

TIME: Temporal-Sensitive Multi-Dimensional Instruction Tuning and Robust Benchmarking for Video-LLMs

Yunxiao Wang^{1*}, Meng Liu^{2†}, Wenqi Liu¹, Xuemeng Song³, Bin Wen^{4†}, Fan Yang⁴, Tingting Gao⁴, Di Zhang⁴, Guorui Zhou⁴, Liqiang Nie¹

¹Shandong University, Jinan, China

²Shandong Jianzhu University, Jinan, China

³Southern University of Science and Technology, Shenzhen, China

⁴Kuaishou Technology, Beijing, China

{yunxiao.wang, liuwq}@mail.sdu.edu.cn, {mengliu.sdu, sxmustc, nieliqiang}@gmail.com, {wenbin, yangfan, lisize, zhangdi08, zhouguorui}@kuaishou.com

Abstract

Video large language models have achieved remarkable performance in tasks such as video question answering, however, their temporal understanding remains suboptimal. To address this limitation, we curate a dedicated instruction fine-tuning dataset that focuses on enhancing temporal comprehension across five key dimensions. In order to reduce reliance on costly temporal annotations, we introduce a multi-task prompt fine-tuning approach that seamlessly integrates temporal-sensitive tasks into existing instruction datasets without requiring additional annotations. Furthermore, we develop a novel benchmark for temporal-sensitive video understanding that not only fills the gaps in dimension coverage left by existing benchmarks but also rigorously filters out potential shortcuts, ensuring a more accurate evaluation. Extensive experimental results demonstrate that our approach significantly enhances the temporal understanding of video-LLMs while avoiding reliance on shortcuts.

Introduction

Recent advancements in video large language models (video-LLMs) (Tang et al. 2025; Cheng et al. 2025a) have demonstrated significant capabilities in video understanding and reasoning. In contrast to image large language models (image-LLMs) (Keye 2025; Lv et al. 2025a), which focus on static visual content analysis, video-LLMs must capture dynamic visual information across frames. This demands robust temporal understanding, a critical research challenge that has garnered substantial attention in recent work (Cheng et al. 2025b; Song et al. 2025).

Despite efforts to enhance temporal understanding in video-LLMs (Liu et al. 2024a; Li et al. 2025b), recent studies indicate that significant challenges remain in tasks requiring advanced temporal reasoning (Wang et al. 2024b; Xiao et al. 2025). These limitations can be attributed to

*Work done during an internship at Kuaishou Technology.

†Corresponding Authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

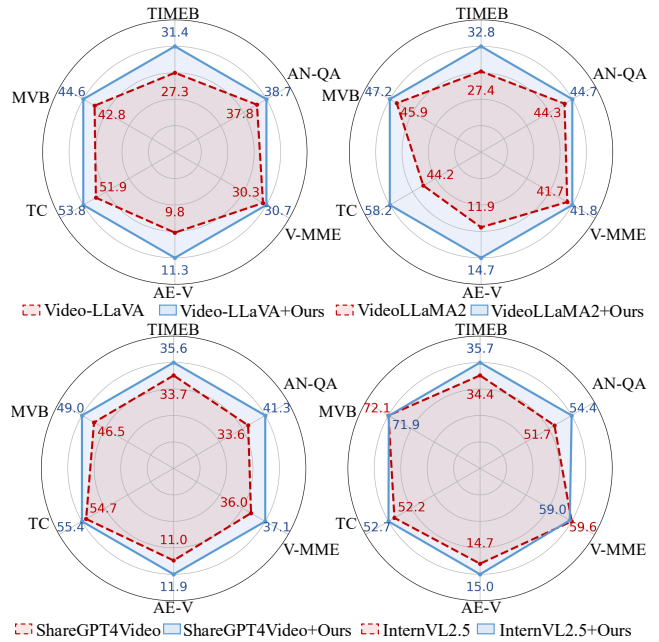


Figure 1: Accuracy of four video-LLMs evaluated across six benchmarks. “+ Ours” suffix denotes models fine-tuned with our approach, achieving substantial improvements on all benchmarks. TIMEBench (TIMEB), MVBench (MVB), TempCompass (TC), and AutoEval-Video (AE-V) focus on temporal scenarios, whereas Video-MME (V-MME) and ActivityNet-QA (AN-QA) evaluate general performance.

the following primary factors: **1) Insufficient temporal instruction fine-tuning data.** Current video instruction tuning datasets (Lin et al. 2024; Chen et al. 2024a) prioritize generalization across common scenarios rather than in-depth temporal comprehension. Although some approaches, such as AVicuna (Tang et al. 2024), generate temporal instruction tuning datasets through video splicing, their applicability remains limited to temporal localization due to the inherent constraints of this method. **2) Exploitation of data**

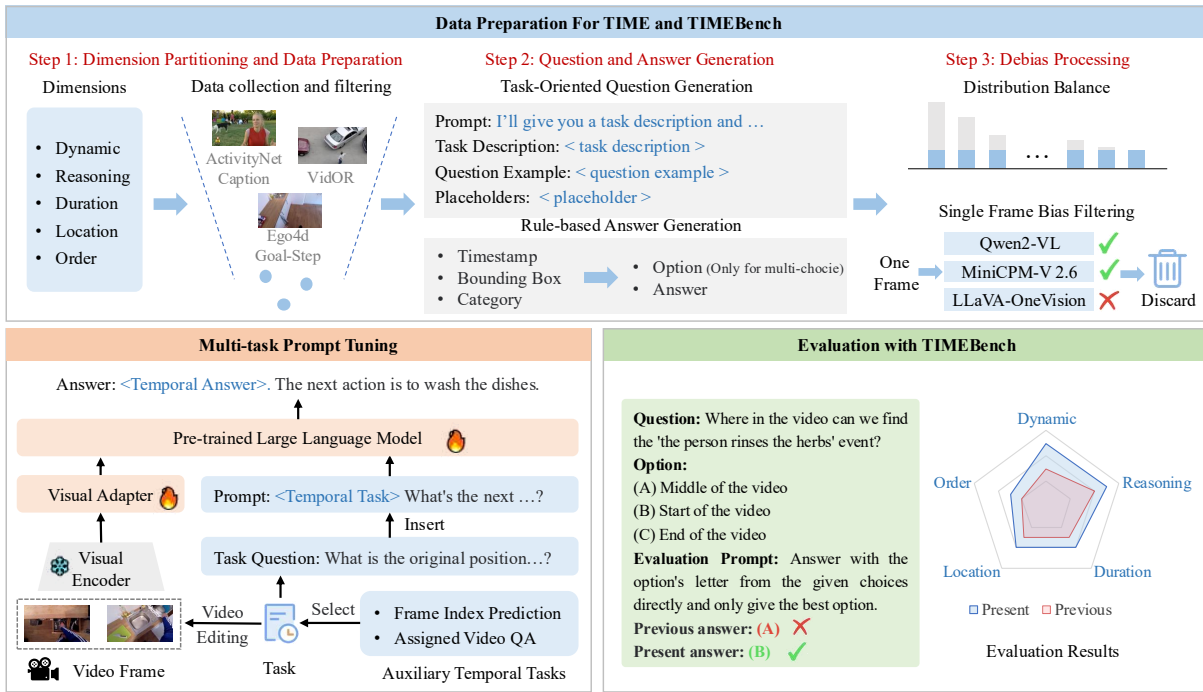


Figure 2: Our framework consists of three components: the TIME instruction tuning dataset, the multi-task prompt tuning method, and the TIMEBench video temporal understanding benchmark. Top panel: the data preparation for TIME and TIMEBench. Bottom left: auxiliary temporal tasks are randomly mixed with the original prompt during multi-task prompt tuning, enabling joint training. Bottom right: TIMEBench uses multiple-choice questions for objective evaluation

shortcuts. For instance, in determining a child’s movement direction (e.g., the *Dynamic* dimension in Figure 3), models may rely on facial orientation instead of conducting a genuine temporal analysis of positional changes, leading to errors when videos are reversed. Notably, temporal benchmarks (Li et al. 2024; Shangguan et al. 2025) are also affected by such shortcuts. Our experimental results (Figure 5) show that performance based on single random frames significantly surpasses random baselines, indicating that many questions can be resolved without true temporal reasoning, thereby overestimating model capabilities.

To address these challenges, we systematically enhance Video-LLMs’ temporal understanding and evaluation through the following improvements. First, **we construct a Temporal-sensitive Multi-dimensional (TIME) instruction-tuning dataset** comprising 34,000 carefully curated samples across five key temporal dimensions (including *Dynamic*, *Reasoning*, *Duration*, *Location*, and *Order*), which serves as a valuable complement to existing general-purpose instruction-tuning datasets. These dimensions capture essential aspects of temporal reasoning, such as the dynamics of events over time and the capacity to predict future occurrences. Simultaneously, we implement rigorous data filtering to remove any potential data shortcuts that might undermine the effectiveness of model tuning. Second, to further mitigate limitations in data volume, **we propose a Multi-Task Prompt tuning (MTP) framework** that augments standard instruction tuning with two types of

temporal-oriented self-supervised tasks: frame-level task to improve single-frame temporal understanding, and segment-level task to enhance cross-segment temporal understanding. As shown in Figure 1, fine-tuning with our approach results in substantial performance gains across most temporal benchmarks for the evaluated video-LLMs. Finally, **we develop a robust Temporal-sensitive Multi-dimensional Benchmark (TIMEBench)**, which covers five dimensions of temporal evaluation and provides a more accurate measure of temporal understanding ability by filtering out potential data shortcuts through a carefully designed filtering mechanism. Experimental results validate the rigor of our evaluation benchmark.

In summation, our contributions are as follows:

- We construct a temporal-sensitive multi-dimensional instruction-tuning dataset intended to augment existing general instruction-tuning corpora, incorporating curated samples that specifically address temporal reasoning challenges in video data.
- We introduce a novel multi-task prompt-tuning approach that enhances existing instruction fine-tuning processes with two temporal-oriented self-supervised tasks.
- We propose a robust temporal-sensitive multi-dimensional benchmark that provides a more accurate measure of temporal understanding ability by filtering out potential data shortcuts.



Figure 3: Examples from the TIME dataset and the TIMEBench benchmark with data distribution across task types. These examples encompass five tasks that target the *Dynamic*, *Reasoning*, *Duration*, *Location*, and *Order* dimensions of temporal understanding. Ground-truth answers are highlighted in orange.

Temporal Sensitive Instruction Tuning

In this section, we present our approach to enhancing temporal understanding in video-LLMs through two key components: the construction of a temporally-sensitive instruction-tuning dataset and the development of a multi-task prompt fine-tuning strategy. These components are designed to augment the model’s capacity for temporal reasoning across a wide range of video-related tasks.

TIME Instruction Tuning Dataset

The limited performance of existing video-LLMs on temporally sensitive tasks can be largely attributed to the absence of a dedicated instruction-tuning dataset tailored for temporal contexts. To address this gap, we construct the TIME instruction-tuning dataset, aimed at enhancing the model’s temporal comprehension capabilities.

Dimension and Data Selection To improve temporal comprehension in video-LLMs, we draw insights from related fields such as video moment localization (Liu et al. 2018a,b) and action anticipation (Hu et al. 2022). We identify tasks that emphasize temporal understanding and categorize temporal reasoning into five key dimensions, as shown in Figure 3. Each dimension is paired with a task designed to strengthen the model’s temporal capabilities.

Dynamic. This dimension evaluates the model’s ability to capture dynamic changes over time, specifically focusing on detecting the direction of object movement. Each video is cropped to highlight the unidirectional movement of a target object. Data is collected from the VidOR (Shang et al.

2019) dataset, originally designed for scene graph generation. To ensure unique identification, video segments containing more than two objects of the same category in a single frame are filtered out. The direction of the object’s movement is determined by calculating the position of the center point of its bounding box.

Reasoning. This dimension evaluates the model’s capacity to forecast future events from observed temporal sequences. Specifically, we focus on action anticipation, where the model predicts the subsequent action based on preceding actions. For this task, we leverage the Ego4D Goal-Step (Song et al. 2023) dataset, which provides annotations for goal-oriented hierarchical activities. In our framework, each step in the sequence is treated as an anticipatory action, while the earlier steps form the observation sequence. To ensure high-quality data, we filter out sequences with fewer than three steps and truncate those exceeding 15 steps. Additionally, only sequences with at least 50% of the steps marked as “essential” are retained to mitigate noise.

Duration. This dimension evaluates the model’s ability to perceive and estimate the duration of events in unedited videos. We challenge the model to determine the length of specific events by categorizing them into one of three segments based on their relative duration within the video. For this task, we use a subset of ActivityNet Captions (Krishna et al. 2017) where events are divided into three parts according to their temporal span, and the model must correctly identify the appropriate category for each event.

Location. This dimension measures the model’s capability

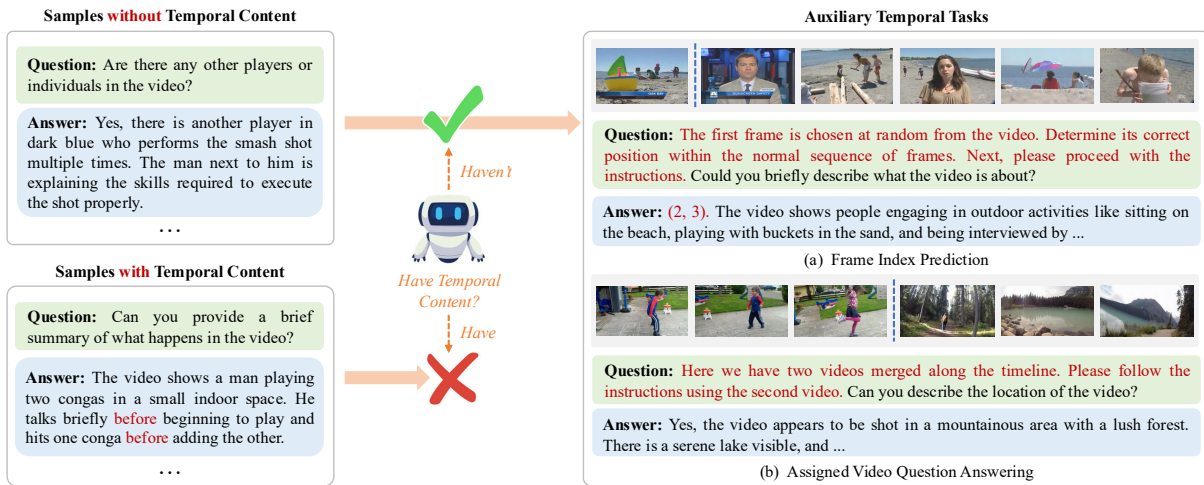


Figure 4: Illustration of our multi-task prompt tuning approach. The LLM first checks if a sample contains a temporal description. Only samples lacking such descriptions undergo auxiliary tasks. For frame index prediction task, a frame is randomly sampled from the sequence, and the model must restore its original position. For assigned videoQA task, a randomly selected video is concatenated with the original, and the model needs to answer the original question according to the correct video.

to pinpoint precise temporal segments within a video by determining when a given event occurs. To construct the dataset for this task, we extract a non-overlapping subset from ActivityNet Captions (Krishna et al. 2017) and uniformly partition each video into three intervals: start, middle, and end. Only activities that both begin and end within the same interval are selected, ensuring the model concentrates on accurately identifying temporal boundaries.

Order. This dimension assesses the model’s ability to understand temporal event sequences by ordering actions accurately. Specifically, the model must arrange multiple non-overlapping actions in the sequence in which they occur. For this task, we use the Charades (Sigurdsson et al. 2016) dataset, partitioning each video into segments comprising three distinct actions performed sequentially. To ensure sequence diversity, segments with identical actions are filtered out so that each sequence consists of unique action types. To ensure sufficient adequate context, we filter out video clips with very short target activities. This step reduces the likelihood of the model erring due to insufficient observational context rather than a lack of temporal understanding.

Question and Answer Generation Since the original datasets lack pre-existing QA pairs, we generate them to build a comprehensive QA dataset. To maximize diversity, we partition the dataset into two parts: open-ended and multiple-choice QA.

Open-ended QA. 1) *Question Generation:* For each task, we manually create an example question and then use ChatGPT (OpenAI 2024) to generate 10 variations based on the task description and example. In questions of *Dynamic*, *Duration*, *Location* and *Reasoning* dimensions, placeholders are inserted and later replaced with specific target objects, activities, or goal descriptions. 2) *Answer Generation:* For *Dynamic* and *Reasoning*, the correct answer typically reflects the direction of movement or the action category.

For *Duration* and *Location*, answers are derived from event-associated timestamps, while for *Order*, the answer is the chronological arrangement of three actions.

Multiple-choice QA. 1) *Question Generation:* We follow the similar process as for open-ended QA, but incorporate additional instructions and answer options. 2) *Instruction Generation:* For each task, we generate 10 extra instructional prompts using ChatGPT (OpenAI 2024). These prompts, inserted before the question, guide video-LLMs to select the correct option from a list of candidates. 3) *Option Generation:* For *Dynamic*, *Duration*, and *Location*, all possible answers are provided as options. For *Order*, random permutations of actions are used to avoid shortcuts based on visual or keyword co-occurrence. For *Reasoning*, incorrect options are sampled from other steps within the same goal to minimize dependence on co-occurrence relationships.

Dataset Debiasing Prior studies (Tong et al. 2024; Yu et al. 2024; Huang et al. 2024) have shown that multimodal language models are prone to biases arising from heuristics such as keyword co-occurrence and inherent distributional patterns in training data, which can lead to hallucinatory behavior in video-LLMs. For example, Otani et al. (Otani et al. 2020) observed that many video moment location datasets exhibit a temporal bias, with events predominantly occurring at the beginning of the video. Consequently, models may over-predict early events, achieving high accuracy on biased test sets without genuine temporal comprehension.

To mitigate these biases and foster a true understanding of temporal events, we implement the following debiasing strategies: 1) We initially conduct a manual verification of the content generated by ChatGPT (OpenAI 2024) to ensure its accuracy. Then, we ensure a balanced distribution of answers across both open-ended and multiple-choice QA formats. Specifically, for *Dynamic*, *Duration*, and *Location*, we maintain roughly equal counts of each answer type to

prevent skewed representations. 2) For *Reasoning* and *Order*, we downsample frequently occurring actions to alleviate long-tailed distributions. 3) For *Dynamic*, we incorporate a reversed version of the video to discourage reliance on visual shortcuts, such as face or vehicle orientation. And 4) we carefully balance the distribution of questions and options across all dimensions.

Multi-task Prompt Tuning

Relying solely on fully annotated datasets for instruction tuning is inherently limited by annotation availability. While existing work (Zhang et al. 2025; Zhu et al. 2025; Lv et al. 2025b) has introduced unsupervised tasks such as masked frame modeling to improve temporal understanding, these approaches are not directly applicable to fine-tuning video-LLMs. To overcome this limitation, we propose a multi-task prompt-tuning approach that incorporates two auxiliary tasks into existing instruction-tuning datasets without requiring additional annotations.

Frame Index Prediction. In this task, we randomly sample a frame from the original video and position it at the beginning of the sequence. A prompt is inserted before the original question to guide the model in predicting the correct position of the displaced frame. The modified question-answer structure is illustrated in the upper-right of Figure 4.

Assigned VideoQA. This task trains the model to identify the relevant video segments required to answer a given question. Specifically, a randomly selected video from the dataset is concatenated with the original video in randomized order, converting the task into a binary video segment localization problem with only two possible choices. Explicit prompt instructions are provided to help the model discern the correct video, without dictating the proportion of the two videos. This setup is shown in the bottom-right of Figure 4.

In both tasks, the temporal content of the original video is minimally altered to ensure that the model can still accurately answer the original question. Moreover, as depicted in Figure 4, we employ Qwen2.5-72B (Team 2024), an open-source LLM with capabilities comparable to ChatGPT (OpenAI 2024), to identify QA pairs that involve temporal content. We refrain from adding new tasks to these pairs to avoid introducing errors from video modification. For each auxiliary task, 10 prompts are generated using ChatGPT (OpenAI 2024). Following the workflow in Figure 2, we tune the video-LLMs with these auxiliary tasks. Specifically, for samples without temporal content, one auxiliary task is randomly selected, the video and QA pair are modified accordingly, and then the video-LLMs are tuned on the modified dataset. This approach opens a new avenue for unsupervised instruction fine-tuning aimed at enhancing the temporal comprehension abilities of LLMs.

TIMEBench

In this section, we introduce TIMEBench, a novel benchmark designed to evaluate temporal understanding in video models, and compare it with existing benchmarks. TIMEBench comprehensively assesses the temporal reasoning capabilities of video models across five dimensions: *Dy-*




Figure 5: Performance comparison of VideoL-LaMA 2 (Cheng et al. 2024) on existing Benchmarks. Acc_s represents the predictions made using only a single frame, Acc_b denotes the accuracy when visual information is excluded, and Acc_r indicates random predictions.

ynamic, Reasoning, Duration, Location, and Order. Following the pipeline illustrated in Figure 2, we systematically collect and filter data from diverse sources and generate multiple-choice questions for objective evaluation. This design not only mitigates data shortcuts but also extends the evaluation to a broader range of temporal aspects, setting TIMEBench apart from existing benchmarks.

Benchmark Construction

As discussed previously, existing benchmarks often fall short of covering the five dimensions outlined in Figure 3 and are susceptible to shortcut issues. To overcome these shortcomings, we develop a benchmark that follows the data preparation and processing pipeline outlined in Figure 2. The key differences are described below.

Data Preparation. We leverage diverse data sources to construct TIMEBench, ensuring robust evaluation on out-of-domain data. For *Reasoning*, we exclusively use the subset of EgoPlan (Chen et al. 2024b) sourced from Epic-Kitchens (Damen et al. 2022) to avoid domain overlap. For *Order*, we use coarse-level annotations of Breakfast (Kuehne, Arslan, and Serre 2014) to avoid overly fine-grained action partitioning, which could impede the model’s ability to distinguish actions. For *Location* and *Duration*, we respectively use TACoS (Regneri et al. 2013) and QVHighlights (Lei, Berg, and Bansal 2021) datasets. We apply random video cropping and generate question-answer options based on the new timestamps, thereby minimizing reliance on prior knowledge or fixed temporal context. For *Dynamic*, we use GOT-10k (Huang, Zhao, and Huang 2021) and flip the videos to prevent data shortcuts.

QA generation. Following the pipeline in Figure 2, we construct QA pairs with an emphasis on multiple-choice formats to ensure objective evaluation. The evaluation prompt is defined as: “Answer with the option’s letter from the given choices directly and only give the best option”. A complete example is provided in the bottom right of Figure 2.

Benchmark Debiasing. Prior work (Tong et al. 2024) as demonstrated that current models often exploit data shortcuts rather than exhibiting true temporal comprehension. To address this, we utilize three advanced open-source multi-modal LLMs, i.e., Qwen2-VL (Wang et al. 2024a), LLaVA-

Model	TIMEBench						MVB	TC	AE-V	V-MME	AN-QA
	LO	DU	DY	OR	RE	AVG					
Random	33.3	33.3	25.0	25.0	25.0	28.3	27.3	30.0	0.0	27.2	0.0
Video-LLaVA + Ours	32.7 37.0	32.3 34.3	22.8 25.5	26.5 34.3	24.8 28.3	27.3 31.4	42.8 44.6	51.9 53.8	9.8 11.3	30.3 30.7	37.8 38.7
VideoLLaMA 2 + Ours	32.3 33.3	30.3 43.0	22.8 28.0	28.0 33.0	24.0 26.8	27.4 32.8	45.9 47.2	44.2 58.2	11.9 14.7	41.7 41.8	44.3 44.7
ShareGPT4Video + Ours	33.7 38.3	46.7 47.3	31.0 31.7	33.8 34.5	26.5 29.5	33.7 35.6	46.5 49.0	54.7 55.4	11.0 11.9	36.0 37.1	33.6 41.3
InternVL 2.5 + Ours	30.3 30.6	45.7 48.3	29.3 29.8	42.0 42.8	26.8 27.0	34.4 35.7	72.1 71.9	52.2 52.7	14.7 15.0	59.6 59.0	51.7 54.4

Table 1: Performance comparison across six benchmarks. The best performance for each part is highlighted in **bold**. TIMEBench, MVB, TC and AE-V focus on temporal scenarios, whereas V-MME and AN-QA address general scenario. TIMEBench, MVB and V-MME are multiple-choice forms, AE-V and AN-QA are open-ended forms, and TC contains both multiple-choice and open-ended forms.

OneVision (Li et al. 2025a), and MiniCPM-V 2.6 (Yao et al. 2024), as judge agents to evaluate the validity of QA pairs. Specifically, each question is based on a randomly selected frame from the video, and the models are tasked with answering using only that single frame. Given the limited temporal comprehension capabilities of current video-LLMs (Wang et al. 2024b; Li et al. 2024), we adopt a majority voting approach. If two or more judge models answer correctly based solely on a single sampled frame, this indicates potential shortcut reliance, and the corresponding QA pair is filtered out. The final benchmark is constructed from the remaining valid pairs, and resampling is performed to ensure a balanced distribution of answer options.

Benchmark Comparison

To ensure the objectivity of our evaluation, we carefully debias the construction of TIMEBench. To assess the effectiveness of this approach, we conducted an experiment (see Figure 5) comparing the impact of potential biases across various benchmarks. In this experiment, we use the random prediction accuracy Acc_r as a baseline and report Acc_s and Acc_b for comparison. Acc_s represents the accuracy of VideoLLaMA 2 (Cheng et al. 2024) using a single video frame as input, while Acc_b corresponds to the blind setting where videos are replaced with fully black images.

The results demonstrate that on MVBench (Li et al. 2024), TempCompass (Liu et al. 2024b), and AutoEval-Video (Chen et al. 2023), the Acc_b consistently surpasses Acc_r even in the absence of visual information. Moreover, with the addition of just one frame from the videos, Acc_s on these benchmarks significantly exceeds both Acc_r and Acc_b , indicating that many questions in these benchmarks can be answered without relying on temporal information. In contrast, on TIMEBench, both Acc_s and Acc_b fall below Acc_r , with a small gap between them, indicating that TIMEBench is less prone to shortcut biases and offers a more rigorous evaluation of temporal understanding. It can be attributed to our meticulous data filtering and debiasing procedures.

Experiments

Experiment Settings. For evaluation, we used TIMEBench and five additional benchmarks: MVBench (Li et al. 2024), TempCompass (Liu et al. 2024b) and AutoEval-Video (Chen et al. 2023), which are tailored for video temporal understanding, as well as Video-MME (Fu et al. 2025) and ActivityNet-QA (Yu et al. 2019), which serve as general benchmarks for video comprehension. To validate the effectiveness of our method, we fine-tuned Video-LLaVA (Lin et al. 2024), VideoLLaMA 2 (Cheng et al. 2024), ShareGPT4Video (Chen et al. 2024a) and InternVL 2.5 (Chen et al. 2024c) using our TIME dataset and our MTP approach¹. Following the original model settings, the entire parameters of the LLMs were fine-tuned for Video-LLaVA and VideoLLaMA 2, whereas LoRA was employed for ShareGPT4Video and InternVL 2.5. Except for ShareGPT4Video, the visual encoders were kept frozen in all other models.

Performance Comparison. As illustrated in Table 1, fine-tuning with our method leads to significant improvements for all video-LLMs across most benchmarks, especially on the four benchmarks that specifically assess video temporal understanding. Simultaneously, performance on the two general benchmarks remains stable or shows slight improvements, indicating that our method enhances temporal reasoning without compromising overall performance. Notably, the results on TIMEBench are closer to random performance, underscoring the stricter evaluation criteria of our benchmark in assessing temporal dimensions.

Ablation Study. We conducted ablation experiments on VideoLLaMA 2 (Cheng et al. 2024) to dissect the contributions of our approach. Table 2 shows that both TIME and

¹For Video-LLaVA, VideoLLaMA 2 and ShareGPT4Video, we combined their respective original instruction datasets with our TIME dataset. In the case of InternVL 2.5, since the the original instruction dataset is unavailable, we directly fine-tuned the model using our TIME dataset on the pre-trained checkpoint.

Model	TIMEBench						MVB	TC	AE-V	V-MME	AN-QA
	LO	DU	DY	OR	RE	AVG					
Baseline	32.3	30.3	22.8	28.0	24.0	27.4	45.9	44.2	11.9	41.7	44.3
+ TIME	34.3	41.0	26.5	34.8	26.3	32.6	<u>47.1</u>	<u>57.4</u>	<u>14.0</u>	41.6	<u>44.6</u>
+ MTP	35.3	32.0	24.5	31.0	24.8	29.5	45.9	49.7	13.8	41.9	44.5
+ ALL	33.3	43.0	28.0	<u>33.0</u>	26.8	32.8	47.2	58.2	14.7	<u>41.8</u>	44.7

Table 2: Ablation study of TIME dataset and MTP method. The best performance is highlighted in **bold**, and the second-best result is indicated with underlines.

Method	LO	DU	DY	OR	RE	AVG
Baseline	32.3	30.3	22.8	28.0	24.0	27.4
Mixing	34.3	41.0	26.5	34.8	26.3	32.6
After	31.3	30.3	23.0	25.3	<u>25.3</u>	26.6
Before	<u>34.2</u>	<u>35.3</u>	<u>23.3</u>	<u>29.8</u>	24.0	<u>28.7</u>

Table 3: Ablation study on the strategy for incorporating the TIME dataset. The best performance is highlighted in **bold**, and the second-best result is indicated with underlines.

Method	LO	DU	DY	OR	RE	AVG
100%	32.7	30.3	25.5	<u>29.5</u>	23.3	27.9
75%	34.0	<u>31.0</u>	<u>26.3</u>	28.0	22.5	27.9
50%	<u>35.3</u>	30.3	22.5	30.5	27.5	<u>28.8</u>
25%	38.0	31.3	26.5	26.5	<u>25.0</u>	28.9
0%	32.3	30.3	22.8	28.0	24.0	27.4

Table 4: Ablation study on frame index prediction task ratio.

MTP independently improve the model’s temporal understanding, while their combination achieves the best balance between temporal tasks and general tasks.

In Table 3, we compared different strategies for integrating TIME data into the training process. Our findings indicate that mixing TIME data with the original instruction fine-tuning dataset concurrently yields superior performance compared to sequentially adding TIME data after or before the original fine-tuning. This improvement is likely due to mitigating the distribution gap between temporal tasks and general tasks by integrating them together.

Table 4 and Table 5 explore the impact of different proportions of auxiliary temporal tasks. We observe that introducing any proportion of auxiliary tasks improves performance on temporal tasks, confirming their effectiveness. However, increasing the proportion of auxiliary tasks beyond a certain threshold does not yield further improvements and may even degrade performance. Based on our experiments, the optimal mix is 50% assigned video question-answering tasks and 25% frame index prediction tasks.

Qualitative Analysis. Figure 6 compares the predictions of the original VideoLLaMA 2 (Cheng et al. 2024) and the fine-tuned version on two TIMEBench samples. In the *Location* example, the fine-tuned model better captures subtle temporal location differences, while in the *Dynamic* sample (with

Method	LO	DU	DY	OR	RE	AVG
100%	<u>35.0</u>	<u>35.3</u>	25.3	<u>32.0</u>	<u>25.4</u>	<u>30.1</u>
75%	36.7	31.0	26.3	30.8	24.3	29.3
50%	34.3	41.0	27.5	32.8	25.5	31.6
25%	33.7	30.7	<u>26.8</u>	29.5	23.8	28.5
0%	32.3	30.3	22.8	28.0	24.0	27.4

Table 5: Ablation study on the assigned videoQA task ratio.

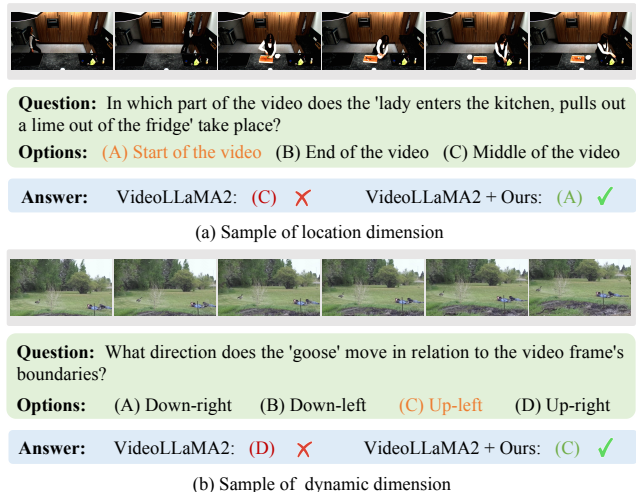


Figure 6: Visualization of partial sample prediction results in TIMEBench. The correct answer is highlighted in orange.

reversed video), the fine-tuned model relies less on priors like face orientation compared to the original model.

Conclusion

In this work, we present a comprehensive framework to enhance temporal understanding in video-LLMs by delineating five key dimensions of temporal reasoning. Building on these dimensions, we introduced the TIME instruction fine-tuning dataset and a novel multi-task prompt-tuning approach that overcomes the limitations of traditional data labeling. Furthermore, our proposed TIMEBench benchmark exposes the temporal shortcomings of existing video-LLMs while validating the effectiveness of our approach in advancing temporal reasoning capabilities.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China, No.:62376140, and No.:U23A20315; the Science and Technology Innovation Program for Distinguished Young Scholars of Shandong Province Higher Education Institutions, No.:2023KJ128; the Special Fund for Taishan Scholar Project of Shandong Province; and Kuaishou Technology, No.: 1480024026.

References

- Chen, L.; Wei, X.; Li, J.; Dong, X.; Zhang, P.; Zang, Y.; Chen, Z.; Duan, H.; Bin, L.; Tang, Z.; Yuan, L.; Qiao, Y.; Lin, D.; Zhao, F.; and Wang, J. 2024a. ShareGPT4Video: Improving Video Understanding and Generation with Better Captions. In *Proceedings of the Advances in Neural Information Processing Systems*.
- Chen, X.; Lin, Y.; Zhang, Y.; and Huang, W. 2023. AutoEval-Video: An Automatic Benchmark for Assessing Large Vision Language Models in Open-Ended Video Question Answering. In *Proceedings of the European Conference on Computer Vision*, 1–14.
- Chen, Y.; Ge, Y.; Ge, Y.; Ding, M.; Li, B.; Wang, R.; Xu, R.; Shan, Y.; and Liu, X. 2024b. EgoPlan-Bench: Benchmarking Multimodal Large Language Models for Human-Level Planning. *arXiv*, abs/2312.06722.
- Chen, Z.; Wu, J.; Wang, W.; Su, W.; Chen, G.; Xing, S.; Zhong, M.; Zhang, Q.; Zhu, X.; Lu, L.; Li, B.; Luo, P.; Lu, T.; Qiao, Y.; and Dai, J. 2024c. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.
- Cheng, C.; Guan, J.; Wu, W.; and Yan, R. 2025a. Scaling Video-Language Models to 10K Frames via Hierarchical Differential Distillation. *arXiv*, abs/2506.15220.
- Cheng, Z.; Hu, J.; Liu, Z.; Si, C.; Li, W.; and Gong, S. 2025b. V-STaR: Benchmarking Video-LLMs on Video Spatio-Temporal Reasoning. *arXiv*, abs/2503.11495.
- Cheng, Z.; Leng, S.; Zhang, H.; Xin, Y.; Li, X.; Chen, G.; Zhu, Y.; Zhang, W.; Luo, Z.; Zhao, D.; and Bing, L. 2024. VideoLLaMA 2: Advancing Spatial-Temporal Modeling and Audio Understanding in Video-LLMs. *arXiv*, abs/2406.07476.
- Damen, D.; Doughty, H.; Farinella, G. M.; Furnari, A.; Kazakos, E.; Ma, J.; Moltisanti, D.; Munro, J.; Perrett, T.; Price, W.; and Wray, M. 2022. Rescaling Egocentric Vision: Collection, Pipeline and Challenges for EPIC-KITCHENS-100. *International Journal of Computer Vision*, 130(1): 33–55.
- Fu, C.; Dai, Y.; Luo, Y.; Li, L.; Ren, S.; Zhang, R.; Wang, Z.; Zhou, C.; Shen, Y.; Zhang, M.; Chen, P.; Li, Y.; Lin, S.; Zhao, S.; Li, K.; Xu, T.; Zheng, X.; Chen, E.; Shan, C.; He, R.; and Sun, X. 2025. Video-MME: The First-Ever Comprehensive Evaluation Benchmark of Multi-modal LLMs in Video Analysis. In *Proceeding of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24108–24118.
- Hu, X.; Dai, J.; Li, M.; Peng, C.; Li, Y.; and Du, S. 2022. Online human action detection and anticipation in videos: A survey. *Neurocomputing*, 491: 395–413.
- Huang, L.; Zhao, X.; and Huang, K. 2021. GOT-10k: A Large High-Diversity Benchmark for Generic Object Tracking in the Wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(5): 1562–1577.
- Huang, Q.; Dong, X.; Zhang, P.; Wang, B.; He, C.; Wang, J.; Lin, D.; Zhang, W.; and Yu, N. 2024. OPERA: Alleviating Hallucination in Multi-Modal Large Language Models via Over-Trust Penalty and Retrospection-Allocation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13418–13427.
- Keye, K. 2025. Kwai Keye-VL Technical Report. *arXiv*, abs/2507.01949.
- Krishna, R.; Hata, K.; Ren, F.; Fei-Fei, L.; and Niebles, J. C. 2017. Dense-Captioning Events in Videos. In *Proceedings of the IEEE International Conference on Computer Vision*, 706–715.
- Kuehne, H.; Arslan, A. B.; and Serre, T. 2014. The Language of Actions: Recovering the Syntax and Semantics of Goal-Directed Human Activities. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 780–787.
- Lei, J.; Berg, T. L.; and Bansal, M. 2021. Detecting Moments and Highlights in Videos via Natural Language Queries. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, volume 34, 11846–11858.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; and Li, C. 2025a. LLaVA-OneVision: Easy Visual Task Transfer. *Transactions on Machine Learning Research*.
- Li, K.; Wang, Y.; He, Y.; Li, Y.; Wang, Y.; Liu, Y.; Wang, Z.; Xu, J.; Chen, G.; Lou, P.; Wang, L.; and Qiao, Y. 2024. MVBench: A Comprehensive Multi-modal Video Understanding Benchmark. In *Proceeding of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22195–22206.
- Li, L.; Liu, Y.; Yao, L.; Zhang, P.; An, C.; Wang, L.; Sun, X.; Kong, L.; and Liu, Q. 2025b. Temporal Reasoning Transfer from Text to Video. In *Proceedings of the International Conference on Learning Representations*.
- Lin, B.; Ye, Y.; Zhu, B.; Cui, J.; Ning, M.; Jin, P.; and Yuan, L. 2024. Video-LLaVA: Learning United Visual Representation by Alignment Before Projection. In *Proceeding of the Conference on Empirical Methods in Natural Language Processing*.
- Liu, M.; Wang, X.; Nie, L.; He, X.; Chen, B.; and Chua, T. 2018a. Attentive Moment Retrieval in Videos. In *Proceedings of the International ACM SIGIR Conference on Research & Development in Information Retrieval*, 15–24.
- Liu, M.; Wang, X.; Nie, L.; Tian, Q.; Chen, B.; and Chua, T. 2018b. Cross-modal Moment Localization in Videos. In *Proceedings of the ACM International on Multimedia*, 843–851.

- Liu, R.; Li, C.; Tang, H.; Ge, Y.; Shan, Y.; and Li, G. 2024a. ST-LLM: Large Language Models Are Effective Temporal Learners. In *Proceeding of the European Conference on Computer Vision*, 1–18.
- Liu, Y.; Li, S.; Liu, Y.; Wang, Y.; Ren, S.; Li, L.; Chen, S.; Sun, X.; and Hou, L. 2024b. TempCompass: Do Video LLMs Really Understand Videos? In *Findings of the Association for Computational Linguistics*, 8731–8772.
- Lv, S.-L.; Chen, Y.-Y.; Zhou, Z.; Yang, M.; and Guo, L.-Z. 2025a. BMIP: Bi-directional Modality Interaction Prompt Learning for VLM. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 5887–5895.
- Lv, S.-L.; Zhu, R.; Li, Y.-F.; and Guo, L.-Z. 2025b. Unlabeled Data or Pre-trained Model: Rethinking Semi-Supervised Learning and Pretrain-Finetuning. *arXiv*, abs/2505.13317.
- OpenAI. 2024. GPT-4 Technical Report. *arXiv*, abs/2303.08774.
- Otani, M.; Nakashima, Y.; Rahtu, E.; and Heikkilä, J. 2020. Uncovering Hidden Challenges in Query-Based Video Moment Retrieval. In *Proceedings of the British Machine Vision Conference*.
- Regneri, M.; Rohrbach, M.; Wetzel, D.; Thater, S.; Schiele, B.; and Pinkal, M. 2013. Grounding Action Descriptions in Videos. *Transactions of the Association for Computational Linguistics*, 1: 25–36.
- Shang, X.; Di, D.; Xiao, J.; Cao, Y.; Yang, X.; and Chua, T. 2019. Annotating Objects and Relations in User-Generated Videos. In *Proceedings of the International Conference on Multimedia Retrieval*, 279–287.
- Shangguan, Z.; Li, C.; Ding, Y.; Zheng, Y.; Zhao, Y.; Fitzgerald, T.; and Cohan, A. 2025. TOMATO: Assessing Visual Temporal Reasoning Capabilities in Multimodal Foundation Models. In *Proceedings of the International Conference on Learning Representations*.
- Sigurdsson, G. A.; Varol, G.; Wang, X.; Farhadi, A.; Laptev, I.; and Gupta, A. 2016. Hollywood in Homes: Crowdsourcing Data Collection for Activity Understanding. In *Proceedings of the European Conference on Computer Vision*, volume 9905, 510–526.
- Song, E.; Chai, W.; Xu, W.; Xie, J.; Liu, Y.; and Wang, G. 2025. Video-MMLU: A Massive Multi-Discipline Lecture Understanding Benchmark. *arXiv*, abs/2504.14693.
- Song, Y.; Byrne, E.; Nagarajan, T.; Wang, H.; Martin, M.; and Torresani, L. 2023. Ego4D Goal-Step: Toward Hierarchical Understanding of Procedural Activities. In *Proceedings of the Annual Conference on Neural Information Processing Systems*.
- Tang, C.; Li, Y.; Yang, Y.; Zhuang, J.; Sun, G.; Li, W.; Ma, Z.; and Zhang, C. 2025. video-SALMONN 2: Captioning-Enhanced Audio-Visual Large Language Models. *arXiv*, abs/2506.15220.
- Tang, Y.; Shimada, D.; Bi, J.; and Xu, C. 2024. AVicuna: Audio-Visual LLM with Interleaver and Context-Boundary Alignment for Temporal Referential Dialogue. *arXiv*, abs/2403.16276.
- Team, Q. 2024. Qwen2.5: A Party of Foundation Models. Accessed on 2025-6-15. Available at: <https://qwenlm.github.io/blog/qwen2.5/>.
- Tong, S.; Liu, Z.; Zhai, Y.; Ma, Y.; LeCun, Y.; and Xie, S. 2024. Eyes Wide Shut? Exploring the Visual Shortcomings of Multimodal LLMs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9568–9578.
- 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. 2024a. Qwen2-VL: Enhancing Vision-Language Model’s Perception of the World at Any Resolution. *arXiv*, abs/2409.12191.
- Wang, Y.; Wang, Y.; Zhao, D.; Xie, C.; and Zheng, Z. 2024b. VideoHalluciner: Evaluating Intrinsic and Extrinsic Hallucinations in Large Video-Language Models. *arXiv*, abs/2406.16338.
- Xiao, J.; Huang, N.; Qin, H.; Li, D.; Li, Y.; Zhu, F.; Tao, Z.; Yu, J.; Lin, L.; Chua, T.; and Yao, A. 2025. VideoQA in the Era of LLMs: An Empirical Study. *International Journal of Computer Vision*, 133(7): 3970–3993.
- Yao, Y.; Yu, T.; Zhang, A.; Wang, C.; Cui, J.; Zhu, H.; Cai, T.; Li, H.; Zhao, W.; He, Z.; Chen, Q.; Zhou, H.; Zou, Z.; Zhang, H.; Hu, S.; Zheng, Z.; Zhou, J.; Cai, J.; Han, X.; Zeng, G.; Li, D.; Liu, Z.; and Sun, M. 2024. MiniCPM-V: A GPT-4V Level MLLM on Your Phone. *arXiv*, abs/2408.01800.
- Yu, Q.; Li, J.; Wei, L.; Pang, L.; Ye, W.; Qin, B.; Tang, S.; Tian, Q.; and Zhuang, Y. 2024. HalluciDoctor: Mitigating Hallucinatory Toxicity in Visual Instruction Data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12944–12953.
- Yu, Z.; Xu, D.; Yu, J.; Yu, T.; Zhao, Z.; Zhuang, Y.; and Tao, D. 2019. ActivityNet-QA: A Dataset for Understanding Complex Web Videos via Question Answering. In *Proceedings of the Innovative Applications of Artificial Intelligence Conference*, 9127–9134.
- Zhang, H.; Liu, M.; Li, Z.; Wen, H.; Guan, W.; Wang, Y.; and Nie, L. 2025. Spatial Understanding from Videos: Structured Prompts Meet Simulation Data. In *Proceedings of the Advances in Neural Information Processing Systems*, 1–16.
- Zhu, R.; Lv, S.-L.; Wang, Z.-K.; and Guo, L.-Z. 2025. Bi-CoG: Bi-Consistency-Guided Self-Training for Vision-Language Models. *arXiv*, abs/2510.20477.