

TEMPLE: Incentivizing Temporal Understanding of Video Large Language Models via Progressive Pre-SFT Alignment

Shicheng Li¹, Lei Li², Kun Ouyang¹, Shuhuai Ren¹, Yuanxin Liu¹, Yuanxing Zhang³,
Fuzheng Zhang⁴, Lingpeng Kong², Qi Liu², Xu Sun¹

¹National Key Laboratory for Multimedia Information Processing, School of Computer Science, Peking University

²The University of Hong Kong

³Kling Team

⁴Kuaishou Technology

{lisc99,xusun}@pku.edu.cn

Abstract

Video Large Language Models (Video LLMs) have achieved significant success by adopting the paradigm of large-scale pre-training followed by supervised fine-tuning (SFT). However, existing approaches struggle with temporal reasoning due to **weak temporal correspondence in the data** and **over-reliance on the next-token prediction paradigm**, which collectively result in the absence of temporal supervision. To address these limitations, we propose **TEMPLE (TEMporal Preference LEarning)**, a systematic framework that enhances temporal reasoning capabilities through Direct Preference Optimization (DPO). To address temporal information scarcity in data, we introduce an automated pipeline for systematically constructing temporality-intensive preference pairs comprising three steps: selecting temporally rich videos, designing video-specific perturbation strategies, and evaluating model responses on clean and perturbed inputs. Complementing this data pipeline, we provide additional supervision signals via preference learning and propose a novel Progressive Pre-SFT Alignment strategy featuring two key innovations: a curriculum learning strategy which progressively increases perturbation difficulty to maximize data efficiency; and applying preference optimization before instruction tuning to incentivize fundamental temporal alignment. Extensive experiments demonstrate that our approach consistently improves Video LLM performance across multiple benchmarks with a relatively small set of self-generated DPO data. Our findings highlight TEMPLE as a scalable and efficient complement to SFT-based methods, paving the way for developing reliable Video LLMs.

Code — <https://github.com/lscpku/TEMPLE>

1 Introduction

Video LLMs have made significant strides in recent years, achieving remarkable success in various video understanding tasks (Gemini Team 2024; OpenAI 2024). These advancements follow the standard multi-modal LLM paradigm: First, the model learns visual knowledge and vision-text alignment through large-scale pretraining on image/video-text pairs (Liu et al. 2023b), and then it acquires instruction following capabilities via supervised fine-tuning (SFT) on curated instruction data (Liu et al. 2023a;

Tong et al. 2024; Zhang et al. 2024c). This two-stage approach has enabled impressive improvements in video-language modeling, allowing models to generate coherent, contextually relevant responses conditioned on video input (Fu et al. 2024; Li et al. 2023a; Ren et al. 2023b).

Despite this progress, current Video LLMs still exhibit poor temporal understanding, primarily due to two fundamental limitations. First, constructing high-quality video datasets is inherently challenging due to the complex visual and temporal structures in videos. Existing datasets often suffer from weak video-text correspondence, visual shortcuts, or insufficient temporal reasoning tasks (Lei, Berg, and Bansal 2023; Mangalam, Akshulakov, and Malik 2023; Li et al. 2023b). Second, both pretraining and SFT rely on next-token prediction using ground-truth text, which does not explicitly enforce dynamic temporal understanding. As a result, models frequently neglect subtle yet significant temporal details, over-rely on partial visual or textual cues, and struggle with temporal relationships between events, leading to inconsistent or flawed responses.

To address these limitations, we propose **TEMPLE (TEMporal Preference LEarning)**, a systematic framework that tackles both challenges through: (1) a scalable data pipeline that generates high-quality temporal preference data without relying on external LLMs, and (2) a tailored temporal alignment methodology that goes beyond next-token prediction. TEMPLE is built on Direct Preference Optimization (DPO) (Rafailov et al. 2023), which optimizes the relative quality of model responses by maximizing the reward margin between preferred and non-preferred outputs. Unlike SFT, which requires high-quality, human-verified responses, DPO focuses on relative comparisons rather than absolute correctness, making it particularly effective for video tasks where perfect annotations are difficult to obtain. By learning from temporality-oriented preference pairs, our framework naturally discourages shortcut learning while enabling efficient scaling through automated annotation.

We implement TEMPLE through an automatic preference data pipeline that addresses the scarcity of temporal supervision in video data. Inspired by advances in image-based LLMs (Yu et al. 2024; Zhou et al. 2024b), we generate preference pairs by evaluating model responses under both clean and corrupted visual inputs. Our pipeline consists of three

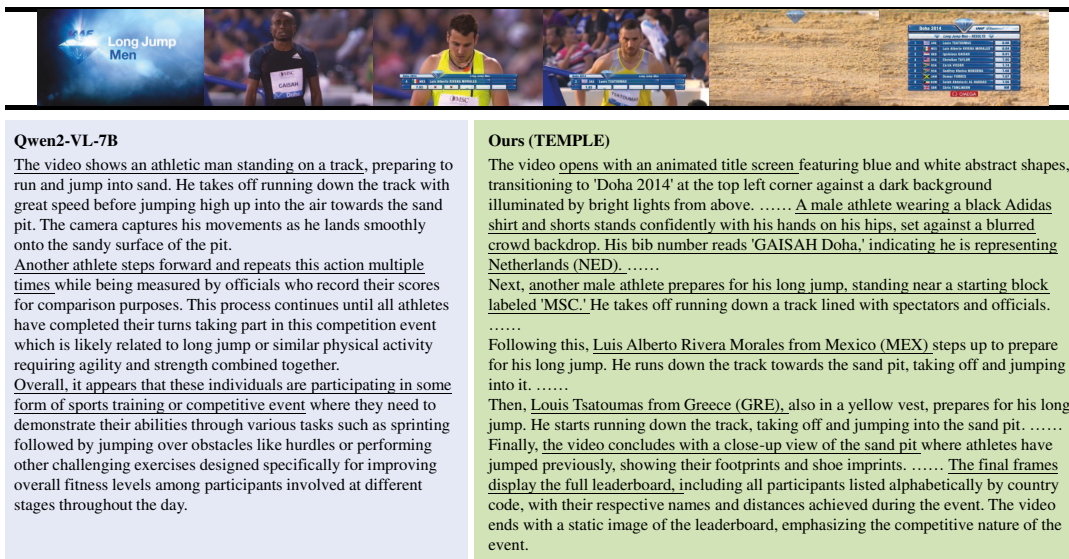


Figure 1: Example of detailed captioning result of the original Qwen2-VL-7B and our approach.

key stages: (1) curating a diverse set of videos that emphasize complex temporal relationships; (2) designing perturbation strategies that precisely target known temporal reasoning weaknesses of current Video LLMs; and (3) constructing preference pairs by contrasting the model’s outputs for clean versus temporally corrupted videos. We also introduce a difficulty factor that quantifies perturbation magnitude to control the sample difficulties. This systematic approach enables us to efficiently generate high-quality contrastive examples that directly highlight temporal understanding capabilities without requiring expensive human annotation.

To leverage the signals in the generated preference data for more effective temporal learning, we further propose a novel Progressive Pre-SFT Alignment strategy tailored for our temporal preference pairs, which consists of two key technical innovations. First, we implement curriculum learning by gradually increasing sample difficulty during preference optimization. This controlled exposure to increasingly challenging temporal contrastive pairs enables more efficient utilization of the preference data. Second, we introduce a novel “Pre-SFT Alignment” strategy, which reverses the conventional order of model alignment. Traditional approaches apply DPO after instruction-tuning (Tunstall et al. 2023), assuming the model has already learned basic capabilities. In contrast, we argue that robust temporal understanding should be established first, before learning to follow diverse instructions. These innovations, coupled with our carefully crafted self-sufficient video filtering and preference generation pipeline, distinguish our work from recent works like Tarsier2 (Yuan et al. 2025) and TPO (Li et al. 2025b), which rely on proprietary LLMs for data filtering and apply DPO in a post-SFT manner.

Through extensive evaluation, we demonstrate that TEMPLE consistently improves model performance across multiple benchmarks including Video-MME (Fu et al. 2024),

MLVU (Zhou et al. 2024a) and Vinoground (Zhang, Mu, and Lee 2024) with only a small number of self-generated DPO pairs. Ablation studies validate the effectiveness of our Progressive Pre-SFT Alignment strategy, which outperforms both the SFT baseline and the conventional SFT-then-DPO approach. Further analysis reveals the transferrability of our DPO data across different model architectures and scales, demonstrating the framework’s flexibility. By introducing this automated, structured preference learning pipeline, we advance the state of Video LLMs and pave the way for more robust and temporally-aware multi-modal models.

2 Method

In this section, we first present a preliminary analysis on the detailed captioning behavior of current Video LLMs, highlighting the need for a more effective alignment strategy. We then introduce our DPO-based approach and propose a three-stage pipeline for automatically generating preference data by comparing model responses on clean and perturbed video inputs. Finally, we present Progressive Pre-SFT Alignment, where we apply DPO using captioning-based preference pairs with increasing difficulty before SFT, enhancing fine-grained temporal alignment in Video LLMs.

2.1 Preliminary Analysis on Detailed Captioning

We conduct a preliminary analysis of the detailed captioning behavior of Qwen2-VL-7B (Wang et al. 2024) by manually annotating events in video keyframes and comparing them to events captured in model-generated captions across 19 videos. The comparison reveals a precision of 74.4% and a recall of merely 35.0%, indicating that the model not only introduces a considerable amount of hallucinated content but also overlooks a substantial portion of events present in the videos (example shown in Figure 1). These findings underscore the critical need for improving Video LLMs’ ability to

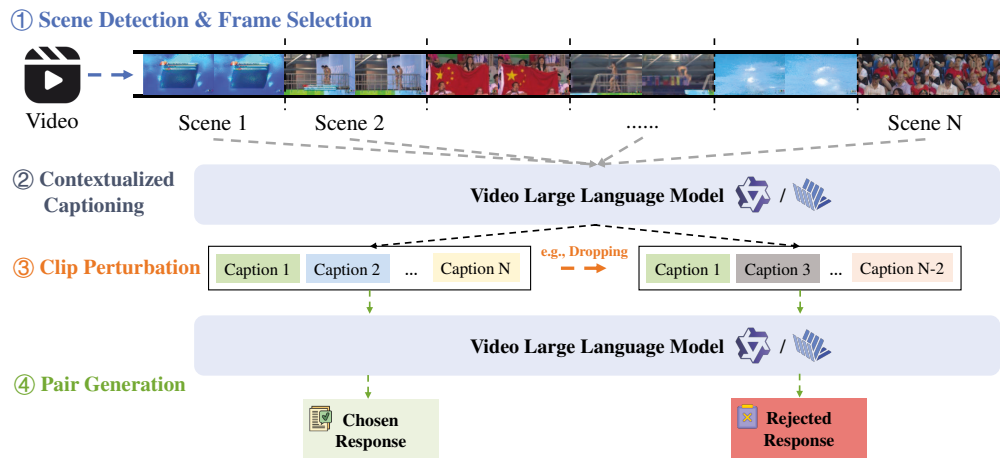


Figure 2: Illustration of our DPO data generation pipeline.

capture and retain temporal information accurately.

2.2 Preference Data Generation Pipeline

Video preference data consists of four key components: video, instruction, chosen response, and rejected response. Inspired by POVID (Zhou et al. 2024b), we generate video DPO data by feeding clean and noised visual inputs into Video LLMs, using the response to the clean input as the chosen response and the response to the corrupted input as the rejected response. To enable preference learning without reliance on external LLMs or APIs, we design a structured pipeline for generating video DPO data, which follows a three-stage process to systematically select videos, craft instructions, and derive response pairs (Figure 2).

Video Selection. Public video datasets often contain relatively static videos with minimal variation in visual content, making them easy for models to process but ineffective for learning temporal structure. Such videos provide limited benefit for DPO and are excluded from our selection process. Instead, we prioritize videos featuring multiple, well-separated events with distinct yet connected visual content. To achieve this, we implement a three-step filtering procedure to identify suitable videos and extract clear keyframes.

In the first step, we perform **scene boundary detection** using TransNetV2 (Soucek and Lokoc 2024), which offers more accurate segmentation than tools like PySceneDetect¹, particularly in handling gradual transitions such as fade-ins and fade-outs. The number of detected scenes serves as an indicator of a video’s temporal complexity, helping us distinguish between dynamic and static content. To refine our selection, we discard monochrome scenes shorter than 0.2 seconds, as these typically correspond to transition frames with little meaningful information. We also filter out videos containing scenes longer than 16 seconds, as such segments often consist of unstructured long shots that are difficult to decompose into discrete events, making them unsuitable for

preference learning. Since current segmentation tools struggle to effectively break down these complex videos, we discard them in favor of those with clearer event boundaries.

In the second step, we further filter out videos that contain frequent scene cuts but little new visual information. This includes cases where multiple shots of the same scene from different angles or repetitive back-and-forth cuts, such as those commonly found in interview footage. To achieve this, we introduce a **similarity grouping** stage, which computes the cosine similarity between middle frames of clips using SigLIP (Zhai et al. 2023) and iteratively group clips with similarity above a predefined threshold. If a video contains too few unique groups, it is deemed to have low information density and is removed. Conversely, if a video contains too many unique groups, it is considered excessively complex, presenting challenges for both computational efficiency and model capability. To balance these factors, we retain only videos with 4 to 32 distinct groups, ensuring sufficient diversity while maintaining manageable complexity.

In the final step, we **extract keyframes** that best represent the visual content of each selected clip. Since we have limited clip length to a maximum of 16 seconds, ensuring that each clip corresponds to a single event, we choose to extract two keyframes per clip, which is generally sufficient to capture its core elements. To improve the clarity of the selected keyframes, we set anchor points at one-third and two-thirds of the clip’s duration and select frames that are both close to these anchors and exhibit high visual clarity. Clarity is assessed using the Laplacian operator (Jain, Kasturi, and Schunck 1995), which measures sharpness to filter out blurry frames. This final step ensures that the extracted keyframes provide distinct and representative snapshots of the video, further enhancing the quality of our DPO data.

Through this three-step process, we systematically filter videos to maximize temporal diversity, eliminate redundant or overly simplistic content, and ensure that selected clips contain clear, well-defined events.

Perturbation Strategy and Task Design. The design of perturbation strategies and task formulation is critical in our

¹<https://github.com/Breakthrough/PySceneDetect>




Figure 3: Illustration of video perturbation strategies controlled by the difficulty factor r .

DPO data construction pipeline, as they directly influence the model behaviors we aim to reinforce. Given that video-based QA datasets often suffer from low quality and unreliable annotations, we opt to use detailed captioning as the sole task for preference learning.

To introduce meaningful preference distinctions, we carefully design perturbation strategies that target the known weaknesses of Video LLMs. One major issue is the model’s tendency to overlook important details in videos or rely on superficial shortcuts in visual or textual inputs. To counter this, we introduce random **clip dropping**, where certain clips are removed from the video before being fed into the model. This forces the model to generate a response with missing information, leading to a rejected response that either omits critical details or makes unreliable inferences based on partial context. This perturbation explicitly penalizes shortcut reliance and encourages the model to develop a more comprehensive understanding of the input.

Another common failure mode is the difficulty in maintaining temporal structure and distinguishing between visually similar but sequentially distinct frames. To address this, we introduce random **clip shuffling** and **clip reversal**, where clips within a video are randomly reordered or flipped temporally. This perturbation preserves the total amount of visual information but disrupts its logical arrangement. Consequently, the model produces a rejected response that misinterprets event chronology or blends together unrelated segments. By optimizing for preference under this perturbation, we encourage the model to develop stronger temporal comprehension and robustness against structural inconsistencies.

These perturbation strategies serve as adversarial challenges, enabling the model to learn more precise and reliable temporal understanding through preference optimization.

Response Generation. Generating high-quality, detailed captions for videos with multiple events is a challenging task for Video LLMs, as they often struggle to retain and integrate information across extended sequences. To mitigate this issue, we draw inspiration from prior work in video captioning and adopt a contextual captioning strategy that enables the model to capture fine-grained temporal dependencies without overwhelming its processing capacity.

Instead of requesting a single caption for the entire video, we employ a **contextualized clip captioning** approach (Chen et al. 2024a). Specifically, at each step, we




Figure 4: SFT loss and gradient norm. Pre-SFT Alignment leads to lower loss and more stable gradients.

provide the model with two consecutive clips (corresponding to four frames) and prompt it to generate a caption for the latter while considering the contextual information from both. This localized approach ensures that captions remain detailed while maintaining continuity across the video.

Once individual clip captions are generated, we aggregate them to form a global video-level caption by feeding all clip captions into the model and prompt it to generate a comprehensive summary of the entire video. This multi-stage process ensures that the final caption preserves both fine-grained event details and high-level narrative coherence.

For response generation in the clean input setting, we directly use the model’s output given the full sequence of clip captions. For perturbed inputs, we modify the clip captions according to our designed perturbations. The model is then prompted with the altered captions to generate a corresponding response, which serves as the rejected response in our preference learning pipeline. By structuring response generation in this manner, we ensure that our preference pairs reflect meaningful variations in video comprehension, reinforcing the model’s ability to process complex temporal relationships and mitigate common failure cases.

2.3 Progressive Pre-SFT Alignment

Progressive Difficulty Scheduling. To systematically examine the impact of perturbing difficulty in DPO, we introduce a difficulty factor r that controls the magnitude of noise applied during preference data generation. This parameter enables us to adjust the degree of disruption in a fine-grained manner, ensuring a structured and progressive learning process for the model.

Given a video with N clips, we regulate difficulty as follows: (i) For clip dropping, we retain only $\lceil N/r \rceil$ clips while discarding the rest, reducing the available visual information. (ii) For clip shuffling and reversal, we partition the clips into $\lceil N/r \rceil$ groups, treating each group as an indivisible unit during perturbation. This ensures that lower values of r introduce finer-grained perturbations while preserving more of the original structure, whereas higher values of r cause more drastic disruptions.

Figure 3 illustrates how the videos are perturbed for different difficulty factors. Intuitively, a smaller r results in weaker perturbations, meaning more information and temporal consistency are preserved, making it harder to differ-

entiate between clean and perturbed inputs. Conversely, a larger r leads to stronger perturbations, making distinctions more apparent and simplifying the preference learning task.

Based on this difficulty factor, we adopt a curriculum learning strategy, where r is gradually decreased throughout training. This allows the model to first learn from easier cases with clearer distinctions and progressively refine its ability to handle more subtle perturbations. By increasing difficulty over time, we ensure that the model develops a more nuanced understanding of video content, leading to improved alignment with fine-grained temporal structures.

Pre-SFT Alignment. In the conventional training paradigm, DPO is typically applied after SFT. This ordering ensures that the model first acquires the fundamental ability to perform tasks and follow instructions before refining its alignment with human preferences. However, considering our objective of improving the model’s fine-grained understanding of video content and its temporal structure, we propose an alternative strategy: performing DPO before SFT. Our key insight is that the primary limitation of Video LLMs lies in their inability to establish precise correspondences between visual input and textual descriptions. Without a strong foundation in fine-grained video-text alignment, the model is prone to relying on static shortcuts and generating responses that are only loosely connected to the true dynamics of the visual content. By applying DPO before SFT, we enable the model to first internalize the intricate relationships between video frames and textual descriptions before moving on to more general instruction-following training. As shown in Figure 4, we observe that models trained with Pre-SFT Alignment exhibit significantly lower loss values and more stable gradient norms during the subsequent SFT phase. This suggests that early-stage temporal alignment enhances the model’s ability to process visual information effectively, leading to improved convergence and overall training stability. By shifting alignment to an earlier stage in training, our approach ensures that Video LLMs develop a robust foundation in temporal understanding, making them better equipped to handle complex multi-modal reasoning tasks.

3 Experiment

3.1 Setting

We conduct our main experiments using the Qwen2-VL model family (Wang et al. 2024), focusing on two model scales: 2B and 7B parameters, due to its strong performance in open-source Video LLMs. We also provide additional evaluation results on Qwen2.5-VL-7B and MiMo-VL-7B in the Appendix.

Data Generation. The videos used for preference data generation are sourced from a subset of the open-domain LLaVA-Video-178K dataset (Zhang et al. 2024c), i.e., Youtube videos within the duration of 1-3 minutes. After our data pipeline, the final data contains 6,385 preference pairs for each difficulty level, resulting in a total of 25,540 training samples.

Training Details. For DPO training, we adopt a four-stage progressive difficulty schedule, where the model is trained sequentially on increasingly challenging data. Beginning with the easiest perturbation setting ($r = 16$), the model gradually adapts to more complex data at $r = 8$, $r = 4$, and finally $r = 2$. For SFT, we use a dataset of 60,000 samples randomly drawn from LLaVA-Video-178K (Zhang et al. 2024c). To ensure efficient training, we apply LoRA (Hu et al. 2022) in both DPO and SFT stages. Detailed hyperparameters can be found in the Appendix.

Evaluation Datasets For evaluation, we assess models trained with and without DPO on four widely used benchmarks. VideoMME (Fu et al. 2024) serves as a comprehensive benchmark for general video understanding; MLVU (Zhou et al. 2024a) focuses on long-video comprehension; Vinoground (Zhang, Mu, and Lee 2024) and Temp-Compass (Liu et al. 2024a) specifically evaluate fine-grained temporal reasoning. For all tasks, we use a maximum pixel per frame of 250,000 and limit the number of frames to 64.

Baselines We validate the effectiveness of our method by comparing its performance against two baselines: the original model without additional tuning, and the model trained directly with SFT. We also report the performance of recent Video LLMs as reference points, including a range of open-source models (Cheng et al. 2024a; Chen et al. 2024b; Li et al. 2025a; Zhang et al. 2024a; Lin et al. 2023b; Li et al. 2024a; Cheng et al. 2024b; Lin et al. 2023a) and closed-source models (Gemini Team 2024; OpenAI 2024).

3.2 Results

Table 1 summarizes the results of our main experiments. On the four benchmarks evaluated, our TEMPLE consistently improves video understanding performance compared to the original model and the model trained solely using SFT. This demonstrates that performing temporal alignment prior to SFT effectively boosts model performance, yielding greater gains compared to the SFT approach. The improvements are particularly pronounced on dimensions related to temporal understanding. On the VideoMME benchmark, TEMPLE delivers substantial gains in temporally-sensitive metrics. For instance, using the Qwen2-VL-7B model, TEMPLE achieves significant improvements over the original model: +3.6 points (70.9 \rightarrow 74.5) in Temporal Perception and +2.8 points (42.4 \rightarrow 45.2) in Temporal Reasoning. Furthermore, on the temporally-focused Vinoground benchmark, TEMPLE enhances performance by 3.2 points, representing a relative increase of nearly 10%. These empirical results collectively validate the effectiveness of the TEMPLE framework in enhancing models’ temporal alignment capabilities and their overall video comprehension performance, leading to consistent gains across diverse benchmarks, especially in temporal understanding tasks.

3.3 Analysis

Ablation Study on Pre-SFT Alignment. To showcase the effectiveness of Pre-SFT Alignment, we conduct an ablation study comparing the performance of multiple different tuning orders. Figure 5 shows the performance differences

Model	# Params	VideoMME (TP/TR)	MLVU	Vinoground_txt	TempCompass_mc	Average
Gemini-1.5-pro	-	75.0	-	35.8	63.9	-
GPT-4o	-	77.2	64.6	54.0	80.0	67.7
VideoLLaMA2	72B	62.4	61.2	36.2	-	-
TimeMarker	7B	57.3	63.9	-	65.8	-
Video-T3	7B	55.0	58.1	31.6	61.2	51.5
LongVA	7B	52.6	56.3	21.2	56.1	46.6
LLaMA-3-VILA1.5	8B	45.9	22.4	-	56.4	-
LLaVA-NeXT-Video	34B	60.2	-	23.0	68.7	-
LLaVA-OneVision	7B	58.2	64.7	41.6	64.8	57.3
MiniCPM-V-2.6	8B	60.9	-	32.6	63.0	-
Video-LLaVA	7B	39.9	47.3	24.8	45.6	39.4
Qwen2-VL-2B	2B	55.2 (54.5/34.5)	57.7	26.4	64.6	51.0
+ SFT	2B	54.6 (65.5/32.8)	60.0	27.4	65.0	51.8
+ TEMPLE (Ours)	2B	55.1 (63.6/35.0)	60.2	28.0	64.6	52.0
Qwen2-VL-7B	7B	60.4 (70.9/42.4)	64.5	34.0	69.6	57.1
+ SFT	7B	60.7 (70.9/44.1)	64.6	34.6	70.8	57.7
+ TEMPLE (Ours)	7B	61.0 (74.5/45.2)	65.8	37.2	71.4	58.9

Table 1: Benchmark accuracy of our approach on VideoMME, MLVU, Vinoground (Text), and TempCompass (Multi-Choice). TP and TR denotes Temporal Perception and Temporal Reasoning, respectively, two subtasks of VideoMME.

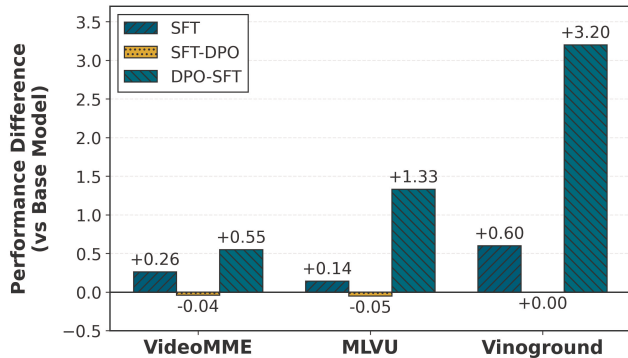


Figure 5: Performance comparison between different tuning strategies on VideoMME, MLVU, and Vinoground (Text).

of Qwen2-VL-7B, comparing its results after tuning with various strategies against the original base model. The Pre-SFT Alignment strategy consistently outperforms other approaches across all three benchmarks. This demonstrates its ability to effectively combine the benefits of both realms: enhancing the model’s temporal understanding while preserving high levels of instruction-following capability. Furthermore, we provide a comparison between our method against previous post-SFT video alignment methods including TPO (Li et al. 2025b) and LLaVA-Hound (Zhang et al. 2024b) in the Appendix, which demonstrates the superiority of our method and the complementarity of both strategies. This ablation study underscores the advantages of aligning the model with fine-grained temporal knowledge before instruction fine-tuning, enabling better overall performance.

Analysis of Progressive Difficulty Scheduling. To investigate the impact of DPO data difficulty on model performance, we first conduct a preliminary comparison by tun-

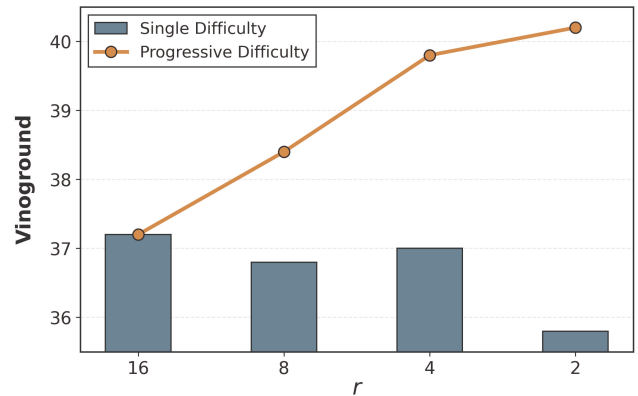


Figure 6: Vinoground (Text) score for models trained with DPO data of different difficulty levels.

ing Qwen2-VL-7B with a single split of DPO data containing the same difficulty level and evaluating the models on Vinoground. As shown by the bars in Figure 6, models tuned with easier DPO data exhibit better performance. This suggests that, due to the initial inability of the model to capture temporal structure, preference optimization with easier data proves more beneficial in the early stages of alignment. This observation motivates our curriculum learning strategy, which formulates DPO as a staged process, progressively increasing the difficulty of DPO data throughout the course. As shown by the line in Figure 6, this progressive tuning approach leads to steady improvements in model performance, validating the efficacy of our strategy in helping the model gradually refine its temporal alignment.

Transferrability across Model Architectures. In our main experiment, we adopt a self-generation approach, us-

	Video	ML	Vinoground		
	MME	VU	Text	Video	Group
Qwen2.5-VL-7B	63.1	64.0	41.6	34.6	17.6
+ SFT	63.1	65.5	42.8	28.8	14.6
+ TEMPLE	63.8	67.3	44.2	25.8	12.4
InternVL2.5-8B	64.4	67.8	42.6	23.8	10.2
+ SFT	64.7	67.4	42.8	24.4	10.2
+ TEMPLE	64.3	67.9	45.0	24.4	10.4

Table 2: Experiment results on the transferrability of our approach across model architectures.

	Video	ML	Vinoground		
	MME	VU	Text	Video	Group
Qwen2-VL-7B	60.4	64.5	34.0	26.6	10.6
+ SFT	60.7	64.6	34.6	26.6	9.6
+ TEMPLE 2B	60.9	64.3	37.4	28.2	12.4
Qwen2-VL-2B	55.2	57.7	26.4	24.6	6.2
+ SFT	54.6	60.0	27.4	26.0	5.4
+ TEMPLE 7B	52.9	59.5	26.0	24.6	6.2

Table 3: Experiment results on the transferrability of our approach across model sizes.

ing the same model for both generating DPO data and performing DPO. To assess the transferrability of our DPO data across different architectures, we conduct an additional experiment where the DPO data generated by Qwen2-VL-7B is used to train two other models: Qwen2.5-VL-7B (Bai et al. 2025) and InternVL2.5-8B (Chen et al. 2024c). The results, presented in Table 2, show that applying DPO on these models using externally generated data still yields improvements over standard SFT on certain metrics, indicating some level of transferability. However, the gains are neither as consistent nor as pronounced as those observed with self-generated data. This suggests that while DPO data can generalize to some extent, aligning the data generation process with the target model leads to more effective learning.

Transferrability across Model Sizes. We further investigate the transferability of our approach across model sizes. Specifically, we evaluate the impact of tuning a larger model using DPO data generated by a smaller model and vice versa. As shown in Table 3, when applying the DPO data from Qwen2-VL-2B to fine-tune Qwen2-VL-7B, the model demonstrates improvements on Vinoground, slight gains on VideoMME, but a minor performance drop on MLVU compared to the SFT-only baseline. In contrast, when tuning Qwen2-VL-2B with DPO data generated by Qwen2-VL-7B, the model consistently underperforms relative to the SFT model. These results emphasize the importance of self-generated DPO data, as it better aligns with the model’s specific capabilities and learning patterns. The findings suggest that while some benefits can be retained when transferring data across sizes, DPO is most effective when the data is tailored to the model’s own behavior and representations.

4 Related Work

Video Large Language Models builds upon early efforts to integrate visual understanding into powerful language models (Liu et al. 2023b,a; Zhu et al. 2023), typically consisting of three key components: a visual encoder, a multi-modal projector, and a language model (Tang et al. 2023). This design enables the model to perform video understanding by generating textual responses conditioned on the video input. Efforts to improve Video LLMs include expanding the scale and improving the quality of video datasets (Zhang et al. 2024c; Chen et al. 2024a,c; Yue et al. 2025), and extending the applicability of Video LLMs by enabling them to handle long-duration videos (Zhang et al. 2024a; Ren et al. 2023b,a), high-resolution content (Wang et al. 2024; Bai et al. 2025; Cheng et al. 2024b; Li et al. 2024a), or streaming video inputs (Lin et al. 2024). In contrast, our work addresses a more fundamental challenge in Video LLMs: the acquisition of temporal knowledge. Prior studies have frequently highlighted the deficiency of these models in capturing fine-grained temporal structure, which we directly tackle through our proposed approach.

Direct Preference Optimization was originally introduced in the text domain to align language model behavior with human preferences. By optimizing an implicit reward model derived from preference data, DPO provides an effective approach to post-training alignment. More recently, research (Zhou et al. 2024b; Liu et al. 2024b; Li et al. 2024b; Zhang 2024) has adapted DPO and other alignment strategies to multi-modal LLMs by developing carefully curated data generation pipelines. In the video domain, some works (Zhang et al. 2024b; Li et al. 2025b; Yuan et al. 2025) have also explored DPO for Video LLMs based on model responses to corrupt visual inputs. However, our approach distinguishes itself through several key improvements. First, our method introduces a carefully designed pipeline that automatically filters temporally rich videos and generates preference data in a fully self-sufficient manner, without relying on external proprietary models. Furthermore, beyond dataset curation, we present a novel Progressive Pre-SFT Alignment strategy for more effective temporal alignment, setting our approach apart from concurrent work in video DPO.

5 Conclusion

In this work, we propose a novel framework to enhance the temporal alignment of Video LLMs through DPO, addressing the lack of supervision signals in current training paradigms. Our method leverages a carefully designed data generation pipeline to curate temporality-oriented video preference data and a novel Progressive Pre-SFT Alignment strategy to inject temporal understanding abilities into Video LLMs. We demonstrate the effectiveness of our approach through extensive experiments on multiple benchmarks, showing substantial gains in temporal understanding. These results open the door for further improvements in Video LLMs, particularly in refining their ability to understand and reason about temporal dynamics in videos.

Acknowledgments

This research was partially supported by the National Natural Science Foundation of China under Grant No. 92470205 and No. 62176002. Xu Sun is the corresponding author.

References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; Zhong, H.; Zhu, Y.; Yang, M.; Li, Z.; Wan, J.; Wang, P.; Ding, W.; Fu, Z.; Xu, Y.; Ye, J.; Zhang, X.; Xie, T.; Cheng, Z.; Zhang, H.; Yang, Z.; Xu, H.; and Lin, J. 2025. Qwen2.5-VL Technical Report. *CoRR*, abs/2502.13923.
- Chen, L.; Wei, X.; Li, J.; Dong, X.; Zhang, P.; Zang, Y.; Chen, Z.; Duan, H.; Lin, B.; Tang, Z.; Yuan, L.; Qiao, Y.; Lin, D.; Zhao, F.; and Wang, J. 2024a. ShareGPT4Video: Improving Video Understanding and Generation with Better Captions. *ArXiv preprint*, abs/2406.04325.
- Chen, S.; Lan, X.; Yuan, Y.; Jie, Z.; and Ma, L. 2024b. TimeMarker: A Versatile Video-LLM for Long and Short Video Understanding with Superior Temporal Localization Ability. *CoRR*, abs/2411.18211.
- Chen, Z.; Wang, W.; Cao, Y.; Liu, Y.; Gao, Z.; Cui, E.; Zhu, J.; Ye, S.; Tian, H.; Liu, Z.; Gu, L.; Wang, X.; Li, Q.; Ren, Y.; Chen, Z.; Luo, J.; Wang, J.; Jiang, T.; Wang, B.; He, C.; Shi, B.; Zhang, X.; Lv, H.; Wang, Y.; Shao, W.; Chu, P.; Tu, Z.; He, T.; Wu, Z.; Deng, H.; Ge, J.; Chen, K.; Dou, M.; Lu, L.; Zhu, X.; Lu, T.; Lin, D.; Qiao, Y.; Dai, J.; and Wang, W. 2024c. Expanding Performance Boundaries of Open-Source Multimodal Models with Model, Data, and Test-Time Scaling. *CoRR*, abs/2412.05271.
- Cheng, Z.; Leng, S.; Zhang, H.; Xin, Y.; Li, X.; Chen, G.; Zhu, Y.; Zhang, W.; Luo, Z.; Zhao, D.; and Bing, L. 2024a. VideoLLaMA 2: Advancing Spatial-Temporal Modeling and Audio Understanding in Video-LLMs. *CoRR*, abs/2406.07476.
- Cheng, Z.; Leng, S.; Zhang, H.; Xin, Y.; Li, X.; Chen, G.; Zhu, Y.; Zhang, W.; Luo, Z.; Zhao, D.; and Bing, L. 2024b. VideoLLaMA 2: Advancing Spatial-Temporal Modeling and Audio Understanding in Video-LLMs. *CoRR*, abs/2406.07476.
- Fu, C.; Dai, Y.; Luo, Y.; Li, L.; Ren, S.; Zhang, R.; Wang, Z.; Zhou, C.; Shen, Y.; Zhang, M.; et al. 2024. Video-MME: The First-Ever Comprehensive Evaluation Benchmark of Multi-modal LLMs in Video Analysis. *ArXiv preprint*, abs/2405.21075.
- Gemini Team. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *ArXiv preprint*, abs/2403.05530.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*.
- Jain, R. C.; Kasturi, R.; and Schunck, B. G. 1995. *Machine vision*. McGraw-Hill. ISBN 978-0-07-032018-5.
- Lei, J.; Berg, T.; and Bansal, M. 2023. Revealing Single Frame Bias for Video-and-Language Learning. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 487–507.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Li, Y.; Liu, Z.; and Li, C. 2024a. LLaVA-OneVision: Easy Visual Task Transfer. *ArXiv preprint*, abs/2408.03326.
- Li, K.; Wang, Y.; He, Y.; Li, Y.; Wang, Y.; Liu, Y.; Wang, Z.; Xu, J.; Chen, G.; Luo, P.; et al. 2023a. Mvbench: A comprehensive multi-modal video understanding benchmark. *ArXiv preprint*.
- Li, L.; Liu, Y.; Yao, L.; Zhang, P.; An, C.; Wang, L.; Sun, X.; Kong, L.; and Liu, Q. 2025a. Temporal Reasoning Transfer from Text to Video. In *ICLR*. OpenReview.net.
- Li, L.; Xie, Z.; Li, M.; Chen, S.; Wang, P.; Chen, L.; Yang, Y.; Wang, B.; Kong, L.; and Liu, Q. 2024b. VLFeedback: A Large-Scale AI Feedback Dataset for Large Vision-Language Models Alignment. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 6227–6246.
- Li, R.; Wang, X.; Zhang, Y.; Wang, Z.; and Yeung-Levy, S. 2025b. Temporal Preference Optimization for Long-Form Video Understanding. *CoRR*, abs/2501.13919.
- Li, S.; Li, L.; Ren, S.; Liu, Y.; Liu, Y.; Gao, R.; Sun, X.; and Hou, L. 2023b. VITATECS: A Diagnostic Dataset for Temporal Concept Understanding of Video-Language Models. *ArXiv preprint*, abs/2311.17404.
- Lin, B.; Zhu, B.; Ye, Y.; Ning, M.; Jin, P.; and Yuan, L. 2023a. Video-LLaVA: Learning United Visual Representation by Alignment Before Projection. *ArXiv preprint*, abs/2311.10122.
- Lin, J.; Fang, Z.; Chen, C.; Wan, Z.; Luo, F.; Li, P.; Liu, Y.; and Sun, M. 2024. StreamingBench: Assessing the Gap for MLLMs to Achieve Streaming Video Understanding. *CoRR*, abs/2411.03628.
- Lin, J.; Yin, H.; Ping, W.; Lu, Y.; Molchanov, P.; Tao, A.; Mao, H.; Kautz, J.; Shoeybi, M.; and Han, S. 2023b. VILA: On Pre-training for Visual Language Models. *arXiv:2312.07533*.
- Liu, H.; Li, C.; Li, Y.; and Lee, Y. J. 2023a. Improved Baselines with Visual Instruction Tuning.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023b. Visual Instruction Tuning. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Liu, Y.; Li, S.; Liu, Y.; Wang, Y.; Ren, S.; Li, L.; Chen, S.; Sun, X.; and Hou, L. 2024a. TempCompass: Do Video LLMs Really Understand Videos? In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics ACL 2024*, 8731–8772.
- Liu, Z.; Zang, Y.; Dong, X.; Zhang, P.; Cao, Y.; Duan, H.; He, C.; Xiong, Y.; Lin, D.; and Wang, J. 2024b. MIA-DPO: Multi-Image Augmented Direct Preference Optimization For Large Vision-Language Models. *CoRR*, abs/2410.17637.
- Mangalam, K.; Akshulakov, R.; and Malik, J. 2023. EgoSchema: A Diagnostic Benchmark for Very Long-form Video Language Understanding. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- OpenAI. 2024. GPT-4o System Card.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization: Your Language Model is Secretly a Reward Model. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Ren, S.; Chen, S.; Li, S.; Sun, X.; and Hou, L. 2023a. TESTA: Temporal-Spatial Token Aggregation for Long-form Video-Language Understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, 932–947. Singapore: Association for Computational Linguistics.

- Ren, S.; Yao, L.; Li, S.; Sun, X.; and Hou, L. 2023b. TimeChat: A Time-sensitive Multimodal Large Language Model for Long Video Understanding. *ArXiv preprint*, abs/2312.02051.
- Soucek, T.; and Lokoc, J. 2024. TransNet V2: An Effective Deep Network Architecture for Fast Shot Transition Detection. In *ACM Multimedia*, 11218–11221. ACM.
- Tang, Y.; Bi, J.; Xu, S.; Song, L.; Liang, S.; Wang, T.; Zhang, D.; An, J.; Lin, J.; Zhu, R.; et al. 2023. Video understanding with large language models: A survey. *ArXiv preprint*, abs/2312.17432.
- Tong, S.; Brown, E.; Wu, P.; Woo, S.; Middepogu, M.; Akula, S. C.; Yang, J.; Yang, S.; Iyer, A.; Pan, X.; Wang, A.; Fergus, R.; LeCun, Y.; and Xie, S. 2024. Cambrian-1: A Fully Open, Vision-Centric Exploration of Multimodal LLMs. *arXiv preprint arXiv:2406.16860*.
- Tunstall, L.; Beeching, E.; Lambert, N.; Rajani, N.; Rasul, K.; Belkada, Y.; Huang, S.; von Werra, L.; Fourrier, C.; Habib, N.; et al. 2023. Zephyr: Direct distillation of lm alignment. *ArXiv preprint*, abs/2310.16944.
- Wang, P.; Bai, S.; Tan, S.; Wang, S.; Fan, Z.; Bai, J.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Fan, Y.; Dang, K.; Du, M.; Ren, X.; Men, R.; Liu, D.; Zhou, C.; Zhou, J.; and Lin, J. 2024. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution. *arXiv preprint arXiv:2409.12191*.
- Yu, T.; Zhang, H.; Yao, Y.; Dang, Y.; Chen, D.; Lu, X.; Cui, G.; He, T.; Liu, Z.; Chua, T.-S.; et al. 2024. RLAIIF-V: Aligning MLLMs through Open-Source AI Feedback for Super GPT-4V Trustworthiness. *ArXiv preprint*, abs/2405.17220.
- Yuan, L.; Wang, J.; Sun, H.; Zhang, Y.; and Lin, Y. 2025. Tarsier2: Advancing Large Vision-Language Models from Detailed Video Description to Comprehensive Video Understanding. *CoRR*, abs/2501.07888.
- Yue, Z.; Lin, Z.; Song, Y.; Wang, W.; Ren, S.; Gu, S.; Li, S.; Li, P.; Zhao, L.; Li, L.; Bao, K.; Tian, H.; Zhang, H.; Wang, X.; Zhu, D.; Cici, He, C.; Ye, B.; Shen, B.; Zhang, Z.; Jiang, Z.; Zheng, Z.; Song, Z.; Luo, Z.; Yu, Y.; Wang, Y.; Tian, Y.; Tu, Y.; Yan, Y.; Huang, Y.; Wang, X.; Xu, X.; Song, X.; Zhang, X.; Yong, X.; Zhang, X.; Deng, X.; Yang, W.; Ma, W.; Lv, W.; Zhuang, W.; Liu, W.; Deng, S.; Liu, S.; Chen, S.; Yu, S.; Liu, S.; Wang, S.; Ma, R.; Wang, Q.; Wang, P.; Chen, N.; Zhu, M.; Zhou, K.; Zhou, K.; Fang, K.; Shi, J.; Dong, J.; Xiao, J.; Xu, J.; Liu, H.; Xu, H.; Qu, H.; Zhao, H.; Lv, H.; Wang, G.; Zhang, D.; Zhang, D.; Zhang, D.; Ma, C.; Liu, C.; Cai, C.; and Xia, B. 2025. MiMo-VL Technical Report. *CoRR*, abs/2506.03569.
- Zhai, X.; Mustafa, B.; Kolesnikov, A.; and Beyer, L. 2023. Sigmoid Loss for Language Image Pre-Training. In *ICCV*, 11941–11952. IEEE.
- Zhang, G. 2024. VL-RLHF: A RLHF Infrastructure for Vision-Language Model. <https://github.com/TideDra/VL-RLHF>.
- Zhang, J.; Mu, C.; and Lee, Y. J. 2024. Vinoground: Scrutinizing LMMs over Dense Temporal Reasoning with Short Videos. *arXiv preprint arXiv:2410.02763*.
- Zhang, P.; Zhang, K.; Li, B.; Zeng, G.; Yang, J.; Zhang, Y.; Wang, Z.; Tan, H.; Li, C.; and Liu, Z. 2024a. Long Context Transfer from Language to Vision. *ArXiv preprint*, abs/2406.16852.
- Zhang, R.; Gui, L.; Sun, Z.; Feng, Y.; Xu, K.; Zhang, Y.; Fu, D.; Li, C.; Hauptmann, A.; Bisk, Y.; and Yang, Y. 2024b. Direct Preference Optimization of Video Large Multimodal Models from Language Model Reward. *ArXiv preprint*, abs/2404.01258.
- Zhang, Y.; Wu, J.; Li, W.; Li, B.; Ma, Z.; Liu, Z.; and Li, C. 2024c. Video Instruction Tuning With Synthetic Data. *arXiv:2410.02713*.
- Zhou, J.; Shu, Y.; Zhao, B.; Wu, B.; Xiao, S.; Yang, X.; Xiong, Y.; Zhang, B.; Huang, T.; and Liu, Z. 2024a. MLVU: A Comprehensive Benchmark for Multi-Task Long Video Understanding. *ArXiv preprint*, abs/2406.04264.
- Zhou, Y.; Cui, C.; Rafailov, R.; Finn, C.; and Yao, H. 2024b. Aligning Modalities in Vision Large Language Models via Preference Fine-tuning. *arXiv:2402.11411*.
- Zhu, D.; Chen, J.; Shen, X.; Li, X.; and Elhoseiny, M. 2023. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. *ArXiv preprint*, abs/2304.10592.