yu et al., 2023), the RLAIF allows for *scalable oversight* with minimal human intervention. In particular, we improve AI feedback by leveraging detailed video descriptions as a context during the generation of preference feedback, focusing on improved clarity in video content, a process we refer to as *context-aware reward modeling* (Sec. 3.1). In addition, to compensate for the limited multimodal instruction-tune data for training the SFT, we augment it with a human-labeled video question answering and an object-centric multimodal instruction-tune dataset. Further, to effectively utilize this expanded dataset, we propose a simple curriculum training strategy for enhancing the alignment between video and text modality (Sec. 3.2).

We call our proposed framework of training an VLMM with AI feedback as Video Large Multimodal model with RLAIF or **VLM-RLAIF** for short. Our empirical studies show that our aligned VLM-RLAIF exhibits superior performance compared to state-of-the-art VLMMs across a wide

array of video benchmarks, as illustrated in Fig. 1.

We summarize our contributions as follows:

- We propose a novel alignment method for video large multimodal models (VLMMs), utilizing Reinforcement Learning from AI feedback (RLAIF) to align video-text modalities effectively.
- We enhance AI’s feedback by proposing context-aware reward modeling, focusing on improved clarity and comprehension in video.
- We enrich the SFT model’s training by introducing additional instruction-tune data and applying a simple curriculum strategy.
- We demonstrate the effectiveness of our proposed VLM-RLAIF on various video understanding benchmarks by a noticeable margin.

2 Related Work

Multimodal large model. Recently, there have been significant advances for LLMs to go beyond natural language understanding, extending into the realm of multimodal comprehension. The goal is to develop LLMs capable of understanding various modalities, *e.g.*, image (Liu et al., 2023a), video (Li et al., 2023b; Lin et al., 2023), 3D point-cloud (Guo et al., 2023) and *etc.*

To make the LLMs multimodal, most of the work utilize a pretrained encoder, such as CLIP (Radford et al., 2021a), Q-former (Li et al., 2022) or ImageBind (Girdhar et al., 2023), to extract each modality’s representations from data. These representations are then projected into the token embedding space of the language model. Then, the models undergo supervised fine-tuning (SFT) with synthetically generated, modality-specific instruction-following datasets. These approaches, adopted in LLaVA (Liu et al., 2023a), Video-LLaVA (Lin

et al., 2023) or Point-LLM (Guo et al., 2023), facilitate the development of proficient conversations grounded in additional modality.

Reinforcement learning from feedback. To operate the model safely and in accordance with human intentions, Reinforcement Learning from Human Feedback (RLHF) has been proposed as a viable solution (Ouyang et al., 2022; Sun et al., 2023a). By collecting preferences from human evaluators, it usually trains the reward model that gives a high reward to the preferred output of the model. However, a significant challenge in this process is the annotation cost associated with selecting the preference. To mitigate this issue, Reinforcement Learning from AI Feedback (RLAIF) was proposed (Bai et al., 2022; Lee et al., 2023; Sun et al., 2023b). RLAIF capitalizes on the inherent ability of Large Language Models (LLMs) to rate or evaluate the generated responses from the SFT model, allowing the LLM itself to assign preferences.

3 VLM-RLAIF Framework

To overcome the limited scalability of human feedback in RLHF, we use AI’s feedback to align multimodality between video and text, reducing the reliance on exhaustive human-annotated preferences (Ouyang et al., 2022; Sun et al., 2023a). In particular, we improve the feedback process by using detailed video descriptions, thereby achieving better contextual clarity in video content. The training procedure of VLM-RLAIF can be summarized into three stages as follows:

Supervised fine-tuning (SFT). We first fine-tune an LLM, e.g., Vicuna, using supervised learning on synthetically generated video-text instruction-tune data (Muhammad Maaz and Khan, 2023). This involves the integration of a vision encoder with two linear layers and additional learnable parameters using LoRA (Hu et al., 2022), into the training process. This fine-tuning allows the model to better follow the instructions (Muhammad Maaz and Khan, 2023; Su et al., 2023). Additionally, we improve the SFT process by expanding the instruction-tune data and introducing simple curriculum learning (Sec. 3.2). We refer to this fine-tuned model as the Video Large Multimodal model with SFT or **VLM-SFT** for short.

Reward modeling with AI feedback. A key aspect of the RLAIF involves leveraging a pre-trained

AI model to generate human-like preferences between different responses generated from the same input (Bai et al., 2022; Sun et al., 2023b; Lee et al., 2023). To obtain human-like preference, we employ the VLM-SFT as a judge to assess preferences. Once preferences are judged, we train a reward model (RM) based on preferences using a cross-entropy loss, following the Bradley-Terry model for estimating score functions from pairwise preferences (Ouyang et al., 2022; Sun et al., 2023a). We describe the training procedure for collecting preferences and training the reward model in Sec. 3.1. The RM give higher score reward to the better response and lower score reward to the less appropriate one in a pair of responses (see examples in Appendix Fig. 12), thus guiding the policy model using reinforcement learning (RL).

Reinforcement learning from AI feedback. We finally fine-tune a supervised policy model, initialized from the VLM-SFT, aiming to optimize the scalar reward output of the trained RM by reinforcement learning. Specifically, we use the Proximal Policy Optimization (PPO) algorithm (Schulman et al., 2017), following (Ouyang et al., 2022; Sun et al., 2023b,a).

3.1 Context-Aware Reward Modeling

For VLM-SFT to select preference grounded on the video, we argue that a detailed understanding of video content is necessary for more accurate and contextually relevant decisions by the VLM-SFT. However, the current video encoder presents challenges in accurately encoding the temporal details of videos as they are based on the image encoder (Radford et al., 2021b).

Context-aware preference selection. We propose to explicitly integrate detailed video descriptions, referred to as *context*, into the preference selection workflow, thereby imparting additional contextual clarity to the VLMM, as illustrated in Figures 3-(2) and 4. Specifically, we break down the video into smaller clips, then generate descriptive captions for each video segment (refer to Appendix Sec. D for more details). Subsequently, these individual captions are concatenated, which we call a *narrative* of the video. The narrative is then provided to a judge model, i.e., VLM-SFT, for better preference selection. The context not only improves the VLM-SFT’s ability to obtain a comprehensive view of the video content, but also enables it to identify the most suitable response

Context-Aware Reward Modeling

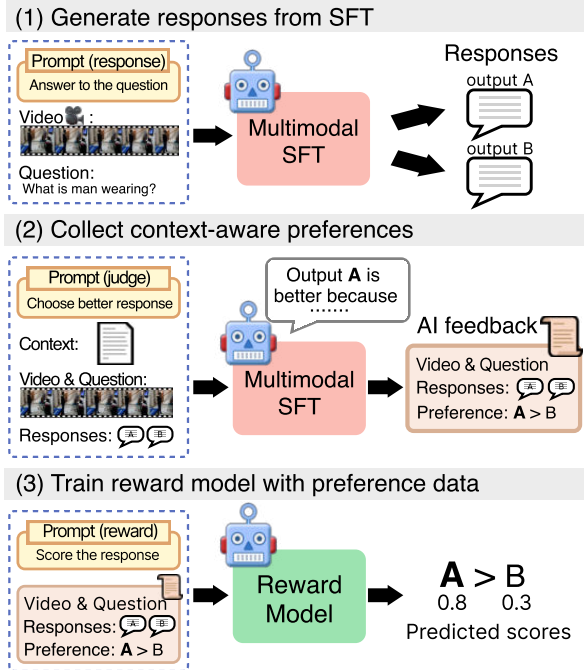


Figure 3: **The three stages of the proposed context-aware reward modeling.** The work flow of each stage is as follows: (1) The SFT model produces two candidate responses from the provided video and question. (2) With the video, question and responses at hand, the SFT model utilize context information and guiding prompt to evaluate the responses. (3) The RM is trained using the preference pairs generated in the previous step as indicated in orange box. Dotted box in each stage denotes a model’s input: the first is for generating responses using the SFT model, the second is for the judge model to evaluate and choose the superior response between options A and B, and the third is for training the RM. Each input includes a task-specific prompt, denoted by the yellow box, tailored to guide the model’s operation within its respective function (more in Appx. Sec. C).

for the video (see Sec. 4.3 for empirical results). Integrating the context with instruction inputs with a specific prompt (rules for generating preferences as illustrated in Appendix Fig. 9), marked in dotted boxes in Fig. 3-(2), allows us to collect context-aware preferences.

Training the reward model. We design the reward model (RM) to assign higher scores to responses considered *better* and lower scores to those considered *worse* in quality. Starting from VLM-SFT with the final linear layer removed, we train a model to take an input prompt and response, *i.e.*, marked in dotted boxes of Fig. 3-(3), and output a scalar reward value. Using the preference dataset produced from the VLM-SFT, we train the RM with a cross-entropy loss. Specifically,

we use 13B VLM-SFT to train the reward model, as it gives slightly better performance than 7B VLM-SFT (see the quantitative comparison discussed in Appendix Sec. A). Note that, after training RLAIIF using the RM, the 7B VLM-RLAIIF significantly surpass the 13B VLM-SFT, thus validating the effectiveness of our proposed framework in aligning video and text modalities (refer to Sec. 4.2 for more details).

3.2 Two-stage Curriculum SFT

During the SFT process, we initially train the LLM with a open-sourced video-text instruction-tune dataset (Muhammad Maaz and Khan, 2023). To improve the VLMM, we not only augment our training with additional video-text instruction-tune datasets but also propose a novel curriculum learning strategy.

Augmenting video instruction-tune dataset.

To improve the video understanding ability, we first augment the video-text instruction-tune dataset (Muhammad Maaz and Khan, 2023) with existing human-annotated video question answering datasets (Xiao et al., 2021; Li et al., 2020). In particular, we focus on obtaining instruction-tune dataset that encompass both visual and temporal intricacies for video comprehension. To obtaining such dataset, we use Video Localized Narratives (Voigtlaender et al., 2023), a dataset that provides detailed descriptions of the appearance and complex actions of all objects in the video. We transform the narrative dataset into an object-centric instruction-tune dataset format using ChatGPT similar to prior work (Muhammad Maaz and Khan, 2023). More details about the generated instruction-tune dataset can be found in the Appendix Sec. B.

Curriculum-based supervised fine-tuning.

Inspired by the human learning process, we argue the importance of the training sequence, the curriculum learning principles (Chang et al., 2021) for learning VLMMs. We propose a simple two-stage curriculum learning strategy for SFT, aiming to optimize the instruction-following proficiency of VLMMs.

For the curriculum based structured learning, we divide the SFT dataset based on complexity into two segments: ‘easy’ and ‘hard’. The model is first trained on ‘easy’ tasks to learn basic concepts, then progressed to ‘hard’ tasks for advanced challenges. For the difficulty or easiness of the data, given that

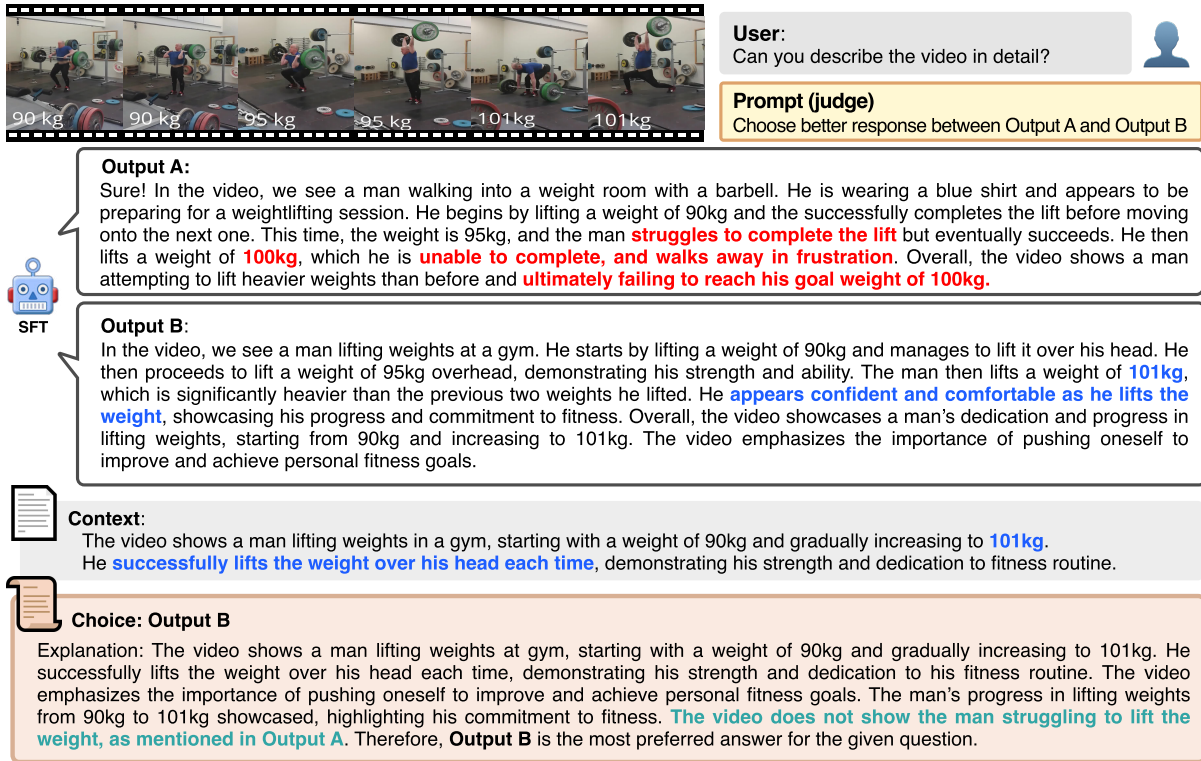


Figure 4: **An example of context-aware preference selection.** We demonstrate our model’s ability to generate preference feedback, *i.e.*, preferred choice and explanation marked in orange box, on given instruction input, prompt, two responses and context. **Red** color denotes an incorrect response, while **blue** color signifies a correctly grounded response with respect to the visual input. The rationale behind selecting ‘Output B’ as the preferred choice is indicated in **green**.

longer answers often require deeper comprehension of the context and enhanced proficiency in creating syntactically complex sentences (Xu et al., 2020; Agrawal and Singh, 2023; Ranaldi et al., 2023), we use *answer length* as our criterion for sample difficulty, *i.e.*, the longer the answer sentence, the more difficult the task is.

4 Experiments

4.1 Experimental Setup

Model details. We initiate training of the VLM-SFT, building on a pre-trained image-text model (Liu et al., 2023a), with various video-text instruction-tune datasets. In particular, we employ a video projection layer consisting of two linear layers with ReLU activation function between them. Upon establishing the VLM-SFT, we train the RM using the VLM-SFT for both its foundation and the generation of preference data. Subsequently, we train the RM using reinforcement learning (VLM-RLAIF). The policy model is initialized from the VLM-SFT, while the value model is initialized from the RM.

Dataset details. For the SFT dataset, we utilize the open-source video-text instruction-tune dataset (80k) (Muhammad Maaz and Khan, 2023; Li et al., 2023b) and video question answering datasets (67k) (Xiao et al., 2021; Li et al., 2020). More importantly, we generate object-centric narrative video-text instruction-tune dataset (180k) for training the VLM-SFT (Sec. 3.2). For the two-stage curriculum learning, we divide the instruction-tune data for SFT into two groups based on the difficulty; easy (214k) and hard (113k) data. To train the RM, we first generate responses from existing instruction-tune data (Muhammad Maaz and Khan, 2023) and generate preferences using them (40k). Then, we again use the existing instruction-tune dataset (100k) (Muhammad Maaz and Khan, 2023) for RL with the trained RM.

Training details. For the video input, we uniformly sample 50 frames from each video and extract spatial and temporal features from them using CLIP visual encoder, similar to (Muhammad Maaz and Khan, 2023). In the two-stage SFT, we set both the LoRA rank and α to 32, respectively, and train the VLM-SFT for one epoch at each stage.

Methods	LLM Size	Video-based Generative Performance				
		Correctness	Detail	Context	Temporal	Consistency
VideoChat (Li et al., 2023b)	7B	2.23	2.50	2.53	1.94	2.24
LLaMA-Adapter (Zhang et al., 2023b)	7B	2.03	2.32	2.30	1.98	2.15
VideoLLaMA (Zhang et al., 2023a)	7B	1.96	2.18	2.16	1.82	1.79
Video-ChatGPT (Muhammad Maaz and Khan, 2023)	7B	2.40	2.52	2.62	1.98	2.37
Valley (Luo et al., 2023)	7B	2.43	2.13	2.86	2.04	2.45
BT-Adapter (Liu et al., 2023b)	7B	2.68	2.69	3.27	2.34	2.46
VTimeLLM (Huang et al., 2023)	7B	2.78	3.10	3.40	2.49	2.47
Video-LLaVA [†] (Lin et al., 2023)	7B	2.84	2.86	3.44	2.46	2.57
VideoChat2 (Li et al., 2024)	7B	3.02	2.88	3.51	2.66	2.81
LLaMA-VID (Li et al., 2023d)	7B	2.96	3.00	3.53	2.46	2.51
LLaMA-VID (Li et al., 2023d)	13B	3.07	3.05	3.60	2.58	2.63
GPT-4V (OpenAI et al., 2023)	-	3.85	3.45	3.84	3.63	2.8
VLM-SFT	7B	2.79	2.82	3.37	2.28	2.49
VLM-RLAIF	7B	3.63	3.25	4.00	3.23	3.32
Δ (RLAIF - SFT)	-	+0.84	+0.43	+0.63	+0.95	+0.83

Table 1: **Quantitative comparison between different VLMs on video-based generative performance benchmark.** Our approach, VLM-RLAIF, shows a performance improvement over previous approaches, with the exception of GPT-4V which requires much more computational resource than ours, and demonstrates noticeable enhancements across five criteria when compared to the VLM-SFT. Here, Δ (RLAIF - SFT) indicates the improvement of RLAIF model over SFT model. [†] denotes reproduced results using the author’s implementation.

For RL, we use QLoRA (Detmers et al., 2023), following (Sun et al., 2023a), setting the rank to 64 and α 16 for computational efficiency and train the policy model for one epoch. All models are trained using $8 \times$ NVIDIA A100 GPUs (80G).

4.2 Quantitative Analysis

We evaluate our proposed VLM-RLAIF on various video benchmarks including video-based generative benchmark, zero-shot video question answering (Muhammad Maaz and Khan, 2023; Lin et al., 2023; Li et al., 2023d), text-to-video retrieval, and action recognition (Li et al., 2023c).

Video-based generative performance. We evaluate VLMs on the video-based generative performance benchmark (Muhammad Maaz and Khan, 2023) that measures five criteria of generated text. In specific, these assess the relevance of the model’s output to the video content, its capacity to capture essential details and contextual information, its understanding of temporal sequences, and the consistency in responding to varied yet related queries. As shown in Tab. 1, the VLM-RLAIF performs *on par with* GPT-4V (OpenAI et al., 2023), which requires much more computational resources than ours (*i.e.*, not a fair comparison), and outperforms previous approaches and the VLM-SFT.

Zero-shot video question answering. To evaluate the reasoning ability of VLLMs, we conduct a quantitative evaluation of video question answering (VideoQA) abilities on three datasets (Xu

et al., 2017; Yu et al., 2019), following (Muhammad Maaz and Khan, 2023).

The results, as shown in Table 2, indicate that the VLM-RLAIF significantly outperforms previous approaches, including VLM-SFT. Notably, VLM-RLAIF exceeds VLM-SFT by 9.2%, 10.6%, and 13.2% in accuracy and by 0.4, 0.4, and 0.3 in score across all datasets. We believe that the better visually-aligned response generated from the VLM-RLAIF improves the performance (see quantitative analysis in Figures 6 and 11).

Zero-shot text-to-video retrieval. For this task, we follow the procedure proposed in (Li et al., 2023c), which compares CLIP score between generated description and ground-truth caption. Table 3 illustrates the summarized performance comparison to various VLMs.

In the two datasets, *i.e.*, MSVD and MSRVT, the proposed VLM-RLAIF clearly outperforms other methods including our VLM-SFT by the help of better alignment by the proposed components.

Zero-shot action recognition. Following the VLMs evaluation procedure proposed in (Li et al., 2023c), we conduct zero-shot action recognition task using two benchmark datasets, *e.g.*, UCF101 and HMDB51. We summarize results of various VLMs in Tab. 3. In two datasets, the proposed VLM-RLAIF again clearly outperforms other methods including our VLM-SFT.

Methods	LLM Size	MSVD-QA		MSRVTT-QA		ActivityNet-QA	
		Acc.	Score	Acc.	Score	Acc.	Score
FrozenBiLM (Yang et al., 2022)	1B	32.2	-	16.8	-	24.7	-
VideoChat (Li et al., 2023b)	7B	56.3	2.8	45.0	2.5	26.5	2.2
LLaMA-Adapter (Zhang et al., 2023b)	7B	54.9	3.1	43.8	2.7	34.2	2.7
VideoLLaMA (Zhang et al., 2023a)	7B	51.6	2.5	29.6	1.8	12.4	1.1
Video-ChatGPT (Muhammad Maaz and Khan, 2023)	7B	64.9	3.3	49.3	2.9	35.2	2.7
Valley (Luo et al., 2023)	7B	60.5	3.3	51.1	2.9	45.1	3.2
BT-Adapter (Liu et al., 2023b)	7B	67.5	3.7	57.0	3.2	45.7	3.2
Video-LLaVA (Lin et al., 2023)	7B	70.7	3.9	59.2	3.5	45.3	3.3
VideoChat2 (Li et al., 2024)	7B	70.0	3.9	54.1	3.3	49.1	3.3
LLaMA-VID (Li et al., 2023d)	7B	69.7	3.7	57.7	3.2	47.4	3.3
LLaMA-VID (Li et al., 2023d)	13B	70.0	3.7	58.9	3.2	47.5	3.3
VLM-SFT	7B	67.2	3.6	52.4	3.0	44.1	3.2
VLM-RLAIF	7B	76.4	4.0	63.0	3.4	57.3	3.5
Δ (RLAIF - SFT)	-	+9.2%	+0.4	+10.6%	+0.4	+13.2%	+0.3

Table 2: **Quantitative comparison between different VLMMs on zero-shot video question answering benchmark.** VLM-RLAIF outperforms previous work across three video-question answering benchmarks.

Methods	LLM Size	T2V Retrieval				Action Recognition			
		MSVD		MSRVTT		UCF101		HMDB51	
		R@1	R@5	R@1	R@5	Top-1	Top-5	Top-1	Top-5
Video-ChatGPT [†] (Muhammad Maaz and Khan, 2023)	7B	26.03	51.25	14.60	33.80	51.49	79.25	37.10	63.97
Video-LLaVA [†] (Lin et al., 2023)	7B	29.34	55.35	18.70	38.60	52.33	80.86	36.64	64.03
LLaMA-VID [†] (Li et al., 2023d)	7B	27.28	53.40	17.00	35.10	56.58	82.79	38.85	65.27
VLM-SFT	7B	26.65	54.27	13.10	30.50	53.03	80.34	38.58	62.37
VLM-RLAIF	7B	36.03	63.40	21.00	40.70	62.83	85.86	44.75	68.37
Δ (RLAIF - SFT)	-	+9.38	+9.13	+7.90	+10.2	+9.80	+5.52	+8.11	+6.00

Table 3: **Quantitative comparison between different VLMMs on zero-shot text-to-video (T2V) retrieval and action recognition.** Following (Li et al., 2023c), we evaluate our proposed VLM-RLAIF on zero-shot T2V retrieval and action recognition. †: reproduced by the authors’ implementation.

SFT datasets			Curr. learning	Video-based Generative Performance				
[A]	[B]	[C]		Corr.	Det.	Cont.	Temp.	Cons.
✓	✗	✗	✗	2.32	2.53	3.03	2.16	2.23
✓	✓	✓	✗	2.43	2.56	3.09	2.19	2.19
✓	✓	✓	✓	2.79	2.82	3.37	2.28	2.49

Table 4: **In-depth analysis for the VLM-SFT training procedure.** ‘[A]’ indicate the multimodal instruction-tune dataset proposed in (Muhammad Maaz and Khan, 2023; Li et al., 2023b). ‘[B]’ represents the use of a human-labeled video question answering dataset (Xiao et al., 2021; Li et al., 2020), while ‘[C]’ refers to the use of an object-centric video narrative instruction-tune dataset (Appendix Sec. B). ‘Curr. learning’ indicates the curriculum learning (Sec. 3.2).

4.3 Detailed Analysis

For a detailed analysis, we use the video-based generative benchmark (Muhammad Maaz and Khan, 2023) specifically, as it is well suited to evaluate the wide-ranging capabilities of VLMM, *i.e.*, focusing on response relevance, detail and context capture, temporal understanding, and consistency across queries.

RLAIF	Context Info.	#Clips	Video-based Generative Performance				
			Corr.	Det.	Cont.	Temp.	Cons.
✗	✗	-	2.79	2.82	3.37	2.28	2.49
✓	✗	-	3.26	3.11	3.74	2.78	3.14
✓	✓	1	3.44	3.20	3.89	2.97	3.36
✓	✓	3	3.63	3.25	4.00	3.23	3.32

Table 5: **Effect of context information on video-based generative performance benchmark.** We investigate the efficacy of using context information for reward modeling (Sec. 3.1). ‘Context Info.’ indicates the use of context in preference selection. ‘# Clips’ denotes the number of segments into which we divide the video to generate the context information.

In-depth analysis of SFT training. We first empirically support the effectiveness of augmenting the SFT dataset with additional instruction-following dataset (Sec. 3.2). The first and second rows of Tab. 4 illustrate the benefits of incorporating this additional dataset in improving performance. On top of that, the application of curriculum learning significantly improves performance, implying the efficacy of curriculum learning for the SFT process (the third row of Tab. 4).




Figure 5: **Effect of preference data size on video-based generative benchmark.** VLM-RLAIF’s performance improves across five metrics as the amount of collected preference data increases. The metrics evaluate whether the model generates correct descriptions of the video, includes detailed explanations, remains contextual, demonstrates temporal understanding, and provides consistent responses to variations of the same question.

Effect of preference data size. Our method’s strength lies in generating synthetic preference feedback in large quantities. To study the benefit of large-sized data, we sweep the size of preference data from 10k – the same quantity utilized in the multimodal image-text RLHF framework (Sun et al., 2023a) – to 40k. As expected, we observe monotonic increases in performance by the increase in data size, as shown in Figure 5.

Context-aware preference labeling. Table 5 presents a study on the impact of context information. This reveals that the use of context information during preference selection improves performance, as demonstrated in the second and third rows in Tab. 5. Moreover, segmenting the video into multiple clips, specifically three, and using their aggregated captions as context further improves performance (see the third and fourth rows). We conjecture that segmenting video into clips with detailed captions enables the model to capture detailed nuances, thereby enriching the contextual clarity for reward modeling.

4.4 Qualitative Analysis

We now qualitatively compares the performance of VLM-SFT and VLM-RLAIF, highlighting their multimodal understanding capabilities in Figure 6. VLM-RLAIF consistently yields more accurate answers than VLM-SFT, as shown in the detailed recognition of the attire of an athlete in the first example, marked in red and blue. The second example further affirms VLM-RLAIF’s benefit in generating better grounded responses to the visual input, where VLM-SFT falls short. More examples are in the Appendix Fig. 11 for the space sake.

5 Conclusion

We propose a novel alignment strategy for VLMs, termed VLM-RLAIF, that uses reinforcement

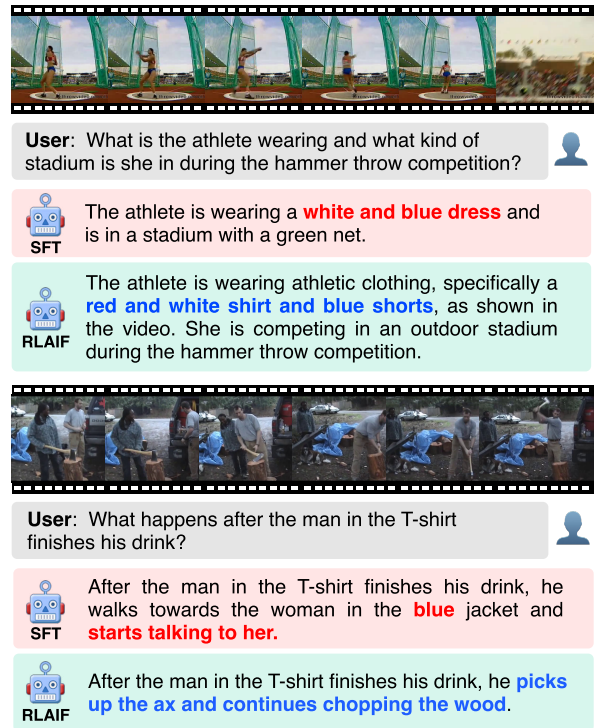


Figure 6: **Samples of VLM-RLAIF and VLM-SFT in video understanding.** Two samples demonstrate better capability of the VLM-RLAIF. Red color indicates visually un-grounded expressions considering video content and blue color well grounded to the video.

learning from AI feedback. To improve multimodal alignment, we propose a context-aware reward modeling, enabling AI to generate feedback for self-improvement. In addition, we expand the instruction-tune dataset for SFT and adopt a curriculum-based training approach, that is particularly effective in the gradual learning of complex video-text relationships.

In our empirical validations, the VLM-RLAIF significantly outperforms previous models in multiple multimodal video-text understanding benchmarks, which implies good generalization performance across tasks.

6 Limitations

Given that our approach utilizes feedback synthesized by the AI model, the effectiveness of our proposed VLM-RLAIF largely depends on the quality of the AI model’s generated responses. In light of recent studies exploring the use of artificially generated data (Koo et al., 2023; Das et al., 2024), we believe that there needs further research to enhance the quality of synthetically generated data, thereby establishing a more reliable RLAIIF system.

In addition, although we have evaluate our model across a range of benchmarks for VLMs, e.g., videoQA, video-based generative tasks, retrieval, and recognition tasks, there are other tasks such as temporal reasoning (Liang et al., 2022) that are necessary for VLMs to be effectively applied in real-world scenarios. Application of our method to these tasks would be a great future research avenue.

References

Ameeta Agrawal and Suresh Singh. 2023. [Corpus complexity matters in pretraining language models](#). In *Proceedings of The Fourth Workshop on Simple and Efficient Natural Language Processing (SustaiNLP)*, pages 257–263, Toronto, Canada (Hybrid). Association for Computational Linguistics.

Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. [Constitutional ai: Harmlessness from ai feedback](#). *arXiv preprint arXiv:2212.08073*.

Ernie Chang, Hui-Syuan Yeh, and Vera Demberg. 2021. Does the order of training samples matter? improving neural data-to-text generation with curriculum learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion

Stoica, and Eric P. Xing. 2023. [Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality](#). 525
526
527

Debarati Das, Karin De Langis, Anna Martin-Boyle, Jaehyung Kim, Minhwa Lee, Zae Myung Kim, Shirley Anugrah Hayati, Risako Owan, Bin Hu, Ritik Parkar, Ryan Koo, Jonginn Park, Aahan Tyagi, Libby Ferland, Sanjali Roy, Vincent Liu, and Dongyeop Kang. 2024. [Under the surface: Tracking the artificiality of llm-generated data](#). 528
529
530
531
532
533
534

Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#). In *NeurIPS*. 535
536
537

Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Man- nat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. 2023. [Imagebind: One embedding space to bind them all](#). In *CVPR*. 538
539
540
541

Ziyu Guo, Renrui Zhang, Xiangyang Zhu, Yiwen Tang, Xianzheng Ma, Jiaming Han, Kexin Chen, Peng Gao, Xianzhi Li, Hongsheng Li, and Pheng-Ann Heng. 2023. [Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following](#). 542
543
544
545
546
547

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*. 548
549
550
551
552

Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. 2023. [Vtimellm: Empower llm to grasp video moments](#). *arXiv preprint arXiv:2311.18445*. 553
554
555

Dohwan Ko, Ji Lee, Woo-Young Kang, Byungseok Roh, and Hyunwoo Kim. 2023. [Large language models are temporal and causal reasoners for video question answering](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4300–4316, Singapore. Association for Computational Linguistics. 556
557
558
559
560
561
562

Ryan Koo, Minhwa Lee, Vipul Raheja, Jong Inn Park, Zae Myung Kim, and Dongyeop Kang. 2023. [Benchmarking cognitive biases in large language models as evaluators](#). 563
564
565
566

Harrison Lee, Samrat Phatale, Hassan Mansoor, Thomas Mesnard, Johan Ferret, Kellie Lu, Colton Bishop, Ethan Hall, Victor Carbune, Abhinav Rastogi, and Sushant Prakash. 2023. [Rlaif: Scaling reinforcement learning from human feedback with ai feedback](#). *arXiv preprint arXiv:2309.00267*. 567
568
569
570
571
572

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. [BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR. 573
574
575
576
577
578
579

580	Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In <i>ICML</i> .	633
581		634
582		635
583		636
584	KunChang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2023b. Videochat: Chat-centric video understanding. <i>arXiv preprint arXiv:2305.06355</i> .	637
585		638
586		639
587		640
588	Kunchang Li, Yali Wang, Yanan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. 2024. <i>Mvbench: A comprehensive multi-modal video understanding benchmark</i> . <i>arXiv</i> .	641
589		642
590		643
591		644
592		645
593	Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. 2020. <i>HERO: Hierarchical encoder for Video+Language omni-representation pre-training</i> . In <i>Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)</i> , pages 2046–2065, Online. Association for Computational Linguistics.	646
594		647
595		648
596		649
597		650
598		651
599		652
600	Shuailin Li, Yuang Zhang, Yucheng Zhao, Qiuyue Wang, Fan Jia, Yingfei Liu, and Tiancai Wang. 2023c. <i>Vlm-eval: A general evaluation on video large language models</i> .	653
601		654
602		655
603		656
604	Yanwei Li, Chengyao Wang, and Jiaya Jia. 2023d. <i>Llama-vid: An image is worth 2 tokens in large language models</i> . <i>arXiv preprint arXiv:2311.17043</i> .	657
605		658
606		659
607	Chen Liang, Wenguan Wang, Tianfei Zhou, and Yi Yang. 2022. Visual abductive reasoning. In <i>IEEE/CVF International Conference on Computer Vision and Pattern Recognition (CVPR)</i> .	660
608		661
609		662
610		663
611	Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. 2023. Video-llava: Learning united visual representation by alignment before projection. <i>arXiv preprint arXiv:2311.10122</i> .	664
612		665
613		666
614		667
615	Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023a. Visual instruction tuning. In <i>NeurIPS</i> .	668
616		669
617	Ruyang Liu, Chen Li, Yixiao Ge, Ying Shan, Thomas H. Li, and Ge Li. 2023b. <i>One for all: Video conversation is feasible without video instruction tuning</i> . <i>arXiv preprint arXiv:2309.15785</i> .	670
618		671
619		672
620		673
621	Ruipu Luo, Ziwang Zhao, Min Yang, Junwei Dong, Minghui Qiu, Pengcheng Lu, Tao Wang, and Zhongyu Wei. 2023. <i>Valley: Video assistant with large language model enhanced ability</i> .	674
622		675
623		676
624		677
625	Salman Khan Muhammad Maaz, Hanoona Rasheed and Fahad Khan. 2023. Video-chatgpt: Towards detailed video understanding via large vision and language models. <i>ArXiv 2306.05424</i> .	678
626		679
627		680
628		681
629	OpenAI, :, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin,	682
630		683
631		684
632		685
	Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mo Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki	686
		687
		688
		689
		690
		691
		692
		693
		694
		695

696	Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lillian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2023. Gpt-4 technical report .	
716	Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In <i>NeurIPS</i> .	
724	Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021a. Learning transferable visual models from natural language supervision. In <i>ICML</i> .	
730	Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021b. Learning transferable visual models from natural language supervision . In <i>Proceedings of the 38th International Conference on Machine Learning</i> , volume 139 of <i>Proceedings of Machine Learning Research</i> , pages 8748–8763. PMLR.	
739	Leonardo Ranaldi, Giulia Pucci, and Fabio Massimo Zanzotto. 2023. Modeling easiness for training transformers with curriculum learning . In <i>Proceedings of the 14th International Conference on Recent Advances in Natural Language Processing</i> , pages 937–948, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.	
746	John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms . <i>arXiv preprint arXiv:1707.06347</i> .	
750	Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. 2023. Pandagpt: One model to instruction-follow them all . <i>arXiv preprint arXiv:2305.16355</i> .	
	Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, Kurt Keutzer, and Trevor Darrell. 2023a. Aligning large multimodal models with factually augmented rlhf . <i>arXiv preprint arXiv:2309.14525</i> .	754 755 756 757 758 759
	Zhiqing Sun, Yikang Shen, Hongxin Zhang, Qinhong Zhou, Zhenfang Chen, David Cox, Yiming Yang, and Chuang Gan. 2023b. Salmon: Self-alignment with principle-following reward models . <i>arXiv preprint arXiv:2310.05910</i> .	760 761 762 763 764
	Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model . https://github.com/tatsu-lab/stanford_alpaca .	765 766 767 768 769
	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models . <i>arXiv preprint arXiv:2302.13971</i> .	770 771 772 773 774 775
	Paul Voigtlaender, Soravit Changpinyo, Jordi Pont-Tuset, Radu Soricut, and Vittorio Ferrari. 2023. Connecting Vision and Language with Video Localized Narratives . In <i>IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> .	776 777 778 779 780
	Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2021. Finetuned language models are zero-shot learners . <i>CoRR</i> , abs/2109.01652.	781 782 783 784 785
	Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. 2021. Next-qa: Next phase of question-answering to explaining temporal actions . In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)</i> , pages 9777–9786.	786 787 788 789 790 791
	Benfeng Xu, Licheng Zhang, Zhendong Mao, Quan Wang, Hongtao Xie, and Yongdong Zhang. 2020. Curriculum learning for natural language understanding . In <i>Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics</i> , pages 6095–6104, Online. Association for Computational Linguistics.	792 793 794 795 796 797 798
	Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. 2017. Video question answering via gradually refined attention over appearance and motion . In <i>ACM Multimedia</i> .	799 800 801 802 803
	Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. 2022. Zero-shot video question answering via frozen bidirectional language models . In <i>NeurIPS</i> .	804 805 806 807

808 Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng
809 Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao
810 Zheng, Maosong Sun, and Tat-Seng Chua. 2023.
811 RLhf-v: Towards trustworthy mllms via behavior
812 alignment from fine-grained correctional human feed-
813 back. *arxiv*.

814 Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yuet-
815 ing Zhuang, and Dacheng Tao. 2019. Activitynet-qa:
816 A dataset for understanding complex web videos via
817 question answering. In *AAAI*, pages 9127–9134.

818 Hang Zhang, Xin Li, and Lidong Bing. 2023a. [Video-
819 llama: An instruction-tuned audio-visual language
820 model for video understanding](#). *arXiv preprint
821 arXiv:2306.02858*.

822 Renrui Zhang, Jiaming Han, Chris Liu, Peng Gao, Ao-
823 jun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hong-
824 sheng Li, and Yu Qiao. 2023b. Llama-adapter: Effi-
825 cient fine-tuning of language models with zero-init
826 attention. *arXiv preprint arXiv:2303.16199*.

827	Appendix		
828	A Quantitative Comparison between		
829	Different SFT Model Sizes		
830	We assess the performance of the VLM-SFT using		
831	varying LLM sizes, specifically 7B and 13B,		
832	across video question answering (VideoQA) and		
833	video-based generative benchmarks, detailed in		
834	Table 6. The results indicate similar outcomes in		
835	the zero-shot VideoQA tasks across both model		
836	sizes. However, in video-based generative tasks,		
837	a clear trend emerges: larger LLM sizes lead to		
838	improved performance. This suggests that the in-		
839	creased model capacity inherent in larger LLMs		
840	enhances the model’s ability to capture and gen-		
841	erate more complex and nuanced video content		
842	representations. Thus, we adopt the 13B model		
843	for the reward model (Sec. 3.1), with empirical		
844	evidence provided in Tables 7 and 8.		
845	B Details About Object-Centric		
846	Instruction-Tune Data		
847	To enhance the alignment of Large Language Mod-		
848	els (LLMs) with video content, we prioritize ac-		
849	quiring more video-text instruction-tune dataset		
850	that captures both the visual and temporal complex-		
851	ities of videos. For this, we leverage the Video		
852	Localized Narratives dataset (Voigtlaender et al.,		
853	2023), which contains a comprehensive description		
854	of each object’s appearance and action, along with		
855	their tracked coordinates within videos. To utilize		
856	the best of the rich semantics contained in narrative		
857	dataset, we feed all ground-truth captions to the		
858	model, <i>i.e.</i> , ChatGPT, and prompt it to generate		
859	question and answer pairs, as shown in the tem-		
860	plate in Fig. 7-(a). We aim to transform all object’s		
861	descriptions into instruction-tune data format, <i>i.e.</i> ,		
862	question and answer pair, which demands an un-		
863	derstanding of the visual specifics of each object		
864	and its surroundings, ensuring the questions and		
865	answers are anchored in the video. An example of		
866	the generated instruction-tune data is depicted in		
867	Fig.7-(b), showcasing the approach’s effectiveness		
868	in creating contextually rich instructional content.		
869	C Input Prompts for Reward Modeling		
870	We leverage three types of input prompts for		
871	context-aware reward modeling, as discussed in		
872	Sec. 3.1. Figure 8 presents a designed input prompt		
873	given to our method, resulting in two responses.		
874	Furthermore, Figure 9 illustrates the detailed input		
		prompt used to select a preference between two re-	875
		sponses. Additionally, Figure 10 displays the input	876
		prompt employed for generating a reward score.	877
	D Details About Generating Contextual		
	Information		
	For effective utilization of VLM-SFT in prefer-		
	ence selection, a nuanced understanding of video is		
	essential. To achieve this, we equip the VLM-SFT		
	with detailed video descriptions, referred to as ‘con-		
	text’, during preference selection process, to en-		
	hance contextual clarity in the video. Specifically,		
	we start by segmenting the video into small clips,		
	each containing up to 20 frames, and then employ		
	the VLM-SFT to generate a detailed video descrip-		
	tion for each segment with an input prompt, ‘De-		
	scribe this video in detail’. We then aggregate		
	the descriptions of each segment by concatenating		
	them, thereby generating <i>context</i> of the video.		
	E Various LLM Size for Reward Model		
	In Tables 7 and 8, we assess the performance of		
	VLM-RLAIF using varying reward model(RM)		
	sizes, each initialized from VLM-RLAIF-7B and		
	VLM-RLAIF-13B. In both settings, we use 7B		
	VLM-SFT to initialize the policy model. For both		
	sizes of RM, our RLAIF method showed a perfor-		
	mance gain with a large margin from VLM-SFT		
	across all evaluation methods. Specifically, RLAIF		
	with 7B RM acquired 5-12% improvement on zero-		
	shot video question answering. Scaling up the RM		
	from 7B to 13B result in more improved perfor-		
	mance, except for text-to-video retrieval task R@5		
	metric on MSR-VTT.		

Methods	LLM Size	Video Question Answering						Video-based Generative Perf.				
		MSVD		MSRVTT		ActivityNet		Corr.	Det.	Cont.	Temp.	Cons.
		Acc.	Score	Acc.	Score	Acc.	Score					
VLM-SFT	7B	66.88	3.58	44.61	3.19	49.13	2.91	2.32	2.53	3.03	2.16	2.23
VLM-SFT	13B	65.83	3.56	45.35	2.84	51.59	2.95	2.64	2.73	3.28	2.38	2.44

Table 6: **Quantitative comparison between different size of the VLM-SFT.** We assess the performance of VLM-SFT with varying LLM sizes, specifically 7B and 13B, across zero-shot video question answering and video-based generative benchmarks. We conduct this evaluation without the integration of augmented instruction-tune data and the implementation of two-stage curriculum learning. ‘Corr.’, ‘Det.’, ‘Cont.’, ‘Temp.’ and ‘Cons.’ indicate ‘Correctness’, ‘Detail’, ‘Context’, ‘Temporal’ and ‘Consistency’, respectively.

Methods	LLM Size		Video Question Answering						Video-based Generative Perf.				
	Policy Model	Reward Model	MSVD		MSRVTT		ActivityNet		Corr.	Det.	Cont.	Temp.	Cons.
			Acc.	Score	Acc.	Score	Acc.	Score					
LLaMA-VID (7B)	-	-	69.7	3.7	57.7	3.2	47.4	3.3	2.96	3.00	3.53	2.46	2.51
VLM-SFT	7B	-	67.2	3.6	52.4	3.0	44.1	3.2	2.79	2.82	3.37	2.28	2.49
VLM-RLAIF	7B	7B	75.1	3.9	61.0	3.3	56.1	3.4	3.47	3.14	3.87	3.05	3.30
VLM-RLAIF	7B	13B	76.4	4.0	63.0	3.4	57.3	3.5	3.63	3.25	4.00	3.23	3.32

Table 7: **Quantitative comparison between different size of policy model and reward model for the VLM-RLAIF.** We evaluate the VLM-RLAIF with different model size for policy model and reward model on zero-shot video question answering and video-based generative benchmark.

Methods	LLM Size		T2V Retrieval				Action Recognition			
	Policy Model	Reward Model	MSVD		MSRVTT		UCF101		HMDB51	
			R@1	R@5	R@1	R@5	Top-1	Top-5	Top-1	Top-5
LLaMA-VID (7B)	-	-	27.28	53.40	17.00	35.10	56.58	82.79	38.85	65.27
VLM-SFT	7B	-	26.65	54.27	13.10	30.50	53.03	80.34	38.58	62.37
VLM-RLAIF	7B	7B	33.73	61.95	20.80	42.90	61.09	85.15	43.86	65.88
VLM-RLAIF	7B	13B	36.03	63.40	21.00	40.70	62.83	85.86	44.75	68.37

Table 8: **Quantitative comparison between different size of policy model and reward model for the VLM-RLAIF.** We evaluate the VLM-RLAIF with different model size for policy model and reward model on zero-shot text-to-video retrieval and action recognition tasks.


(a) Prompt for generating object-centric instruction-tune data

```
You play two roles: a human asking questions related to a video and an intelligent chatbot designed to help people find information from a given video. Your task is to first play the role of a human who asks questions related to a video and then play the role of an AI assistant that provides information based on the video content.
-----
##TASK:
Users will provide a narrative caption of a video describing each objects in the video denoted as "Narrations", and you will generate a set of conversation-like questions and answers related to the narration as "Outputs".
The questions should be designed to extract information directly from the given information, so that the provided information or parts of it can serve as the answers.
Generate FIVE different conversational style questions and answers based on the given information.
-----
##GUIDELINES:
- The questions must be like a human conversation and based on the events in the video.
- The questions should be designed to extract information DIRECTLY from the given information, so that it or parts of it can serve as the answers.
- The answers must be in complete sentence, and they should directly reference the information provided.
- The questions can be related to the appearance, motion, trajectory, and reasoning.
-----
##PRINCIPLES:
- Please generate the response in the form of a Python list of dictionary string with keys 'Q' for question and 'A' for answer. Each corresponding value should be the question and answer text respectively.
-----
##SINGLE EXAMPLE:
Narrations:
- Baby girl : A baby girl on the left side wearing a grey t-shirt is carrying an egg then she throws the egg at the head of the man, then the egg falls on the ground and it breaks on a grey surface.
- Man : A man wearing a red t-shirt sitting on a brown mat is talking with the baby girl on a grey surface.
- background : In the background, there is a grey car, a grey surface, a brown mat, and people speaking and crying sounds are audible.

Outputs:
[{"Q": "Who throws the egg at the man?", "A": "A baby girl is throwing the egg at the man."}, {"Q": "At what part of the man is the egg thrown?", "A": "It is thrown to his head."}, {"Q": "Which side of the screen is the baby girl on?", "A": "She is on the left side."}, {"Q": "What is the man wearing?", "A": "The man is wearing red t-shirt."}, {"Q": "Where is the man sitting?", "A": "The man is sitting on the brown mat."}]
-----
Now It's your turn.
Narrations:
{NARRATIONS}

Outputs:
```

(b) Example data



Q: How many parrots are in the group? A: There are seven parrots in the group.
Q: What colors are the parrots in the group? A: There are two grey parrots, three yellow parrots, one green parrot, and one white parrot in the group.
Q: What are the parrots doing? A: They are eating food from the bowl.
Q: What is the white parrot doing? A: The white parrot is walking toward the red-white water feeding container.
Q: What can be seen in the background? In the background, there is a large iron grid with a red-white water feeding container, a steel bowl with food, and white walls with pillars.

Figure 7: **Designed input prompt for object-centric instruction-tune dataset generation from video narrative dataset and the data sample.** In (a), the prompt we used to generate instruction-following dataset is displayed in a white box. It includes the system prompt, a task definition, guidelines, principles, and an example set of input and desired output. Using this prompt with ChatGPT-3.5-turbo, we create an additional instruction following datasets. The resulting example is visualized in (b). The question covers details of the video scenes, such as the number and appearance of parrots and the action of each parrot.

Prompt (response)

A chat between a curious human and an artificial intelligence assistant.
The assistant gives helpful, detailed, and polite answers to the human's questions.

Figure 8: **Designed input prompt fed to an VLM-SFT to generate responses.**

Prompt (judge)

Considering given video, question and context information, select the most preferred (least wrong) answer for the question.

Inputs:

1. Question, the question queried to the AI system
2. Answer1, the first answer prediction from the AI system
3. Answer2, the second answer prediction from the AI system

And, we will provide the AI system with context information about the given video in the form of a paragraph as follows: {context_captions}

Considering the provided definitions and context information, please select the preferred answer (Answer1 or Answer2) for the given instruction (Question) and provide a concise explanation for choosing it as the correct one. When generating the explanation, please provide valid justifications without simply mentioning the answer itself.

Your answer should be formatted as:

Choice:

Explanation:

Now provide your answer in this situation:

Question: {question}

Answer1: {response A}

Answer2: {response B}

When generating output, you should consider the visual signal and the context information provided simultaneously.

Figure 9: **Designed input prompt fed to an VLM-SFT to generate AI's preference feedback.**

Prompt (reward)


USER: Please evaluate the quality of your last response. There are several dimensions you should consider in your evaluation:

1. Accurate: The AI should provide factual and accurate information from the video, and refrain from making statements that are not supported by the video or inconsistent with the video.
2. Helpful: The AI's response should precisely serve the user's needs and interests, while grounding the response in the video.
3. Language Natural: The AI should employ language that flows smoothly and is free from repetitive or awkward constructs.
4. Concise: The AI should efficiently address the task or answer the question, communicating the necessary information with brevity and clarity.

A good response should be accurate, helpful, language natural, and concise.

ASSISTANT: Following your definitions, the quality score of my last response is


Figure 10: **Designed input prompt fed to a reward model to generate scalar reward score.**



User: What is being taped to the gift in the video?

SFT: The person in the video is taping a piece of **paper** to the gift.


RLAIF: The video clearly shows that **money** is being taped to the gift.



User: What tools did the man use to clean the leather sofa in the video?

SFT: The man used a **vacuum cleaner** to clean the leather sofa in the video.


RLAIF: The man used a **brush**, a **sponge**, and a **cleaning solution** to clean the leather sofa in the video.



User: What is the technician doing with the tattoo gun ?

SFT: The technician is using the tattoo gun to draw a tattoo on the **woman's face**.


RLAIF: The technician is using the tattoo gun to create the design on the **woman's ear**.



User: Can you describe what happens in the video ?

SFT: The man is wearing a brown jacket and backpack. He is standing on the road and **smoking a cigarette**.


RLAIF: Yes, the video shows a man playing a **harmonica** on the street and interacting with others.



User: Can you describe what happens in the video ?

SFT: In the video, a man wearing a white t-shirt is seen standing and speaking. He is holding a glass of water and a glass of **orange juice**. He is also **holding a yellow object and a white object**. In the background, there is a black background with a graphical design.


RLAIF: The video shows a man discussing and demonstrating a **weight loss product**. He talks about the product and its benefits, and also shows the product in action.



User: What is happening at the beginning of the video?

SFT: At the beginning of the video, a girl is **seen standing in a store** and talking to a woman.


RLAIF: At the beginning of the video, a woman is **getting her ears pierced** at a piercing station in a store.



User: What safety gear does the little boy wear while skateboarding in the video?

SFT: The little boy is seen wearing a helmet while skateboarding in the video.

RLAIF: The little boy is wearing a helmet and **knee pads** while skateboarding in the video.



User: What are the boys wearing in the video, and what activity are they doing?

SFT: The boys are wearing green shirts and are doing jumping exercises in the video.

RLAIF: The boys are wearing **Santa costumes and green shirts**, and are performing a **series of acrobatic and playful movements**, including jumping, flipping, and rolling around.

Figure 11: **More examples of the comparative results between VLM-SFT and VLM-RLAIF.** We present more qualitative examples of questions and answers generated by the VLM-SFT and VLM-RLAIF. The VLM-RLAIF yields responses that are more accurately grounded in the visual input, as highlighted in **blue** for accurate responses and **red** for less accurate ones from VLM-SFT.



User: Can you describe the video in detail?



Output A: In the video, the bartender prepares a Bloody Mary drink for a customer. The bartender starts by filling a highball glass with ice and adding a few shots, such as **vodka, gin, and tequila**. The bartender then adds a pinch of salt and **squeezes in freshly squeezed lime juice**. Next, the bartender adds a **straw and a slice of lemon** to the glass, **followed by a small amount** of worcesterchire sauce and **hot sauce**. The bartender then **garnishes the drink with a pickle slice and a sprig of parsley**. Finally, the bartender pours in **tomato juice** and **stirs the drink with a long spoon**. Overall, the video shows the process of making a Bloody Mary drink from scratch, including the selection of ingredients, the mixing of the drink, and the final presentation of the drink.

Output B: Sure, the video begins with a bartender standing behind a bar. He is seen preparing a drink by putting ice in a glass. He then pours **shots** into the glass. After that, the bartender adds **tomato juice** to the drink and **stirs it with a spoon**. The final product is a Bloody Mary cocktail. The video showcases the step-by-step process of making a Bloody Mary drink, from ice to final presentation. The bartender's actions in the video are precise and deliberate as he carefully prepares the drink to perfection. The video gives viewers a glimpse into the art of mixology and the careful craftsmanship that goes into preparing a delicious and visually appealing cocktail. Overall, the video is an excellent representation of the skill and precision that goes into creating a Bloody Mary drink.

Reward A: 0.4982 < B: 1.0323



User: Can you describe the video in detail?



Output A: Sure, the video shows a man sitting in a canoe while holding a paddle. The **camera is facing him, and he is talking to the camera man**. As the man starts to paddle, he moves his arms back and forth along the water. The movement of the paddle can be seen in the water, creating ripples and waves. The man continues paddling along the water, and the camera captures the scenery around him. The sound of the water splashing around the man and the paddle can be heard throughout the video. The video seems to be a peaceful and tranquil moment, with the man enjoying his time on the water while the camera captures the beautiful scenery around him. Overall, the video provides a serene and relaxing experience for the viewer, allowing them to escape from their daily routine and enjoy the beauty of nature.

Output B: The video begins with a man sitting on a kayak in a lake, paddling the kayak with his hands while looking forward. The **scene then transitions** to a person in a **red kayak** riding down a river, navigating through the waves and enjoying the thrill of the river. The man in the kayak is shown paddling through the water, maneuvering the kayak with ease and control. The video captures the beauty and serenity of the lake and river scenery, with the sound of the water and the rushing river providing a soothing and peaceful ambiance. The kayaker's skills and technique in navigating the kayak are impressive, as he confidently navigates through the rushing river with ease and control. Overall, the video highlights the thrill and beauty of kayaking in the great outdoors.

Reward A: 0.7730 > B: 0.2276




User: Can you describe the video in detail?



Output A: Sure, the video is about a man who is playing the bagpipe in the middle of a square. There are many cars passing by him, and at one point, he **stops playing and does a handstand on the ground. As soon as he lands, he looks at the camera and smiles**. The man's outfit consists of a traditional **Scottish kilt** and a **hooded jersey**, which is commonly worn for bagpipe performances. The square, where the man is playing, is surrounded by shops and cafes, and there are many people walking around. The background music of the bagpipe is soothing and melodious, and it seems to be blending well with the urban setting. Overall, the video shows a unique and entertaining performance of a traditional Scottish instrument in a modern city setting.

Output B: Sure! The video shows a man playing bagpipes on the street while people are walking by on the sidewalk. He wears a **black jacket and a red knee-length socks**. A few people are seen **stopping to watch him play**. He plays with **bags on the sidewalk**. There are several **cars and a traffic light on the street** while the man plays. The video captures a lively street scene where the man playing bagpipes attracts the attention of passersby. His outfit, which includes a skirt, is unconventional and adds to the overall charm of the video. The traffic light and cars in the background create a bustling atmosphere, and the people walking on the sidewalk add to the overall sense of movement and energy in the video. Overall, the video captures a vibrant street scene with a skilled musician entertaining the passersby.

Reward A: 0.4481 < B: 0.9563



User: Can you describe what happens in the video?



Output A: Sure, in the video, a man can be seen speaking to the camera while standing **in front of a window**. He then proceeds to **clean the window** while still speaking to the camera. The man seems to be using a **cleaning solution** and a squeegee to wipe down the window. Throughout the video, he **pauses at different points to address the camera** and continue talking. The man's movements are deliberate and methodical, suggesting that he is experienced in window cleaning. The video seems to be focused on showcasing the process of window cleaning rather than the man himself. Overall, the video provides a **detailed look at the process of window cleaning and the equipment used**.

Output B: In the video, we see a man standing in a room with a glass door in front of him. He is holding a glass and can be seen wiping the glass with a towel. He is wearing a **red shirt and black pants** and is standing near a window in the living room. The man seems to be enjoying his time in the room, as he **takes a sip from his glass** and continues to wipe the door. In the distance, we can see a **white car parked outside**, and a white car parked in the parking lot is visible as well. There are no other people in the room, and the lighting is bright, indicating that it is daytime. **The man seems to be in a relaxed mood, and the room looks cozy and comfortable**. Overall, the video portrays a **peaceful and serene scene of a man enjoying his time in a well-lit room** in his home.

Reward A: 0.9241 > B: 0.2293

Figure 12: Predictions of SFT model and the predicted reward. We present qualitative examples of questions and two responses generated from the VLM-SFT, and the reward given to each answer using the trained RM. The reward is higher for the answer with more accurately grounded answers, as highlighted in blue for accurate responses and red for less accurate contents.