

SlowFast-LLaVA-1.5: A Family of Token-Efficient Video Large Language Models for Long-Form Video Understanding

Mingze Xu[°], Mingfei Gao[°], Shiyu Li[°], Jiasen Lu, Zhe Gan, Zhengfeng Lai,
Meng Cao, Kai Kang, Yinfei Yang[†], Afshin Dehghan[†]
Apple

{mingze_xu2, mgao22, shiyu_li, yinfeiy, adehghan}@apple.com

[°]First authors; [†]Senior authors

Abstract

We introduce **SlowFast-LLaVA-1.5** (abbreviated as SF-LLaVA-1.5), a family of video large language models (LLMs) offering a token-efficient solution for long-form video understanding. We incorporate the two-stream SlowFast mechanism into a streamlined training pipeline, and perform joint video-image training on a carefully curated data mixture of only publicly available datasets. Our primary focus is on highly efficient model scales (1B and 3B), demonstrating that even relatively small Video LLMs can achieve state-of-the-art performance on video understanding, meeting the demand for mobile-friendly models. Experimental results demonstrate that SF-LLaVA-1.5 achieves superior performance on a wide range of video and image tasks, with robust results at all model sizes (ranging from 1B to 7B). Notably, SF-LLaVA-1.5 achieves state-of-the-art results in long-form video understanding (e.g., LongVideoBench and MLVU) and excels at small scales across various video benchmarks.

1 Introduction

Video large language models (LLMs) (Maaz et al., 2024b; Lin et al., 2023a; Xu et al., 2024a) integrate video perception into pre-trained LLMs to process videos and generate responses to user commands. Although significant progress has been made, notable limitations remain in existing Video LLMs. *First*, they enhance perception and reasoning over long video sequences by leveraging the LLM’s increasing context length and handling massive input frames (Shen et al., 2024; Chen et al., 2024d; Zhang et al., 2024c). However, the potential for transferring this capability to highly efficient models is underexplored. *Second*, achieving optimal performance typically requires internal datasets and a complex training lifecycle, with selective parameters frozen at each stage (Li et al., 2024a; Zhang et al., 2025a). These intricate designs lead to high computational costs and reproducibility challenges. *Third*, many Video LLMs (Zohar et al., 2024; Li et al., 2024d) are optimized exclusively for video tasks, limiting their effectiveness as joint models for image understanding tasks.

Building upon the success of SlowFast-LLaVA (Xu et al., 2024b), we introduce **SlowFast-LLaVA-1.5**, a new family of Video LLMs for long-form video understanding, focusing on the most efficient model scales (1B and 3B). Our model family is both effective and token-efficient in modeling long-range temporal context. This is achieved by employing the SlowFast mechanism, which balances the trade-off between processing more input frames that significantly increases the token count and computational cost, and reducing tokens per frame that inevitably loses fine-grained details. Specifically, the Slow pathway captures detailed spatial features at a low frame rate, while the Fast pathway operates at a high frame rate with fewer tokens per frame to focus on motion cues. The success of our model also relies on a streamlined training pipeline and a carefully curated data mixture. Our model training consists of only two stages. The first stage is supervised fine-tuning on image-only data, providing a good foundation for general knowledge and reasoning. The second stage conducts video-image joint training to learn spatial and temporal features for video understanding while maintaining strong performance in image understanding. To

ensure seamless reproducibility, all pre-trained weights and training datasets used in this work are publicly accessible.

We comprehensively evaluate our models on various video and image benchmarks. Experimental results demonstrate that SlowFast-LLaVA-1.5 achieves state-of-the-art performance in long-form video understanding. Notably, our 7B model scores 62.5% on LongVideoBench and 71.5% on MLVU, outperforming existing methods by a clear margin. SlowFast-LLaVA-1.5 also excels at smaller model sizes, achieving 56.6% and 60.8% on Video-MME (w/o sub) at the 1B and 3B scales, respectively. As a unified image and video model, it maintains strong image performance despite the simple training recipe.

Our main contributions are as follows. *First*, we introduce SlowFast-LLaVA-1.5, a new family of Video LLMs ranging from 1B to 7B parameters. We demonstrate the effectiveness of incorporating the SlowFast mechanism into a supervised fine-tuning framework, modeling long-range context while maintaining high efficiency. *Second*, our model family provides enhanced reproducibility by using only two training stages and publicly available datasets, distinguishing it from existing methods. *Third*, SlowFast-LLaVA-1.5 achieves the state-of-the-art performance on long-form video understanding. Moreover, our smaller models (1B and 3B) clearly outperform comparable Video LLMs across video benchmarks.

2 Related Work

Image Large Language Models have gained widespread attention (Achiam et al., 2023; Team et al., 2023; Touvron et al., 2023; Chen et al., 2024e; Bai et al., 2025). Significant progress across multiple fronts includes: (i) enhancing data quantity and quality during pre-training (McKinzie et al., 2024; Liu et al., 2024a; Lin et al., 2023b; Li et al., 2024c) and supervised fine-tuning (SFT) (Zhang et al., 2025b; Deitke et al., 2024; Chen et al., 2024a; Wang et al., 2023; Tong et al., 2025); (ii) accommodating images of various high resolutions (Lin et al., 2023c; Zhang et al., 2024b; Wang et al., 2024b); (iii) improving architecture designs, including different visual encoders (Zhai et al., 2023; Tong et al., 2024; Shi et al., 2024) and vision-language connectors (Li et al., 2023a; Cha et al., 2024); and (iv) conducting comprehensive studies for easy deployment (Team et al., 2023; Marafioti et al., 2025). These rapid advancements also establish a strong foundation for related areas such as video understanding (Maaz et al., 2024b; Lin et al., 2023a), referring & grounding (You et al., 2023; 2024), and visual agents (Durante et al., 2024; Yang et al., 2025).

Video Large Language Models have become an active research area (Li et al., 2023b; Song et al., 2024; Chen et al., 2024b; Zhang et al., 2024e; Zohar et al., 2024). Early Video LLMs are developed as specialist models (Zhang et al., 2023; Cheng et al., 2024; Xu et al., 2024a; Ryoo et al., 2024), achieving strong performance on video tasks but with some trade-offs in image understanding. Training-free Video LLMs (Kim et al., 2024; Xu et al., 2024b) offer an efficient alternative by leveraging Image LLMs without fine-tuning on video data, enabling flexible deployment across various applications. Recent models (Zhang et al., 2024g; Liu et al., 2025; Zhang et al., 2025a) are jointly trained on video and image datasets, obtaining superior results in both modalities. Long-form video understanding (Zhou et al., 2024; Wu et al., 2025) gained increasing attentions, addressing hour-long videos (Chen et al., 2024d; Li et al., 2024d) or live streams (Qian et al., 2024; Zhang et al., 2024a) while optimizing the token efficiency (Lee et al., 2024b). The proposed SlowFast-LLaVA-1.5 is a family of Video LLMs designed for modeling long-range temporal context. It enhances SlowFast-LLaVA (Xu et al., 2024b) by implementing the SlowFast design within a unified video-image training framework, achieving state-of-the-art performance with efficient token utilization.

3 SlowFast-LLaVA-1.5

We provide a detailed explanation of SlowFast-LLaVA-1.5 (abbreviated as SF-LLaVA-1.5), which incorporates the SlowFast video projector into a LLaVA-style architecture, improving long-range temporal modeling while optimizing token efficiency. In contrast to its training-free pioneer (Xu et al., 2024b), this paper (i) systematically investigates different instantiations based on the generic SlowFast idea (Sec. 3.1), (ii) designs a compact yet effective training pipeline (Sec. 3.2), and (iii) introduces tailored data mixtures using only publicly available datasets for each training stage (Sec. 3.3).

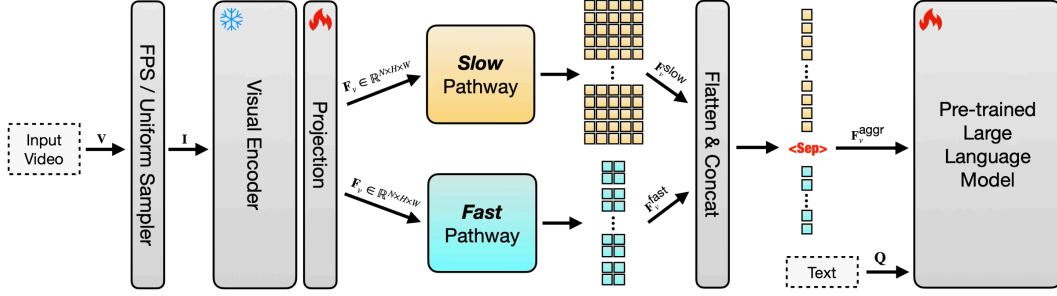


Figure 1: **Visualization of the video understanding pipeline in SlowFast-LLaVA-1.5.** Compared to its training-free pioneer (Xu et al., 2024b), our projector and LLM are fine-tuned throughout the training cycle, while keeping the vision encoder frozen.

3.1 Model Architecture

As shown in Fig. 1, the architecture of SF-LLaVA-1.5 follows the core design principle of SF-LLaVA (Xu et al., 2024b). It takes a video/image V and a question Q as inputs and responds with a textual answer A . For video inputs, we sample N frames, $I = \{I_1, I_2, \dots, I_N\}$, at a fixed frame rate without special frame assembling (N equals 1 for image input). After that, a visual encoder (e.g., OryxViT (Liu et al., 2025)) is used to extract frame-level features $F_v \in \mathbb{R}^{N \times H \times W}$ from the inputs independently, keeping their original aspect ratio. The video and image feature tokens are then fed into different projectors, with video using the two-stream SlowFast projector and image using a two-layer MLP.

The SlowFast projector processes F_v through two pathways, one dedicated to capturing spatial patterns and the other to modeling motion cues.

- **The Slow pathway**, which focuses on capturing detailed spatial semantics, operates at a reduced frame rate by downsampling the total frame count from N to N^{slow} . To further improve the efficiency while preserving sufficient details, it applies spatial pooling over F_v with proper strides of $\sigma_h \times \sigma_w$. The output feature is $F_v^{slow} \in \mathbb{R}^{N^{slow} \times H^{slow} \times W^{slow}}$, where $H^{slow} = H/\sigma_h$ and $W^{slow} = W/\sigma_w$.
- **The Fast Pathway**, which focuses on modeling long-range context, maintains the original frame rate, while downsampling more aggressively on the spatial resolution to $H^{fast} \times W^{fast}$. The output feature is $F_v^{fast} \in \mathbb{R}^{N^{fast} \times H^{fast} \times W^{fast}}$, where $N^{fast} = N$, $H^{fast} \ll H^{slow}$ and $W^{fast} \ll W^{slow}$.

F_v^{slow} and F_v^{fast} are flattened and concatenated together as a token vector F_v^{aggr} , which serves as the final visual input to the LLM. A dedicated special token is typically used to separate F_v^{slow} and F_v^{fast} , assisting the LLM in distinguishing the two sets of features.

3.1.1 Instantiations of SlowFast

Next, we describe two approaches for organizing the Slow and Fast tokens.

- **The Group-based SlowFast (GSF)** places the Slow tokens before the Fast tokens (Appendix Fig. 2 above). This design is inspired by the AnyRes (Zhang et al., 2024f) technique in image understanding, where the Fast tokens provide a global overview of the video and the Slow tokens capture fine-grained spatial details. Notably, SlowFast-LLaVA (Xu et al., 2024b) works effectively only under this setting, as it is a training-free model that benefits from “overfitting” to its image backbone (i.e., LLaVA-NeXT).
- **The Interleaved SlowFast (ISF)** arranges the tokens according to their spatial and temporal order (Appendix Fig. 2 below). Since Slow and Fast frames contain different numbers of tokens, a learnable special token is utilized to separate adjacent frames, allowing the LLM to distinguish which frame a token belongs to. Different from GSF, N^{fast} equals to $N - N^{slow}$ in this approach. ISF balances the presence of both token

types throughout the input sequence, preventing the model from becoming overly focused on one type of information at a time.

Unless noted otherwise, we use GSF by default, as it aligns better with the image pipeline using AnyRes inputs. Interestingly, experiments (Sec. 4.4.1) show that SF-LLaVA-1.5 is not sensitive to this setting, suggesting that the generic SlowFast idea and our training recipe are the main reason for the strong performance on long-form video understanding.

3.2 Training Pipeline

The training pipeline of SF-LLaVA-1.5 is much simpler than most of the existing Video LLMs (Chen et al., 2024e; Zohar et al., 2024; Zhang et al., 2025a; Li et al., 2024d) with only two training stages, as detailed in Table 1.

Settings	Stage I	Stage II
Dataset	Image	Image & Video
Trainable	Projector & LLM	Projector & LLM
Image Projector	MLP w/ GELU	MLP w/ GELU
Video Projector	-	SlowFast
Batch Size	512	512
Learning Rate	$2e^{-5}$	$2e^{-5}$
Context Length	8K	16K
Number of Input Frames	1	1 or 128
Max Image Resolution	1280×1280	1536×1536
Max Video Resolution	-	480×480
Training Steps	1 epoch	1 epoch

Table 1: Training settings for SlowFast-LLaVA-1.5.

Stage I (image understanding) conducts SFT with images to provide a good warmup status for video understanding. For simplicity and efficiency, we do not use any extra pretraining stages (Li et al., 2024a) or image splitting strategy (Lin et al., 2023c), although they have proven to be effective for boosting text-rich results. Instead, we use native resolution inputs following Oryx (Liu et al., 2025), where, for each image $I_i \in \mathbb{R}^{H_i \times W_i}$, we have a low resolution I_i^{low} and a high resolution I_i^{high} input. The low-resolution image is obtained by simply resizing the original image to a base resolution, as in $I_i^{low} = \text{resize}(I_i, H_i^{base} \times W_i^{base})$. For I_i^{high} , we keep its original aspect ratio and resize it to $H_i^{high} \times W_i^{high}$, as in Eq. 1 and 2,

$$\text{scale} = \begin{cases} \sqrt{\theta^{I^{max}} / (H_i \times W_i)}, & \text{if } H_i \times W_i > \theta^{I^{max}} \\ \sqrt{\theta^{I^{min}} / (H_i \times W_i)}, & \text{if } H_i \times W_i < \theta^{I^{min}} \\ 1.0, & \text{otherwise,} \end{cases} \quad (1)$$

$$\begin{aligned} H_i^{high} &= \text{int}(H_i * \text{scale} / p) * p \\ W_i^{high} &= \text{int}(W_i * \text{scale} / p) * p, \end{aligned} \quad (2)$$

where H_i^{high} and W_i^{high} represent the resized heights and widths and p denotes the patch size of the ViT-based vision encoder. Eq. 1 calculates the resizing scale ensuring that the area of I_i^{high} is between two pre-defined, minimum area $\theta^{I^{min}}$ and maximum area $\theta^{I^{max}}$ thresholds.

Eq. 2 makes sure that both H_i^{high} and W_i^{high} are multiples of p . To accommodate different input resolutions, the original position embeddings of the vision encoder are rescaled using bilinear interpolation. After feature projection, the low-resolution and high-resolution image features are concatenated together as the final image feature.

Stage II (joint image & video understanding) performs SFT jointly with images and videos, initialized by the pre-trained checkpoint from Stage I. By default, we keep the image resizing setting the same as Stage I, except that we increase the maximum area threshold, $\theta^{I^{max}}$, to a larger value for better performance. For video, each frame uses a single resolution that is set using the same strategy as in Eq. 1 and 2, where we use $\theta^{V^{min}}$ and $\theta^{V^{max}}$ to denote the corresponding minimum and maximum area thresholds.

3.3 Data Mixture

Our image and video mixtures are detailed in Table 8. Many state-of-the-art models (Li et al., 2024a; Zhang et al., 2025a) achieve superior performance using internal training data that is unavailable to the research community. **To ensure the reproducibility of our models, we only include publicly available datasets in our data mixtures.**

Image Mixture. General, TextRich, and Knowledge are fundamental for developing the reasoning capabilities of a multimodal LLM, which can ultimately benefit both image and video understanding. We begin with datasets from these three categories in MM1.5 (Zhang et al., 2025b) and evaluate additional datasets for each group from LLaVA-OneVision (Li et al., 2024a) and InternVL2.5 (Chen et al., 2024e). Datasets are included in our mixture only if they empirically improve performance. The final mixture contains 4.67M samples.

Video Mixture. We build a diverse set of video instruction-following datasets. We begin with LLaVA-Hound (Zhang et al., 2024e), ShareGPT4Video (Chen et al., 2024b), VideoChatGPT-Plus (Maaz et al., 2024a), and ActivityNet-QA (Yu et al., 2019) to include large-scale video data with caption and QA labels. We add NExT-QA (Xiao et al., 2021) and Perception Test (Pătrăucean et al., 2023) to improve performance on temporal reasoning. Furthermore, we incorporate LLaVA-Video-178K (Zhang et al., 2024g) and Cinepile (Rawal et al., 2024) to enhance long-form video understanding. Finally, we filter out duplicate videos from the same data source and construct our final mixture with 2.01M training samples.

4 Experiments

We evaluate SF-LLaVA-1.5 across multiple video and image QA benchmarks (details will be provided in Appendix A.2). For video, we focus on long-form video understanding, while also reporting the results in general video QA and temporal reasoning. For image, we evaluate the models from general, knowledge, and text-rich perspectives.

4.1 Implementation Details

Model Architecture. We use Oryx-ViT¹ (Liu et al., 2025) with patch size 16 as visual encoder and Qwen2.5² (Bai et al., 2025) series of LLMs at varying scales as the backbone. We employ different projectors for video and image inputs. Specifically, the Group-based SlowFast (GSF) structure is used to aggregate video tokens. For the Slow pathway, we uniformly select $N^{slow} = 32$ frames and apply 2×2 pooling to their extracted features. For the Fast pathway, we use features of all frames (*i.e.*, $N^{fast} = N = 128$) and downsample their features to 4×4 tokens. For the image projector, we use a two-layer MLP with GELU activation function.

Training Details. As summarized in Table 1, we freeze the visual encoder in all stages and only fine-tune the projectors and LLM. We use the same hyperparameters for 1B, 3B, and 7B models, setting the total batch size to 512 and learning rate to $2e^{-5}$. All models are trained on 128 H100-80G GPUs for 1 epoch.

- **Training Stage I** only uses image understanding data. The low resolution image I_i^{low} is fixed at $H_i^{base} \times W_i^{base} = 384 \times 384$ and the high resolution image I_i^{high} is obtained as in Eq. 1 and Eq. 2, where $\theta^{I^{min}} = 0$ and $\theta^{I^{max}} = 1280^2$. The maximum context length is set to 8K. The models trained by this stage are named as **SF-LLaVA-1.5-Image**.
- **Training Stage II** continues training based on SF-LLaVA-1.5-Image by combining our video and image data mixture. For image, the high-resolution image is obtained in the same way as Stage I, except that we increase $\theta^{I^{max}}$ to 1536^2 . For video, we follow prior work (Zohar et al., 2024) and sample frames at 1 FPS. We set the max frame number to 128 and uniformly sample the frames if the number exceeds this upper bound. For each video frame, we set $\theta^{V^{min}} = 288^2$ and $\theta^{V^{max}} = 480^2$. The maximum context length is set to 16K. The models trained by this stage are named as **SF-LLaVA-1.5**.

¹<https://huggingface.co/THUdyh/Oryx-ViT>.

²<https://huggingface.co/Qwen>.

Model	Max Input Frames	Max Input Tokens	General VideoQA		Long-Form Video Understanding			Temporal Reasoning	
			VideoMME (w/o sub)	PercepTest (val)	LongVideoBench (val)	MLVU (m-avg)	LVBench (avg)	TempComp (mc)	NExT-QA (test)
1B Model Comparison									
LLaVA-OV-0.5B (Li et al., 2024a)	32	6K	44.0	49.2	45.8	50.3	32.7 [†]	53.2	57.2
MM1.5-1B (Zhang et al., 2025b)	24	3K	45.7	-	43.9	-	-	-	71.8
LinVT-Mipha-1.6B (Gao et al., 2024)	120	-	44.5	-	49.7	56.2	-	45.2	71.1
Apollo-1.5B (Zohar et al., 2024)	2fps	3K	53.0	61.0	54.1	63.3	-	60.8	-
InternVL2.5-2B (Chen et al., 2024e)	64	16K	51.9	-	52.0	61.4	37.9 [†]	53.4 [†]	77.2 [†]
Qwen2-VL-2B (Wang et al., 2024b)	2fps	16K	55.6	53.9	48.7 [†]	62.7 [†]	39.4 [†]	60.6 [†]	77.2 [†]
SF-LLaVA-1.5-1B	128	9K	56.6	61.9	54.3	64.3	39.7	60.5	76.7
3B Model Comparison									
VILA1.5-3B (Liu et al., 2024e)	8	2K	42.2	49.1	42.9	44.4	-	56.1	-
MM1.5-3B (Zhang et al., 2025b)	24	3K	49.5	-	45.4	-	-	-	74.7
LongVU-3.2B (Shen et al., 2024)	1fps	8K	51.5	-	-	55.9	-	-	-
InternVL2-4B (Chen et al., 2024f)	64	16K	53.9	53.9 [†]	53.0	59.9	35.1 [†]	60.2 [†]	71.1 [†]
LinVT-Blip3-4B (Zohar et al., 2024)	120	-	58.3	-	56.6	67.9	-	59.6	80.1
Apollo-3B (Zohar et al., 2024)	2fps	3K	58.4	65.0	55.1	68.7	-	62.5	-
SF-LLaVA-1.5-3B	128	9K	60.8	65.8	57.3	68.8	43.3	64.0	80.8
7B Model Comparison									
MM1.5-7B (Zhang et al., 2025b)	24	3K	53.5	-	49.4	-	-	-	76.9
Kangaroo-8B (Liu et al., 2024b)	64	10K	56.0	-	54.8	61.0	39.4	62.5	-
Oryx1.5-7B (Liu et al., 2025)	64	14K	58.8	70.0	56.3	67.5	39.0 [†]	58.8 [†]	81.8
LLaVA-OV-7B (Li et al., 2024a)	32	6K	58.2	49.7	56.5	64.7	-	-	79.4
LLaVA-Video-7B (Zhang et al., 2024g)	64	11K	63.3	66.9	58.2	70.8	-	-	83.2
Apollo-7B (Zohar et al., 2024)	2fps	2K	61.3	67.3	58.5	70.9	-	64.9	-
NVILA-8B (Liu et al., 2024e)	256	8K	64.2	65.4 [†]	57.7	70.1	44.0 [†]	69.7 [†]	82.2
InternVL2.5-8B (Chen et al., 2024e)	64	16K	64.2	-	60.0	69.0 [†]	43.2 [†]	68.3 [†]	85.0 [†]
Qwen2-VL-7B (Wang et al., 2024b)	2fps	16K	63.3	62.3	55.6 [†]	69.8 [†]	44.7 [†]	67.9 [†]	81.2 [†]
SF-LLaVA-1.5-7B	128	9K	63.9	69.6	62.5	71.5	45.3	68.8	83.3

Table 2: **Comparison with state-of-the-art models on video understanding.** Bold and underlined are the best and second-best results for each task. [†]denotes reproduced results.

4.2 Video Understanding Results

We mainly compare SF-LLaVA-1.5 with state-of-the-art Video LLMs that are trained on publicly available datasets. Here we highlight some key observations based on Table 2.

First, SF-LLaVA-1.5 achieves state-of-the-art results in long-form video understanding. Specifically, SF-LLaVA-1.5 outperforms existing models on both LongVideoBench and LVBench across all model sizes. For reference, it surpasses InternVL2.5 at both 1B (+2.3% on LongVideoBench and +1.8% on LVBench) and 7B (+2.5% on LongVideoBench and +2.1% on LVBench) scales. SF-LLaVA-1.5 also exhibits leading performance on MLVU. Compared to the state-of-the-art model, Apollo, it achieves +1.0% at the 1B scale and comparable results at other scales. Additionally, SF-LLaVA-1.5 delivers better results even compared to Video LLMs tailored for long videos, such as LongVU. For instance, SF-LLaVA-1.5-3B significantly surpasses LongVU-3.2B by +9.3% on Video-MME and +12.9% on MLVU.

Second, SF-LLaVA-1.5 is the state-of-the-art model at the smaller scales. As edge deployment becomes increasingly important, more models are emerging in the 1B and 3B sizes, including LLaVA-OV, InternVL2.5, Qwen2-VL, and Apollo. For reference, SF-LLaVA-1.5-1B surpasses Qwen2-VL-2B across benchmarks (e.g., 56.6% vs. 55.6% on Video-MME, 61.9% vs. 53.9% on Perception Test, 64.3% vs. 62.7% on MLVU). Compared to Apollo-1.5B, SF-LLaVA-1.5-1B exhibits a +3.6% improvement on Video-MME, while leading in other tasks. Similarly, at the 3B scale, SF-LLaVA-1.5-3B outperforms Apollo-3B by +2.4% on Video-MME for general Video QA and by +1.5% on TempCompass for temporal reasoning.

Third, SF-LLaVA-1.5 optimizes the trade-off between performance and efficiency. SF-LLaVA-1.5 excels in long-form video understanding while using fewer tokens than existing methods. Using Oryx1.5 as an example, SF-LLaVA-1.5 utilizes only ~65% of its input tokens (9K vs. 14K) but processes twice as many frames (128 vs. 64), resulting in better performance on nearly all benchmarks (e.g., 63.9% vs. 58.8% on Video-MME and 71.5% vs. 67.5% on MLVU). Notably, NVILA uses a similar number of input tokens as SF-LLaVA-1.5, yet SF-LLaVA-1.5 surpasses it by +4.8% on LongVideoBench and +1.4% on MLVU. These results demonstrate the advantages of SF-LLaVA-1.5 in modeling long-range context.

Fourth, SF-LLaVA-1.5 exhibits robustness across tasks and model sizes. SF-LLaVA-1.5 consistently achieves strong performance across all benchmarks in Table 2. This demonstrates two key points: *i)* using two-stream SlowFast inputs is beneficial for modeling long-range

Model	Max Input Pixels	Train Stage #	Knowledge				General VQA		TextRich		
			AI2D (test)	SQA (test)	MMM(V (val)	MathV (testmini)	MM-Vet	RW-QA	OCRBench (test)	TextVQA (val)	DocVQA (test)
1B Model Comparison											
Gemini Nano-1 (Team et al., 2023)	-	-	37.9	-	26.3	27.3	-	-	-	62.5	72.2
LLaVA-OV-0.5B (Li et al., 2024a)	5.31M	4	57.1	67.2	31.4	34.8	29.1	55.6	-	-	70.0
MM1.5-1B (Zhang et al., 2025b)	4.52M	3	59.3	82.1	35.8	37.2	37.4	53.3	60.5	72.5	81.0
InternVL2.5-1B Chen et al. (2024e)	9.63M	2	69.3	-	40.9	43.2	48.8	57.5	78.5	72.0	84.8
MolmoE-1B (Deitke et al., 2024)	4.10M	2	86.4	-	34.9	34.0	-	60.4	-	78.8	77.7
SF-LLaVA-1.5-Image-1B	2.36M	1	70.8	87.8	39.3	51.2	41.1	57.1	69.5	70.2	85.2
SF-LLaVA-1.5-1B	2.36M	2	72.8	87.7	40.5	51.0	51.2	59.2	70.0	71.3	85.4
3B Model Comparison											
Gemini Nano-2 (Team et al., 2023)	-	-	51.0	-	32.6	30.6	-	-	-	65.9	74.3
MiniCPM-V2-3B (Yao et al., 2024)	1.81M	6	62.9	80.7	38.2	38.7	38.2	55.8	60.5	74.1	71.9
BLIP3-4B (Xue et al., 2024)	-	5	-	88.3	41.1	39.6	-	60.5	-	71.0	-
MM1.5-3B (Zhang et al., 2025b)	4.52M	3	65.7	85.8	37.1	44.4	41.0	56.9	65.7	76.5	87.7
Phi-3.5-V-4B (Abdin et al., 2024)	-	3	78.1	91.3	43.0	43.9	-	-	-	72.0	-
SF-LLaVA-1.5-Image-3B	2.36M	1	75.8	90.0	43.7	57.0	51.1	61.8	72.3	72.0	87.5
SF-LLaVA-1.5-3B	2.36M	2	77.0	90.3	44.7	58.6	47.5	63.4	73.4	73.0	88.8
7B Model Comparison											
VILA1.5-8B (Lin et al., 2023b)	-	-	76.6	-	38.6	36.7	-	52.7	-	68.5	40.6
Idetics2-8B (Laurençon et al., 2024a)	2.95M	3	-	-	43.0	51.4	-	-	-	73.0	74.0
Cambrian-1-8B (Tong et al., 2025)	-	2	73.0	80.4	42.7	49.0	-	64.2	62.4	71.7	77.8
LLaVA-OV-7B (Li et al., 2024a)	5.31M	4	81.4	96.0	48.8	63.2	57.5	66.3	-	-	87.5
MM1.5-7B (Zhang et al., 2025b)	4.52M	3	72.2	89.6	41.8	47.6	42.2	62.5	63.5	76.5	88.1
Oryx1.5-7B (Liu et al., 2025)	2.36M	3	79.7	-	47.1	-	-	-	71.3	75.7	90.1
InternVL2.5-8B (Chen et al., 2024e)	9.63M	2	84.5	-	56.0	64.4	-	70.1	-	79.1	93.0
Qwen2-VL-7B (Wang et al., 2024b)	-	3	83.0	-	54.1	58.2	62.0	70.1	-	84.3	94.5
SF-LLaVA-1.5-Image-7B	2.36M	1	79.2	91.8	47.0	61.0	50.1	64.6	74.2	75.4	89.7
SF-LLaVA-1.5-7B	2.36M	2	80.4	91.1	49.0	62.5	54.7	67.5	76.4	76.4	90.3

Table 3: **Comparison with state-of-the-art models on image understanding.** This table denotes “MathV” for MathVista and “RW-QA” for RealWorldQA. **Bold** and underlined are the best and second-best results for each task.

temporal context across various video tasks, and *ii*) our proposed training pipeline and data mixture enable seamless generalization from mobile-friendly to large-scale Video LLMs.

4.3 Image Understanding Results

We also compare SF-LLaVA-1.5 against recent multimodal LLMs on image understanding, as shown in Table 3, highlighting the following observations.

First, SF-LLaVA-1.5 excels at smaller model scales. Similar to video, SF-LLaVA-1.5’s 1B and 3B models achieve competitive results across image benchmarks. Specifically, SF-LLaVA-1.5-1B outperforms InternVL2.5-1B by +3.5% on AI2D and +7.8% on MathVista, even though we use less than 30% of their input resolution. When compared to MolmoE-1B, our model clearly wins on MMMU (+5.6%), MathVista (+17.0%) and DocVQA (+7.7%), although MolmoE-1B is a specialist model optimized for image understanding. At the 3B scale, SF-LLaVA-1.5-3B also demonstrates superior results, (*e.g.*, outperforming Phi-3.5-Vision-4B by +1.7% on MMMU, +14.7% on MathVista and +1.0% on TextVQA).

Second, SF-LLaVA-1.5 outperforms strong baselines at the 7B scale, except for InternVL2.5 and Qwen2-VL. Using MM1.5-7B as an example, SF-LLaVA-1.5 achieves better results across benchmarks (*e.g.*, +7.2% on MMMU, +12.5% on MM-Vet, and +12.9% on OCRBench). We are impressed by the superior results of InternVL2.5 and Qwen2-VL, especially on TextRich. We hypothesize it is due to our (*i*) lower input resolution (*e.g.*, 2.36M vs. 9.63M of InternVL2.5), (*ii*) fewer training stages (*e.g.*, 2 vs. 3 of Qwen2-VL) and (*iii*) frozen vision encoder. This aligns with prior findings (Zhang et al., 2025b) that, when the model size gets larger, higher input resolution and more training stages with fully tunable parameters are pivotal for improving the image performance. Given that our model is video-centric and these enhancements significantly increase training costs, we leave their exploration for future work.

Third, SF-LLaVA-1.5’s image capability benefits from joint video-image training. SF-LLaVA-1.5, jointly optimized on video and image data, outperforms SF-LLaVA-1.5-Image on most benchmarks. To confirm the improvements are not solely due to longer training, we conduct a second-stage training for SF-LLaVA-1.5-Image using only image data. However, the performance gap remains, indicating that joint training is the primary factor. Additionally, the improvements are more significant on Knowledge and General benchmarks (*e.g.*, +1.2% on MMMU and +10.1% on MM-Vet at the 1B scale). We hypothesize this is

Structure	Video-MME (w/o sub)				LongVideoBench (val)
	(short)	(med)	(long)	(avg)	
Group-based SlowFast (GSF)	64.4	52.8	46.1	54.4	52.7
Interleaved SlowFast (ISF)	64.7	52.4	45.3	54.1	52.3

Table 4: Comparison between GSF and ISF on video understanding.

Slow Frames N^{slow}	Fast Frames N^{fast}	Total Frames N	Input Token #	Video-MME (w/o sub)				LongVideoBench (val)
				(short)	(med)	(long)	(avg)	
32	0	32	7K	62.0	50.4	44.1	52.1	52.4
48	0	48	10K	64.9	51.1	45.0	53.7	52.5
64	0	64	14K	64.3	51.0	45.5	53.6	52.2
128	0	128	28K	63.0	53.3	46.0	54.1	52.3
0	128	128	2K	59.3	49.7	44.3	51.1	49.7
32	128	128	9K	64.4	52.8	46.1	54.4	52.7

Table 5: Results of SF-LLaVA-1.5 with different design choices on video understanding.

because our video data mainly comes from lifestyle scenarios, which could not directly benefit text-rich tasks. A deeper analysis of joint training will be provided in Sec. 4.4.2.

4.4 Ablation Studies

All ablation studies are conducted on the 1B model with our default settings (Sec. 4.1). To save training costs, models are trained on 1.2M image and 600K video samples, randomly selected from our original data mixture (Appendix A.1). The performance is evaluated on Video-MME and LongVideoBench to cover both short and long videos.

4.4.1 Design Choices of SlowFast

Group-based SlowFast (GSF) vs. Interleaved SlowFast (ISF). We introduced these SlowFast structures in Sec. 3.1.1 and report their video understanding results in Table 4. GSF and ISF perform comparably on Video-MME (54.4% vs. 54.1% on average) and LongVideoBench (52.7% vs. 52.3%), suggesting that SF-LLaVA-1.5 is not sensitive to this design choice. This highlights the general effectiveness of the SlowFast approach in improving long-form video understanding. Since GSF consistently achieves superior performance across benchmarks, we adopt it as the default SlowFast structure in this paper.

Effect of the Slow and Fast Pathways. *First*, we assess the necessity of the Slow and Fast pathways by removing them individually. Table 5 shows that SF-LLaVA-1.5 outperforms both Slow-only (row 1 vs. row 6) and Fast-only (row 5 vs. row 6) models. This is expected since they use fewer input frames or tokens than the full model. *Second*, we test if SlowFast remains more effective when the Slow-only model uses a comparable number of input tokens (e.g., 48 frames with $\sim 10K$ tokens). The results (row 2 vs. row 6) demonstrate that SlowFast outperforms this baseline (e.g., +1.1% on Video-MME long), indicating that the improvements are not merely due to using more information. *Third*, we argue that SlowFast enhances both computational efficiency and long-range temporal modeling. We verify this by comparing SlowFast with the Slow-only model that uses the same number of input frames (e.g., $N^{slow} = N = 128$). The results (row 4 vs. row 6) show that SlowFast maintains superior performance while using only $\sim 30\%$ (9K vs. 28K) of its input tokens.

SlowFast vs. Other Video Projectors. We compare SlowFast with existing video projectors in Table 6. Specifically, we apply 2×2 average pooling in Spatial Pooling and Dynamic Compressor and follow Apollo (Zohar et al., 2024) by using 16 tokens per frame in Q-Former and Perceiver Resampler. All models process up to 128 input frames. Compared to Spatial Pooling and Dynamic Compressor, SlowFast improves runtime by 25% while surpassing them across all benchmarks. It also significantly outperforms Q-Former and Perceiver Resampler, which use fixed-length tokens for information compression, limiting their ability to handle long video sequences. Moreover, Q-Former and Perceiver Resampler introduce additional parameters (e.g., BERT-Base in Q-Former), which restrict their advantage in runtime efficiency. These results demonstrate SlowFast’s effectiveness in balancing strong video performance and computational efficiency.

Video Projector	Input Token #	Runtime (per video)	Video-MME (w/o sub)				LongVideoBench (val)
			(short)	(med)	(long)	(avg)	
Spatial Pooling (Xu et al., 2024a)	28K	2.40s	63.3	51.8	45.7	53.6	51.7
Dynamic Compressor (Liu et al., 2025)	28K	2.45s	63.5	52.4	45.8	53.9	52.3
Qformer (Li et al., 2023a)	2K	1.59s	46.7	43.0	38.4	42.7	45.0
Perceiver Resampler (Jaegle et al., 2021)	2K	1.50s	52.8	45.9	43.0	47.2	48.4
SlowFast	9K	1.79s	64.4	52.8	46.1	54.4	52.7

Table 6: **Comparison between SlowFast and existing video projectors on video understanding.** All models take 128 frames as inputs. The runtime (per video) measures only the model’s forward pass on a single H100-80G GPU, using the LongVideoBench dataset.

Ratio	Video Benchmarks					Image Benchmarks			
	Video-MME (w/o sub)				LongVideoBench (val)	MMMU (val)	RW-QA	OCRBench (test)	TextVQA (val)
	(short)	(med)	(long)	(avg)					
0 : 2	60.6	47.3	41.6	49.8	48.2	38.4	58.4	66.9	69.4
1 : 0	63.4	51.8	44.3	53.2	52.0	39.4	55.8	61.6	64.2
2 : 0	63.7	52.0	44.9	53.5	52.4	39.3	56.6	61.9	64.6
1 : 0.5	65.1	50.1	45.9	53.7	52.3	44.0	59.0	66.2	69.2
1 : 1	64.8	50.5	45.3	53.5	52.1	39.9	58.5	68.3	69.5
1 : 2	64.4	52.8	46.1	54.4	52.7	40.0	59.1	68.2	69.7
1 : 3	63.7	52.3	46.0	54.0	52.5	40.7	58.8	68.0	69.3

Table 7: **Results of using different video-to-image data ratios in joint training.**

4.4.2 Design Choices of Model Training

Effect of Video-to-Image Ratio in Joint Training. We examine the optimal video-to-image ratio by fixing video samples at 600K and evaluating the impact of varying image samples. Specifically, we explore the following ratios $\{0:2, 1:0, 2:0, 1:0.5, 1:1, 1:2, 1:3\}$, where a ratio of “1:0” and “2:0” use only video data. Results are shown in Table 7 with the following findings. *First*, training with only video data clearly decreases the performance in image understanding (row 2 vs. row 4), with a substantial drop on text-rich benchmarks (e.g., -5.0% on TextVQA). *Second*, joint video-image training generally improves SF-LLaVA-1.5’s video capability (row 2 vs. row 6), such as on Video-MME (53.2% vs. 54.4% on average). *Third*, increasing the proportion of image data does not always lead to better video results (row 6 vs. row 7). *Fourth*, a video-to-image ratio of “1:2” achieves the best overall performance in video and image understanding, which we adopt in our final data mixture.

5 Limitations

First, SF-LLaVA-1.5 prefers FPS sampling, but falls back to uniform sampling when the video duration exceeds the maximum frame capacity (i.e., 128 in this paper). This approach may miss some key frames in long-form videos and mislead the model about a video’s playback speed (e.g., A ten-minute video and a one-hour video have the same number of input frames). Developing an efficient memory model to summarize the long-range context is a promising direction (Xu et al., 2021). We can also input extra information (e.g., frame timestamps) to enhance the temporal modeling. *Second*, SF-LLaVA-1.5’s performance can be further improved by tuning all parameters, including the visual encoder. However, we found this is not trivial for Long Video LLMs due to the high GPU memory cost of caching the activation values. Future studies could explore the integration of memory-saving techniques, such as Stochastic BP (Cheng et al., 2022). More analysis will be discussed in Appendix A.4.

6 Conclusion

Building upon the insights of SlowFast-LLaVA (Xu et al., 2024b), we introduce SlowFast-LLaVA-1.5, a new family of token-efficient Video LLMs for long-form video understanding. While SlowFast-LLaVA adapts the two-stream SlowFast inputs into a training-free model, this work explores further improvements by building a supervised fine-tuning pipeline with a high-quality data mixture. Our model family, ranging from 1B to 7B parameters, focuses on developing lightweight models that are both compact for potential edge deployment and powerful for various video tasks. Experimental results demonstrate that SlowFast-LLaVA-1.5 achieves superior performance across video benchmarks while maintaining strong image capabilities. We hope our work inspires the community to develop efficient yet robust Long Video LLMs based on open-source datasets.

Acknowledgments

We thank Yizhe Zhang, Feng Tang, Jesse Allardice, Jiaming Hu, Yihao Qian, Zhe Fu, Hong-You Chen, Wentao Wu, Junting Pan, Bowen Zhang, Yanghao Li for their kind help.

References

- <https://huggingface.co/datasets/lmms-lab/llava-onevision-data/viewer/k12.printing.wendlerc/renderedtext>.
- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv:2404.14219*, 2024.
- Manoj Acharya, Kushal Kafle, and Christopher Kanan. Tallyqa: Answering complex counting questions. In *AAAI*, 2019.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. GPT-4 technical report. *arXiv:2303.08774*, 2023.
- Gilles Baechler, Srinivas Sunkara, Maria Wang, Fedir Zubach, Hassan Mansoor, Vincent Etter, Victor Cărbune, Jason Lin, Jindong Chen, and Abhanshu Sharma. Screenai: A vision-language model for ui and infographics understanding. *arXiv:2402.04615*, 2024.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv:2502.13923*, 2025.
- Jonas Belouadi, Simone Paolo Ponzetto, and Steffen Eger. DeTikZify: Synthesizing graphics programs for scientific figures and sketches with TikZ. In *NeurIPS*, 2024.
- Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gómez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *ICCV*, 2019.
- Jie Cao and Jing Xiao. An augmented benchmark dataset for geometric question answering through dual parallel text encoding. In *COLING*, 2022.
- Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. Honeybee: Locality-enhanced projector for multimodal llm. In *CVPR*, 2024.
- Shuaichen Chang, David Palzer, Jialin Li, Eric Fosler-Lussier, and Ningchuan Xiao. Mapqa: A dataset for question answering on choropleth maps. In *NeurIPS Workshop*, 2022.
- Jiaqi Chen, Tong Li, Jinghui Qin, Pan Lu, Liang Lin, Chongyu Chen, and Xiaodan Liang. Unigeo: Unifying geometry logical reasoning via reformulating mathematical expression. *arXiv:2212.02746*, 2022.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. In *ECCV*, 2024a.
- Lin Chen, Xilin Wei, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Bin Lin, Zhenyu Tang, Li Yuan, Yu Qiao, Dahua Lin, Feng Zhao, and Jiaqi Wang. Sharegpt4video: Improving video understanding and generation with better captions. *arXiv:2406.04325*, 2024b.
- Shimin Chen, Xiaohan Lan, Yitian Yuan, Zequn Jie, and Lin Ma. Timemarker: A versatile video-llm for long and short video understanding with superior temporal localization ability. *arXiv:2411.18211*, 2024c.

- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang. Tabfact : A large-scale dataset for table-based fact verification. In *ICLR*, 2020.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv:1504.00325*, 2015.
- Yukang Chen, Fuzhao Xue, Dacheng Li, Qinghao Hu, Ligeng Zhu, Xiuyu Li, Yunhao Fang, Haotian Tang, Shang Yang, Zhijian Liu, et al. Longvila: Scaling long-context visual language models for long videos. *arXiv:2408.10188*, 2024d.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv:2412.05271*, 2024e.
- Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv:2404.16821*, 2024f.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, et al. Finqa: A dataset of numerical reasoning over financial data. *arXiv:2109.00122*, 2021.
- Feng Cheng and Gedas Bertasius. Tallformer: Temporal action localization with a long-memory transformer. In *ECCV*, 2022.
- Feng Cheng, Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Li, and Wei Xia. Stochastic backpropagation: A memory efficient strategy for training video models. In *CVPR*, 2022.
- Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, and Lidong Bing. VideoLLaMA 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv:2406.07476*, 2024.
- Zhoujun Cheng, Haoyu Dong, Zhiruo Wang, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou, and Dongmei Zhang. Hitab: A hierarchical table dataset for question answering and natural language generation. *arXiv:2108.06712*, 2021.
- Erfei Cui, Yinan He, Zheng Ma, Zhe Chen, Hao Tian, Weiyun Wang, Kunchang Li, Yi Wang, Wenhui Wang, Xizhou Zhu, Lewei Lu, Tong Lu, Yali Wang, Limin Wang, Yu Qiao, and Jifeng Dai. Sharegpt-4o: Comprehensive multimodal annotations with gpt-4o, 2024. URL <https://sharegpt4o.github.io/>.
- Brian Davis, Bryan Morse, Scott Cohen, Brian Price, and Chris Tensmeyer. Deep visual template-free form parsing. In *ICDAR*, 2019.
- Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *arXiv:2409.17146*, 2024.
- Zane Durante, Qiuyuan Huang, Naoki Wake, Ran Gong, Jae Sung Park, Bidipta Sarkar, Rohan Taori, Yusuke Noda, Demetri Terzopoulos, Yejin Choi, et al. Agent AI: Surveying the horizons of multimodal interaction. *arXiv:2401.03568*, 2024.
- Jiajun Fei, Dian Li, Zhidong Deng, Zekun Wang, Gang Liu, and Hui Wang. Video-ccam: Enhancing video-language understanding with causal cross-attention masks for short and long videos. *arXiv:2408.14023*, 2024.

- Chaoyou Fu, Yuhan Dai, Yondong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. *arXiv:2405.21075*, 2024.
- Jiahui Gao, Renjie Pi, Jipeng Zhang, Jiacheng Ye, Wanjun Zhong, Yufei Wang, Lanqing Hong, Jianhua Han, Hang Xu, Zhenguo Li, et al. G-llava: Solving geometric problem with multi-modal large language model. *arXiv:2312.11370*, 2023.
- Lishuai Gao, Yujie Zhong, Yingsen Zeng, Haoxian Tan, Dengjie Li, and Zheng Zhao. Linvt: Empower your image-level large language model to understand videos. *arXiv:2412.05185*, 2024.
- Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *CVPR*, 2018.
- Rujun Han, Peng Qi, Yuhao Zhang, Lan Liu, Juliette Burger, William Yang Wang, Zhiheng Huang, Bing Xiang, and Dan Roth. Robustqa: Benchmarking the robustness of domain adaptation for open-domain question answering. In *ACL Findings*, 2023.
- Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and CV Jawahar. Icdar2019 competition on scanned receipt ocr and information extraction. In *ICDAR*, 2019.
- Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. In *ICML*, 2021.
- Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *CVPR*, 2017.
- Kushal Kafle, Scott Cohen, Brian Price, and Christopher Kanan. Dvqa: Understanding data visualizations via question answering. In *CVPR*, 2018.
- Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Ákos Kádár, Adam Trischler, and Yoshua Bengio. Figureqa: An annotated figure dataset for visual reasoning. *arXiv:1710.07300*, 2017.
- Mehran Kazemi, Hamidreza Alvari, Ankit Anand, Jialin Wu, Xi Chen, and Radu Soricut. Geomverse: A systematic evaluation of large models for geometric reasoning. *arXiv:2312.12241*, 2023.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images. In *ECCV*, 2016.
- Aniruddha Kembhavi, Minjoon Seo, Dustin Schwenk, Jonghyun Choi, Ali Farhadi, and Hannaneh Hajishirzi. Are you smarter than a sixth grader? textbook question answering for multimodal machine comprehension. In *CVPR*, 2017.
- Douwe Kiela, Hamed Firooz, Aravind Mohan, Vedanuj Goswami, Amanpreet Singh, Pratik Ringshia, and Davide Testuggine. The hateful memes challenge: Detecting hate speech in multimodal memes. *NeurIPS*, 2020.
- Geewook Kim, Teakgyu Hong, Moonbin Yim, JeongYeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoo Yun, Dongyoon Han, and Seunghyun Park. Ocr-free document understanding transformer. In *ECCV*, 2022.
- Wonkyun Kim, Changin Choi, Wonseok Lee, and Wonjong Rhee. An image grid can be worth a video: Zero-shot video question answering using a vlm. *arXiv:2403.18406*, 2024.
- Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 2018.

- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. What matters when building vision-language models? *arXiv:2405.02246*, 2024a.
- Hugo Laurençon, Léo Tronchon, and Victor Sanh. Unlocking the conversion of web screenshots into html code with the websight dataset. *arXiv:2403.09029*, 2024b.
- Hosu Lee, Junho Kim, Hyunjun Kim, and Yong Man Ro. Look every frame all at once: Video-ma2mba for efficient long-form video understanding with multi-axis gradient checkpointing. *arXiv:2411.19460*, 2024a.
- Seon-Ho Lee, Jue Wang, Zhikang Zhang, David Fan, and Xinyu Li. Video token merging for long-form video understanding. *arXiv:2410.23782*, 2024b.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. *arXiv:2408.03326*, 2024a.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023a.
- Kunchang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. VideoChat: Chat-centric video understanding. *arXiv:2305.06355*, 2023b.
- Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, Limin Wang, and Yu Qiao. MVBench: A comprehensive multi-modal video understanding benchmark. *arXiv:2311.17005*, 2023c.
- Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong Feng, Lingpeng Kong, and Qi Liu. Multimodal arxiv: A dataset for improving scientific comprehension of large vision-language models. *arXiv:2403.00231*, 2024b.
- Qingyun Li, Zhe Chen, Weiyun Wang, Wenhai Wang, Shenglong Ye, Zhenjiang Jin, et al. Omnicorpus: A unified multimodal corpus of 10 billion-level images interleaved with text. *arXiv:2406.08418*, 2024c.
- Xinhao Li, Yi Wang, Jiashuo Yu, Xiangyu Zeng, Yuhao Zhu, Haian Huang, Jianfei Gao, Kunchang Li, Yinan He, Chenting Wang, et al. Videochat-flash: Hierarchical compression for long-context video modeling. *arXiv:2501.00574*, 2024d.
- Zhuowan Li, Xingrui Wang, Elias Stengel-Eskin, Adam Kortylewski, Wufei Ma, Benjamin Van Durme, and Alan L Yuille. Super-clevr: A virtual benchmark to diagnose domain robustness in visual reasoning. In *CVPR*, 2023d.
- Bin Lin, Bin Zhu, Yang Ye, Munan Ning, Peng Jin, and Li Yuan. Video-LLaVA: Learning united visual representation by alignment before projection. *arXiv:2311.10122*, 2023a.
- Ji Lin, Hongxu Yin, Wei Ping, Yao Lu, Pavlo Molchanov, Andrew Tao, Huizi Mao, Jan Kautz, Mohammad Shoeybi, and Song Han. Vila: On pre-training for visual language models. *arXiv:2312.07533*, 2023b.
- Junming Lin, Zheng Fang, Chi Chen, Zihao Wan, Fuwen Luo, Peng Li, Yang Liu, and Maosong Sun. Streamingbench: Assessing the gap for mllms to achieve streaming video understanding. *arXiv:2411.03628*, 2024.
- Ziyi Lin, Chris Liu, Renrui Zhang, Peng Gao, Longtian Qiu, Han Xiao, Han Qiu, Chen Lin, Wenqi Shao, Keqin Chen, et al. Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models. *arXiv:2311.07575*, 2023c.
- Fangyu Liu, Guy Emerson, and Nigel Collier. Visual spatial reasoning. *ACL*, 2023a.
- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Aligning large multi-modal model with robust instruction tuning. *arXiv:2306.14565*, 2023b.

- Fuxiao Liu, Xiaoyang Wang, Wenlin Yao, Jianshu Chen, Kaiqiang Song, Sangwoo Cho, Yaser Yacoob, and Dong Yu. Mmc: Advancing multimodal chart understanding with large-scale instruction tuning. *arXiv:2311.10774*, 2023c.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv:2310.03744*, 2023d.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023e.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. LLaVA-NeXT: Improved reasoning, ocr, and world knowledge, 2024a. URL <https://llava-vl.github.io/blog/2024-01-30-llava-next/>.
- Jiajun Liu, Yibing Wang, Hanghang Ma, Xiaoping Wu, Xiaoqi Ma, Xiaoming Wei, Jianbin Jiao, Enhua Wu, and Jie Hu. Kangaroo: A powerful video-language model supporting long-context video input. *arXiv:2408.15542*, 2024b.
- Yuanxin Liu, Shicheng Li, Yi Liu, Yuxiang Wang, Shuhuai Ren, Lei Li, Sishuo Chen, Xu Sun, and Lu Hou. Tempcompass: Do video llms really understand videos? *arXiv:2403.00476*, 2024c.
- Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang, Wenwen Yu, Chunyuan Li, Xu-Cheng Yin, Cheng-Lin Liu, Lianwen Jin, and Xiang Bai. Ocrbench: on the hidden mystery of ocr in large multimodal models. *Science China Information Sciences*, 2024d.
- Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, Shang Yang, Haocheng Xi, Shiyi Cao, Yuxian Gu, Dacheng Li, et al. Nvila: Efficient frontier visual language models. *arXiv:2412.04468*, 2024e.
- Zuyan Liu, Yuhao Dong, Ziwei Liu, Winston Hu, Jiwen Lu, and Yongming Rao. Oryx MLLM: On-demand spatial-temporal understanding at arbitrary resolution. *ICLR*, 2025.
- Pan Lu, Ran Gong, Shibiao Jiang, Liang Qiu, Siyuan Huang, Xiaodan Liang, and Song-Chun Zhu. Inter-gps: Interpretable geometry problem solving with formal language and symbolic reasoning. In *ACL*, 2021a.
- Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and Song-Chun Zhu. Iconqa: A new benchmark for abstract diagram understanding and visual language reasoning. In *NeurIPS Track on Datasets and Benchmarks*, 2021b.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafford, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *NeurIPS*, 2022a.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. *arXiv:2209.14610*, 2022b.
- Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. In *ICLR*, 2024.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. VideoGPT+: Integrating image and video encoders for enhanced video understanding. *arXiv:2406.09418*, 2024a.
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-ChatGPT: Towards detailed video understanding via large vision and language models. In *ACL*, 2024b.
- Andrés Marafioti, Orr Zohar, Miquel Farré, Merve Noyan, Elie Bakouch, Pedro Cuenca, Cyril Zakka, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, Vaibhav Srivastav, Joshua Lochner, Hugo Larcher, Mathieu Morlon, Lewis Tunstall, Leandro von Werra, and Thomas Wolf. Smolvlm: Redefining small and efficient multimodal models. 2025.

- U-V Marti and Horst Bunke. The iam-database: an english sentence database for offline handwriting recognition. *IJDAR*, 2002.
- Ahmed Masry, Do Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. ChartQA: A benchmark for question answering about charts with visual and logical reasoning. In *ACL Findings*, 2022.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *CVPR*, 2021.
- Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. Infographicvqa. In *WACV*, 2022.
- Brandon McKinzie, Zhe Gan, Jean-Philippe Fauconnier, Sam Dodge, Bowen Zhang, Philipp Dufter, Dhruvi Shah, Xianzhi Du, Futang Peng, Floris Weers, et al. MM1: Methods, analysis & insights from multimodal llm pre-training. *arXiv:2403.09611*, 2024.
- Nitesh Methani, Pritha Ganguly, Mitesh M. Khapra, and Pratyush Kumar. Plotqa: Reasoning over scientific plots. In *WACV*, 2020.
- Anand Mishra, Karteek Alahari, and CV Jawahar. Scene text recognition using higher order language priors. In *BMVC*, 2012.
- Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. Ocr-vqa: Visual question answering by reading text in images. In *ICDAR*, 2019.
- Harold Mouchere, Christian Viard-Gaudin, Dae Hwan Kim, Jin Hyung Kim, and Utpal Garain. Crohme2011: Competition on recognition of online handwritten mathematical expressions. In *ICDAR*, 2011.
- Jason Obeid and Enamul Hoque. Chart-to-text: Generating natural language descriptions for charts by adapting the transformer model. *arXiv:2010.09142*, 2020.
- OpenAI. Gpt-4v, 2023. URL <https://openai.com/index/gpt-4v-system-card/>.
- OpenAI. Gpt-4o, 2024. URL <https://openai.com/index/hello-gpt-4o/>.
- Panupong Pasupat and Percy Liang. Compositional semantic parsing on semi-structured tables. *arXiv:1508.00305*, 2015.
- Renjie Pi, Jianshu Zhang, Jipeng Zhang, Rui Pan, Zhehai Chen, and Tong Zhang. Image textualization: An automatic framework for creating accurate and detailed image descriptions. *arXiv:2406.07502*, 2024.
- Viorica Pătrăucean, Lucas Smaira, Ankush Gupta, Adrià Recasens Contente, Larisa Markeeva, Dylan Banarse, Skanda Koppula, Joseph Heyward, Mateusz Malinowski, Yi Yang, Carl Doersch, Tatiana Matejovicova, Yury Sulsky, Antoine Miech, Alex Frechette, Hanna Klimczak, Raphael Koster, Junlin Zhang, Stephanie Winkler, Yusuf Aytar, Simon Osindero, Dima Damen, Andrew Zisserman, and João Carreira. Perception test: A diagnostic benchmark for multimodal video models. In *NeurIPS*, 2023. URL <https://openreview.net/forum?id=HYEGXFnPoq>.
- Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Shuangrui Ding, Dahua Lin, and Jiaqi Wang. Streaming long video understanding with large language models. *NeurIPS*, 2024.
- Ruchit Rawal, Khalid Saifullah, Ronen Basri, David Jacobs, Gowthami Somepalli, and Tom Goldstein. Cinepile: A long video question answering dataset and benchmark. *arXiv:2405.08813*, 2024.
- Michael S Ryoo, Honglu Zhou, Shrikant Kendre, Can Qin, Le Xue, Manli Shu, Silvio Savarese, Ran Xu, Caiming Xiong, and Juan Carlos Niebles. xgen-mm-vid (blip-3-video): You only need 32 tokens to represent a video even in vlms. *arXiv:2410.16267*, 2024.
- Minjoon Seo, Hannaneh Hajishirzi, Ali Farhadi, Oren Etzioni, and Clint Malcolm. Solving geometry problems: Combining text and diagram interpretation. In *EMNLP*, 2015.

- Xiaoqian Shen, Yunyang Xiong, Changsheng Zhao, Lemeng Wu, Jun Chen, Chenchen Zhu, Zechun Liu, Fanyi Xiao, Balakrishnan Varadarajan, Florian Bordes, et al. Longvu: Spatiotemporal adaptive compression for long video-language understanding. *arXiv:2410.17434*, 2024.
- Min Shi, Fuxiao Liu, Shihao Wang, Shijia Liao, Subhashree Radhakrishnan, De-An Huang, Hongxu Yin, Karan Sapra, Yaser Yacoob, Humphrey Shi, et al. Eagle: Exploring the design space for multimodal llms with mixture of encoders. *arXiv:2408.15998*, 2024.
- Chenglei Si, Yanzhe Zhang, Zhengyuan Yang, Ruibo Liu, and Diyi Yang. Design2code: How far are we from automating front-end engineering? *arXiv:2403.03163*, 2024.
- Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. Textcaps: a dataset for image captioning with reading comprehension. In *ECCV*, 2020.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *CVPR*, 2019.
- Amanpreet Singh, Guan Pang, Mandy Toh, Jing Huang, Wojciech Galuba, and Tal Hassner. TextOCR: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text. In *CVPR*, 2021.
- Enxin Song, Wenhao Chai, Guanhong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, et al. Moviechat: From dense token to sparse memory for long video understanding. In *CVPR*, 2024.
- Tomasz Stanisławek, Filip Graliński, Anna Wróblewska, Dawid Lipiński, Agnieszka Kaliska, Paulina Rosalska, Bartosz Topolski, and Przemysław Biecek. Kleister: key information extraction datasets involving long documents with complex layouts. In *ICDAR*, 2021.
- Stacey Svetlichnaya. Deepform: Understand structured documents at scale—wandb. ai, 2020.
- Ryota Tanaka, Kyosuke Nishida, and Sen Yoshida. Visualmrc: Machine reading comprehension on document images. In *AAAI*, 2021.
- Benny J. Tang, Angie Boggust, and Arvind Satyanarayan. VisText: A Benchmark for Semantically Rich Chart Captioning. In *ACL*, 2023.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv:2312.11805*, 2023.
- Peter Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Adithya Jairam Vedagiri IYER, Sai Charitha Akula, Shusheng Yang, Jihan Yang, Manoj Middepogu, Ziteng Wang, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *NeurIPS*, 2025.
- Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *CVPR*, 2024.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv:2302.13971*, 2023.
- Bryan Wang, Gang Li, Xin Zhou, Zhouong Chen, Tovi Grossman, and Yang Li. Screen2words: Automatic mobile ui summarization with multimodal learning. In *UIST*, 2021.
- Haibo Wang, Bo Feng, Zhengfeng Lai, Mingze Xu, Shiyu Li, Weifeng Ge, Afshin Dehghan, Meng Cao, and Ping Huang. Streambridge: Turning your offline video large language model into a proactive streaming assistant. *arXiv:2505.05467*, 2025.
- Jiawei Wang, Liping Yuan, Yuchen Zhang, and Haomiao Sun. Tarsier: Recipes for training and evaluating large video description models. *arXiv:2407.00634*, 2024a.

- Junke Wang, Lingchen Meng, Zejia Weng, Bo He, Zuxuan Wu, and Yu-Gang Jiang. To see is to believe: Prompting gpt-4v for better visual instruction tuning. *arXiv:2311.07574*, 2023.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv:2409.12191*, 2024b.
- Weihan Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Xiaotao Gu, Shiyu Huang, Bin Xu, Yuxiao Dong, et al. Lvbenc: An extreme long video understanding benchmark. *arXiv:2406.08035*, 2024c.
- Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. Longvideobench: A benchmark for long-context interleaved video-language understanding. *NeurIPS*, 2025.
- Wenhao Wu. FreeVA: Offline mllm as training-free video assistant. *arXiv:2405.07798*, 2024.
- Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. NExT-QA: Next phase of question-answering to explaining temporal actions. In *CVPR*, 2021.
- Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. PLLaVA: Parameter-free llava extension from images to videos for video dense captioning. *arXiv:2404.16994*, 2024a.
- Mingze Xu, Yuanjun Xiong, Hao Chen, Xinyu Li, Wei Xia, Zhuowen Tu, and Stefano Soatto. Long short-term transformer for online action detection. *NeurIPS*, 2021.
- Mingze Xu, Mingfei Gao, Zhe Gan, Hong-You Chen, Zhengfeng Lai, Haiming Gang, Kai Kang, and Afshin Dehghan. SlowFast-LLaVA: A strong training-free baseline for video large language models. *arXiv:2407.15841*, 2024b.
- Zhiyang Xu, Chao Feng, Rulin Shao, Trevor Ashby, Ying Shen, Di Jin, Yu Cheng, Qifan Wang, and Lifu Huang. Vision-flan: Scaling human-labeled tasks in visual instruction tuning. *arXiv:2402.11690*, 2024c.
- Le Xue, Manli Shu, Anas Awadalla, Jun Wang, An Yan, Senthil Purushwalkam, Honglu Zhou, Viraj Prabhu, Yutong Dai, Michael S Ryoo, et al. xgen-mm (blip-3): A family of open large multimodal models. *arXiv:2408.08872*, 2024.
- Jianwei Yang, Reuben Tan, Qianhui Wu, Ruijie Zheng, Baolin Peng, Yongyuan Liang, Yu Gu, Mu Cai, Seonghyeon Ye, Joel Jang, et al. Magma: A foundation model for multimodal ai agents. *arXiv:2502.13130*, 2025.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, Qianyu Chen, Huarong Zhou, Zhensheng Zou, Haoye Zhang, Shengding Hu, Zhi Zheng, Jie Zhou, Jie Cai, Xu Han, Guoyang Zeng, Dahai Li, Zhiyuan Liu, and Maosong Sun. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv:2408.01800*, 2024.
- Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, et al. Ureader: Universal ocr-free visually-situated language understanding with multimodal large language model. *arXiv:2310.05126*, 2023.
- Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv:2310.07704*, 2023.
- Keen You, Haotian Zhang, Eldon Schoop, Floris Weers, Amanda Swearngin, Jeffrey Nichols, Yinfei Yang, and Zhe Gan. Ferret-UI: Grounded mobile ui understanding with multimodal llms. In *ECCV*, 2024.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. In *ICML*, 2024.

- Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. ActivityNet-QA: A dataset for understanding complex web videos via question answering. In *AAAI*, 2019.
- Ye Yuan, Xiao Liu, Wondimu Dikubab, Hui Liu, Zhilong Ji, Zhongqin Wu, and Xiang Bai. Syntax-aware network for handwritten mathematical expression recognition. *arXiv:2203.01601*, 2022.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *CVPR*, 2024.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, 2023.
- Boqiang Zhang, Kehan Li, Zesen Cheng, Zhiqiang Hu, Yuqian Yuan, Guanzheng Chen, Sicong Leng, Yuming Jiang, Hang Zhang, Xin Li, et al. VideoLLaMA 3: Frontier multimodal foundation models for image and video understanding. *arXiv:2501.13106*, 2025a.
- Chi Zhang, Feng Gao, Baoxiong Jia, Yixin Zhu, and Song-Chun Zhu. Raven: A dataset for relational and analogical visual reasoning. In *CVPR*, 2019.
- Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMA: An instruction-tuned audio-visual language model for video understanding. *arXiv:2306.02858*, 2023.
- Haoji Zhang, Yiqin Wang, Yansong Tang, Yong Liu, Jiashi Feng, Jifeng Dai, and Xiaojie Jin. Flash-vstream: Memory-based real-time understanding for long video streams. *arXiv:2406.08085*, 2024a.
- Haotian Zhang, Mingfei Gao, Zhe Gan, Philipp Dufter, Nina Wenzel, Forrest Huang, Dhruvi Shah, Xianzhi Du, Bowen Zhang, Yanghao Li, et al. MM1. 5: Methods, analysis & insights from multimodal llm fine-tuning. *ICLR*, 2025b.
- Pan Zhang, Xiaoyi Dong, Yuhang Zang, Yuhang Cao, Rui Qian, Lin Chen, Qipeng Guo, Haodong Duan, Bin Wang, Linke Ouyang, et al. Internlm-xcomposer-2.5: A versatile large vision language model supporting long-contextual input and output. *arXiv:2407.03320*, 2024b.
- Peiyuan Zhang, Kaichen Zhang, Bo Li, Guangtao Zeng, Jingkang Yang, Yuanhan Zhang, Ziyue Wang, Haoran Tan, Chunyuan Li, and Ziwei Liu. Long context transfer from language to vision. *arXiv:2406.16852*, 2024c.
- Renrui Zhang, Xinyu Wei, Dongzhi Jiang, Yichi Zhang, Ziyu Guo, Chengzhuo Tong, Jiaming Liu, Aojun Zhou, Bin Wei, Shanghang Zhang, Peng Gao, and Hongsheng Li. Mavis: Mathematical visual instruction tuning. In *arXiv:2407.08739*, 2024d.
- Ruohong Zhang, Liangke Gui, Zhiqing Sun, Yihao Feng, Keyang Xu, Yuanhan Zhang, Di Fu, Chunyuan Li, Alexander Hauptmann, Yonatan Bisk, et al. Direct preference optimization of video large multimodal models from language model reward. *arXiv:2404.01258*, 2024e.
- Shi-Xue Zhang, Xiaobin Zhu, Lei Chen, Jie-Bo Hou, and Xu-Cheng Yin. Arbitrary shape text detection via segmentation with probability maps. *TPAMI*, 2022.
- Yuanhan Zhang, Bo Li, haotian Liu, Yong jae Lee, Liangke Gui, Di Fu, Jiashi Feng, Ziwei Liu, and Chunyuan Li. LLaVA-NeXT: A strong zero-shot video understanding model, 2024f. URL <https://llava-vl.github.io/blog/2024-04-30-llava-next-video/>.
- Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video instruction tuning with synthetic data. *arXiv:2410.02713*, 2024g.
- Victor Zhong, Caiming Xiong, and Richard Socher. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv:1709.00103*, 2017.

Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang, Tiejun Huang, and Zheng Liu. Mlvu: A comprehensive benchmark for multi-task long video understanding. *arXiv:2406.04264*, 2024.

Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. Tat-qa: A question answering benchmark on a hybrid of tabular and textual content in finance. *arXiv:2105.07624*, 2021.

Yuke Zhu, Oliver Groth, Michael Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In *CVPR*, 2016.

Orr Zohar, Xiaohan Wang, Yann Dubois, Nikhil Mehta, Tong Xiao, Philippe Hansen-Estruch, Licheng Yu, Xiaofang Wang, Felix Juefei-Xu, Ning Zhang, et al. Apollo: An exploration of video understanding in large multimodal models. *arXiv:2412.10360*, 2024.

A Appendix

A.1 Details of Data Mixture

Mixture	Data Category	Datasets	# Samples
Image Mixture	General	LLaVA Complex Reasoning (Liu et al., 2023e), LLaVA Conversation (Liu et al., 2023e), ShareGPT-4v (Chen et al., 2024a), Coco Caption (Chen et al., 2015), LLaVA v1.5 VQAv2 OKVQA (Liu et al., 2023d), LLaVA v1.5 GQA (Liu et al., 2023d), LLaVA v1.5 A-OKVQA (Liu et al., 2023d), Pixmo-Ask-Model-Anything (Deitke et al., 2024), Image Textualization (Pi et al., 2024), ShareGPT4o (Cui et al., 2024), Vision FLAN (Xu et al., 2024c), VizWiz (Gurari et al., 2018), TallyQA (Acharya et al., 2019), Visual7W (Zhu et al., 2016), VQARAD (Lau et al., 2018), VSR (Liu et al., 2023a), Hateful Memes (Kiela et al., 2020)	4.67M
	TextRich	OCRQA (Mishra et al., 2019), Synthdog-En (Kim et al., 2022), TextCaps (Sidorov et al., 2020), TextVQA (Singh et al., 2019), DVQA (Kafle et al., 2018), ChartQA (Masry et al., 2022), DocVQA (Mathew et al., 2021), InfoVQA (Mathew et al., 2022), VisualMRC (Tanaka et al., 2021), WikiTQ (Pasupat & Liang, 2015), DeepForm (Svetlichnaya, 2020), KleisterCharity (Stanisławek et al., 2021), TabFact (Chen et al., 2020), ScreenQA (Baechler et al., 2024), TabMWP (Lu et al., 2022b), ST-VQA (Biten et al., 2019), VisText (Tang et al., 2023), HiTab (Cheng et al., 2021), ArxivQA (Li et al., 2024b), WikiSQL (Zhong et al., 2017), Chart2Text (Obeid & Hoque, 2020), RenderedText (ren), FinQA (Chen et al., 2021), TAT-QA (Zhu et al., 2021), Pixmo-Docs (Deitke et al., 2024), PlotQA (Methani et al., 2020), MMC-Instruct (Liu et al., 2023c), ArT (Zhang et al., 2022), NAF (Davis et al., 2019), SROIE (Huang et al., 2019), LRV Chart (Liu et al., 2023b), FigureQA (Kahou et al., 2017), RoBUT SQA (Han et al., 2023), Screen2Words (Wang et al., 2021), HME100K (Yuan et al., 2022), UReader (Ye et al., 2023), Diagram Image2Text (Laurençon et al., 2024a), ChromeWriting (Mouchere et al., 2011), IIIT5K (Mishra et al., 2012), IAM (Marti & Bunke, 2002), TextOCR (Singh et al., 2021), K12 Printing (k12)	
	Knowledge	AI2D (Kembhavi et al., 2016), ScienceQA (Lu et al., 2022a), GeomVerse (Kazemi et al., 2023), CLEVER (Johnson et al., 2017), IconQA (Lu et al., 2021b), RAVEN (Zhang et al., 2019), Inter-GPS (Lu et al., 2021a), WebSight (Laurençon et al., 2024b), DaTikZ (Belouadi et al., 2024), Design2Code (Si et al., 2024), TQA (Kembhavi et al., 2017), MAVIS MCollect (Zhang et al., 2024d; Li et al., 2024a), MAVIS Data Engine (Zhang et al., 2024d; Li et al., 2024a), Geo170K (Gao et al., 2023), Geo170K Align (Gao et al., 2023; Li et al., 2024a), Geometry3K (Lu et al., 2021a), GEOS (Seo et al., 2015), GeoQA+ (Cao & Xiao, 2022), MapQA (Chang et al., 2022), Super-CLEVR (Li et al., 2023d), UniGeo (Chen et al., 2022)	
Video Mixture	General	LLaVA-Hound (Zhang et al., 2024e), ShareGPT4Video (Chen et al., 2024b), VideoChatGPT-Plus (Maaz et al., 2024a), LLaVA-Video-178K (Zhang et al., 2024g), Cinepile (Rawal et al., 2024), ActivityNet-QA (Yu et al., 2019), NExT-QA (Xiao et al., 2021), Perception Test (Pătrăucean et al., 2023)	2.01M

Table 8: Details of our image and video mixtures.

A.2 Benchmarks and Metrics

All evaluations are performed using the `lmms-eval`³ toolkit, where we use the official evaluation metrics to report numbers without any filtering on the prediction outputs.

Category	Benchmark	# Videos	# QAs	Avg Duration (s)
General Video QA	Video-MME (Fu et al., 2024)	900	2700	1010
	Perception Test (val) (Pătrăucean et al., 2023)	5900	19139	23
	ActivityNet-QA (test) (Yu et al., 2019)	800	8000	180
	VCGBench (test) Maaz et al. (2024b)	800	3497	180
Long-Form Video Understanding	LongVideoBench (val) (Wu et al., 2025)	752	1337	473
	MLVU (test) (Zhou et al., 2024)	1730	3102	930
	LVBench (test) (Wang et al., 2024c)	103	1549	4101
Temporal Reasoning	TempCompass (mc) (Liu et al., 2024c)	410	7540	-
	NExT-QA (mc) (Xiao et al., 2021)	1000	8564	44

Table 9: Details of video understanding benchmarks.

³<https://github.com/EvolvingLMs-Lab/lmms-eval>.

A.2.1 Video Benchmarks

We evaluate our model on various video understanding benchmarks in Table 9.

A.2.2 Image Benchmarks

We evaluate our model on the following image understanding benchmarks:

- **Knowledge Image QA** inspects a model’s capability of answering questions requiring knowledge in specific domains. Our model is evaluated on AI2D (Kembhavi et al., 2016) and ScienceQA (Lu et al., 2022a) for science, MathVISTA Lu et al. (2024) for math and MMMU (Yue et al., 2024) for multi-discipline tasks.
- **General Image QA** evaluates the general image capability of our model. We select RealWorldQA⁴ and MMVet (Yu et al., 2024) to serve this purpose, where RealWorldQA examines a model’s capability in real-world scenarios and MMVet assesses a model’s performance for more complicated tasks.
- **TextRich Image QA** contains images embedded with dense texts. To achieve high performance, a model is expected to excel at reasoning over reading. We include OCRBench (Liu et al., 2024d), TextVQA (Singh et al., 2019) and DocVQA (Mathew et al., 2021) measuring OCR, scene text and document understanding, respectively.

A.2.3 Instantiations of SlowFast Cont’d

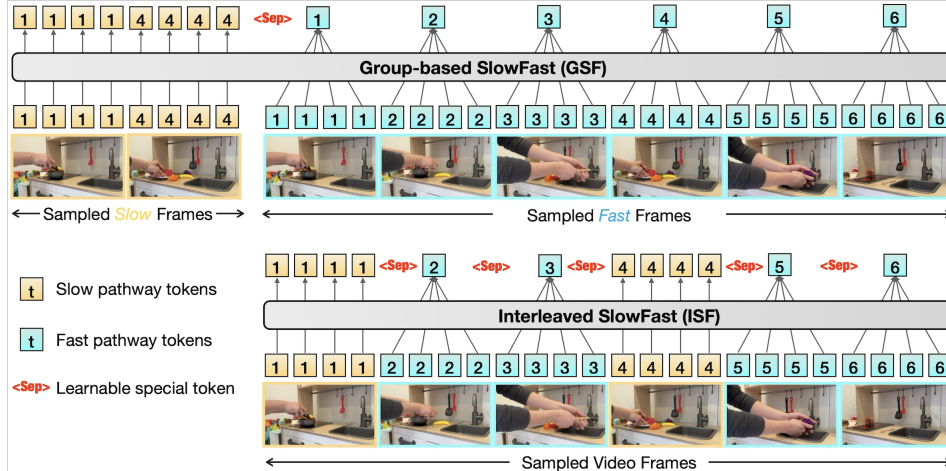


Figure 2: Visualization of Group-based SlowFast (GSF) and Interleaved SlowFast (ISF).

A.3 More Video Understanding Results

We compare with recent Video LLMs as representative examples in Table 2, and here, we include a broader group of models in Table 10. For ActivityNet-QA and VCGBench, we adopt the GPT-assisted evaluation to assess the accuracy. Specifically, we use GPT-3.5-Turbo-0125 as the judge. **It is worth noting that** our model cannot be directly compared with previous work that uses GPT-3.5-Turbo-0613 (deprecated by OpenAI) or an unknown version, since different GPT versions can significantly impact the results (Wu, 2024).

A.4 Effect of Training the Visual Encoder

By default, the visual encoder is frozen in both Stage I and II. We now assess whether training the visual encoder improves the image and video understanding performance.

We start with training Stage I, tuning the visual encoder together with other parameters (named as **SF-LLaVA-1.5-Image-E2E**). We evaluate it on image benchmarks, with results presented in Table 11. We observe that training the visual encoder significantly improves

⁴<https://huggingface.co/datasets/xai-org/RealworldQA>

Model	General VideoQA					Long-Form Video Understanding			Temporal Reasoning	
	VideoMME	VideoMME	PercepTest	ActivityNet-QA	VCGBench	LongVideoBench	MLVU	LVBench	TempComp	NEXT-QA
	(w/o sub)	(w/ sub)	(val)	(test)	(test)	(val)	(dev)	(avg)	(mc)	(test)
<i>Proprietary Models</i>										
GPT-4V (OpenAI, 2023)	59.9	63.3	-	57.0	4.06	61.3	49.2	-	-	-
GPT-4o (OpenAI, 2024)	71.9	77.2	-	-	-	66.7	64.6	30.8	70.9	-
Gemini-1.5-Flash (Team et al., 2023)	70.3	75.0	-	-	-	61.6	-	-	-	-
Gemini-1.5-Pro (Team et al., 2023)	75.0	81.3	-	57.5	-	64.0	-	33.1	69.3	-
<i>1B Model Comparison</i>										
LLaVA-OV-0.5B (Li et al., 2024a)	44.0	43.5	49.2	<u>50.5[†]</u>	<u>3.12[†]</u>	45.8	50.3	32.7 [†]	53.2 [†]	57.2
MM1.5-1B (Zhang et al., 2025b)	45.7	-	-	56.1	<u>3.14</u>	43.9	-	-	-	71.8
Apollo-1.5B (Zohar et al., 2024)	53.0	54.6	<u>61.0</u>	-	-	<u>54.1</u>	<u>63.3</u>	-	60.8	-
LinVT-Mipha-1.6B (Gao et al., 2024)	44.5	46.1	-	<u>47.5[†]</u>	-	49.7	56.2	-	45.2	71.1
InternVL2.5-2B (Chen et al., 2024e)	51.9	54.1	-	-	-	52.0	61.4	37.9 [†]	53.4 [†]	77.2[†]
Qwen2-VL-2B (Wang et al., 2024b)	<u>55.6</u>	60.4	53.9	-	-	48.7 [†]	62.7 [†]	39.4 [†]	<u>60.6[†]</u>	77.2[†]
SF-LLaVA-1.5-1B	56.6	<u>58.1</u>	61.9	<u>52.9</u>	3.27	54.3	64.3	39.7	60.5	<u>76.7</u>
<i>3B Model Comparison</i>										
Blip3-Video-4B (Ryoo et al., 2024)	-	-	-	<u>56.9[†]</u>	-	-	-	-	-	77.1
Phi-3.5-V-4B (Abdin et al., 2024)	51.5	-	-	-	-	-	-	-	-	-
V-Ma ² mba-3.1B (Lee et al., 2024a)	45.2	-	-	51.7	3.03	43.0	-	-	-	-
VILA1.5-3B (Liu et al., 2024e)	42.2	44.2	49.1	<u>50.7[†]</u>	-	42.9	44.4	-	56.1	-
MM1.5-3B (Zhang et al., 2025b)	49.5	-	-	57.9	<u>3.17</u>	45.4	-	-	-	74.7
LongVU-3.2B (Shen et al., 2024)	51.5	-	-	-	-	-	55.9	-	-	-
InternVL2-4B (Chen et al., 2024f)	53.9	57.0	53.9 [†]	-	-	53.0	59.9	<u>35.1[†]</u>	60.2 [†]	71.1 [†]
LinVT-Blip3-4B (Zohar et al., 2024)	58.3	<u>62.4</u>	-	<u>58.9[†]</u>	-	56.6	67.9	-	59.6	<u>80.1</u>
Apollo-3B (Zohar et al., 2024)	<u>58.4</u>	60.6	<u>65.0</u>	-	-	55.1	<u>68.7</u>	-	<u>62.5</u>	-
SF-LLaVA-1.5-3B	60.8	63.1	65.8	<u>55.5</u>	3.32	57.3	68.8	43.3	64.0	80.8
<i>7B Model Comparison</i>										
VideoChatGPT-7B (Maaz et al., 2024b)	-	-	-	35.2	2.42	-	-	-	43.5*	-
VideoLLaVA-7B (Lin et al., 2023a)	39.9*	41.6	-	45.3	-	39.1*	47.3*	-	49.8*	-
MovieChat+-7B (Song et al., 2024)	-	-	-	<u>48.1[†]</u>	<u>2.73[†]</u>	-	-	22.5*	-	54.8
PLLaVA-7B (Xu et al., 2024a)	-	-	-	56.3	3.12	40.2*	-	-	-	-
Tarsier-7B (Wang et al., 2024a)	-	-	-	<u>59.5</u>	-	-	-	-	-	71.6
LLaVA-Next-Video-7B (Zhang et al., 2024f)	-	-	-	<u>53.5[†]</u>	<u>3.26[†]</u>	-	-	-	-	-
VideoChat2-HD-7B (Li et al., 2023c)	45.3	55.7	47.3	-	-	3.10	-	-	48.8*	79.5
VideoLLaMA2-7B (Cheng et al., 2024)	47.9	50.3	51.4	<u>50.2[†]</u>	<u>3.13[†]</u>	-	48.5*	-	-	-
VideoCCAM-9B (Fei et al., 2024)	53.9	56.1	-	<u>59.7[†]</u>	-	-	63.1	-	-	-
Flash-VStream-7B (Zhang et al., 2024a)	-	-	-	<u>51.9[†]</u>	-	-	-	-	-	61.6
VILA-1.5-8B (Lin et al., 2023b)	-	-	41.8	<u>54.3[†]</u>	-	-	-	-	58.8*	-
TimeMaker-8B (Chen et al., 2024c)	57.3	-	-	-	-	56.3	49.2	41.3	60.4	-
LongVA-7B (Zhang et al., 2024c)	52.6	54.3	-	-	<u>3.57[†]</u>	-	56.3	-	57.0*	69.3
LongVILA-7B (Chen et al., 2024d)	60.1	65.1	58.1	<u>59.5[†]</u>	-	57.1	-	-	-	80.7
LongVU-7B (Shen et al., 2024)	60.6	-	-	-	-	-	65.4	-	-	-
XComposer-8B (Zhang et al., 2024b)	55.8	58.8	34.4	-	-	-	37.3	-	62.1*	-
VideoLLaMA2.1-7B (Cheng et al., 2024)	54.9	56.4	54.9	<u>53.0[†]</u>	-	-	57.4	36.2	56.8	75.6
LinVT-Qwen2-VL-7B (Gao et al., 2024)	63.1	63.3	-	<u>60.1[†]</u>	-	57.2	68.9	-	65.8	85.5
MM1.5-7B (Zhang et al., 2025b)	53.5	-	-	60.9	<u>3.22</u>	49.4	-	-	-	76.9
Kangaroo-8B (Liu et al., 2024b)	56.0	57.6	-	-	-	54.8	61.0	39.4	62.5	-
Oryx1.5-7B (Liu et al., 2025)	58.8	64.2	70.0	-	<u>3.62[†]</u>	56.3	67.5	39.0 [†]	58.8 [†]	81.8
LLaVA-OV-7B (Li et al., 2024a)	58.2	61.5	49.7	<u>56.6[†]</u>	<u>3.51[†]</u>	56.5	64.7	-	<u>64.2[†]</u>	79.4
LLaVA-Video-7B (Zhang et al., 2024g)	63.3	<u>69.7</u>	66.9	<u>56.5[†]</u>	<u>3.52[†]</u>	58.2	70.8	-	-	83.2
Apollo-7B (Zohar et al., 2024)	61.3	63.3	67.3	-	-	58.5	<u>70.9</u>	-	64.9	-
NVILA-8B (Liu et al., 2024e)	64.2	70.0	65.4 [†]	60.9	-	57.7	70.1	44.0 [†]	69.7[†]	82.2
InternVL2.5-8B (Chen et al., 2024e)	64.2	66.9	-	-	-	<u>60.0</u>	69.0 [†]	43.2 [†]	68.3 [†]	<u>85.0[†]</u>
Qwen2-VL-7B (Wang et al., 2024b)	63.3	69.0	62.3	-	-	<u>55.6[†]</u>	69.8 [†]	<u>44.7[†]</u>	67.9 [†]	81.2 [†]
SF-LLaVA-1.5-7B	<u>63.9</u>	65.4	<u>69.6</u>	57.0	3.35	62.5	71.5	45.3	68.8	83.3

Table 10: **Comparison with a broader group of Video LLMs on video understanding.**

[†]denotes reproduced results. *denotes results from the benchmark leaderboard. [‡]denotes results evaluated using GPT-3.5-Turbo-0613 or an unknown version, which cannot be directly compared with our results. **Bold** and underlined are the best and second-best results for each task.

the image performance, especially on Text-Rich tasks (row 1 and row 2 of each model scale). For reference, SF-LLaVA-1.5-Image-E2E-3B outperforms SF-LLaVA-1.5-Image-3B by +4.9% on OCRBench and +2.7% on TextVQA.

We move on to Stage II with fully tunable parameters but encounter the out-of-memory issue (even when we train the 1B model with batch size 1 on H100-80G GPUs). This issue arises from caching a large number of activation values from the visual encoder while extracting features from 128 input frames — that is why we do not have this problem in Stage I. Stochastic BP (Cheng et al., 2022) is proposed to solve this problem and is utilized by modern temporal action detectors (Cheng & Bertasius, 2022) for efficient end-to-end

Model	Training Visual Encoder		Knowledge				General VQA		Text-Rich		
	Stage I	Stage II	AI2D (test)	SQA (test)	MMMU (val)	MathV (testmini)	MM-Vet	RW-QA	OCRBench (test)	TextVQA (val)	DocVQA (test)
<i>1B Model Comparison</i>											
SF-LLaVA-1.5-Image-E2E-1B	✓	-	73.9	89.3	38.3	53.0	41.1	60.3	74.0	73.8	87.8
SF-LLaVA-1.5-Image-1B	✗	-	70.8	87.8	39.3	51.2	41.1	57.1	69.5	70.2	85.2
SF-LLaVA-1.5-E2E-1B	✓	✗	70.5	81.7	38.9	41.7	34.3	55.7	48.8	60.1	68.1
SF-LLaVA-1.5-1B	✗	✗	72.8	87.7	40.5	51.0	51.2	59.2	70.0	71.3	85.4
<i>3B Model Comparison</i>											
SF-LLaVA-1.5-Image-E2E-3B	✓	-	77.2	90.0	44.1	61.1	48.0	61.8	77.2	74.7	90.0
SF-LLaVA-1.5-Image-3B	✗	-	75.8	90.0	43.7	57.0	51.1	61.8	72.3	72.0	87.5
SF-LLaVA-1.5-E2E-3B	✓	✗	75.2	84.3	44.2	47.8	38.6	56.9	51.6	64.9	72.9
SF-LLaVA-1.5-3B	✗	✗	77.0	90.3	44.7	58.6	47.5	63.4	73.4	73.0	88.8
<i>7B Model Comparison</i>											
SF-LLaVA-1.5-Image-E2E-7B	✓	-	79.5	91.2	47.1	63.5	47.4	66.9	78.3	75.8	90.7
SF-LLaVA-1.5-Image-7B	✗	-	79.2	91.8	47.0	61.0	50.1	64.6	74.2	75.4	89.7
SF-LLaVA-1.5-E2E-7B	✓	✗	76.7	85.8	44.4	54.0	44.9	60.5	59.6	70.8	78.8
SF-LLaVA-1.5-7B	✗	✗	80.4	91.1	49.0	62.5	54.7	67.5	76.4	76.4	90.3

Table 11: Results of SF-LLaVA-1.5-E2E and SF-LLaVA-1.5-Image-E2E on image benchmarks, which fully train the visual encoder together with the projector and LLM.

Model	Training Visual Encoder		General VideoQA		Long-Form Video Understanding			Temporal Reasoning	
	Stage I	Stage II	VideoMME (w/o sub)	PercepTest (val)	LongVideoBench (val)	MLVU (m-avg)	LVBench (avg)	TempComp (mc)	NEXt-QA (test)
<i>1B Model Comparison</i>									
SF-LLaVA-1.5-E2E-1B	✓	✗	54.1	58.6	51.5	61.7	40.2	59.3	73.9
SF-LLaVA-1.5-1B	✗	✗	56.6	61.9	54.3	64.3	39.7	60.5	76.7
<i>3B Model Comparison</i>									
SF-LLaVA-1.5-E2E-3B	✓	✗	58.4	62.4	53.0	65.0	40.9	63.2	78.6
SF-LLaVA-1.5-3B	✗	✗	60.8	65.8	57.3	68.8	43.3	64.0	80.8
<i>7B Model Comparison</i>									
SF-LLaVA-1.5-E2E-7B	✓	✗	59.2	68.1	59.2	70.3	44.3	67.9	81.0
SF-LLaVA-1.5-7B	✗	✗	63.9	69.6	62.5	71.5	45.3	68.8	83.1

Table 12: Results of SF-LLaVA-1.5-E2E on video benchmarks, which fully trains the visual encoder together with the projector and LLM.

training. However, integrating this memory-saving technique into multimodal LLMs is non-trivial and is left for future exploration.

Finally, we test if tuning the visual encoder only in Stage I and freezing it in Stage II is effective. We train models (named as **SF-LLaVA-1.5-E2E**) based on SF-LLaVA-1.5-Image-E2E, with the visual encoder frozen. The models are evaluated on both image and video benchmarks, as shown in Table 11 and Table 12. The results show that SF-LLaVA-1.5-E2E performs significantly worse than SF-LLaVA-1.5 across all metrics. We argue that tuning the visual encoder in Stage I harms its generalization ability, leading to overfitting on image tasks and conflicts between image and video tasks. We will explore the optimal training strategy for Video LLMs in future work.

A.5 Streaming Video Understanding

Model	Real-Time Visual Understanding	Omni-Source Understanding	Contextual Understanding
SOTA in StreamingBench (Lin et al., 2024)	71.1	38.4	35.0
SF-LLaVA-1.5-7B w/ StreamBridge (Wang et al., 2025)	72.5	33.1	36.0

Table 13: Comparison of model performance on different understanding benchmarks.

We convert SlowFast-LLaVA-1.5-7B into a proactive streaming assistant by adopting the StreamBridge (Wang et al., 2025) technique without fine-tuning. The results on StreamingBench (Lin et al., 2024) are shown in Table 13. We observe that SlowFast-LLaVA-1.5-7B achieves state-of-the-art performance on Real-Time Visual Understanding and Contextual Understanding, showing our effectiveness for streaming video understanding. We fall behind on Omni-Source Understanding since our model is never trained on interleaved datasets with audio.

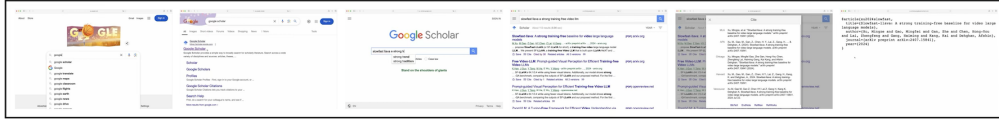
B Qualitative Results



Question 1: Describe the video in detail.

SlowFast-LLaVA-1.5: The video showcases a creative drawing process on a large white sheet of paper attached to a wooden frame. It begins with a hand holding a pink marker, sketching a simple geometric shape that evolves into a house with a triangular roof and a door. The hand then adds details like windows and a chimney, followed by a tree with a curved trunk and branches. A sun is drawn next to the tree, and clouds are added above it. The hand switches to a green marker to color in the clouds, enhancing the scene's vibrancy. The focus shifts to the sun, which is colored in blue, and the hand adds yellow details to the tree's branches, creating a sense of depth and texture. The hand then uses a purple marker to add shading to the tree, making it appear more three-dimensional. Throughout the video, the background remains consistent, featuring the wooden frame and the large white sheet of paper. The video concludes with the hand adding final touches to the tree, completing the vibrant and detailed drawing.

Figure 3: SF-LLaVA-1.5 summarizes a video with detailed caption.



Question 1: According to the video, if I want to cite the paper titled 'Slowfast-llava', what should I do?

SlowFast-LLaVA-1.5: To cite the paper titled 'Slowfast-llava', you should first search for it on Google Scholar. Then, click on the 'Cite' button next to the paper's title. This will open a citation dialog box displaying the citation details in various formats, such as MLA, APA, and Harvard. You can choose the format you need and copy the citation text to use in your work.

Question 2: What is the full name of the Slowfast-llava paper?

SlowFast-LLaVA-1.5: The full name of the Slowfast-llava paper is 'Slowfast-llava: A strong training-free baseline for video large language models.'

Question 3: What is the author list of the Slowfast-llava paper?

SlowFast-LLaVA-1.5: The author list is Xu, Mingze and Gao, Mingfei and Gan, Zhe and Chen, Hongyou and Lai, Zhengfei and Gang, Haiming and Kang, Kai and Dehghan, Afshin.

Question 4: How many citations does the Slowfast-llava paper have?

SlowFast-LLaVA-1.5: The Slowfast-llava paper has 38 citations.

Figure 4: SF-LLaVA-1.5 learns the process from the video and captures text-rich details.



Question 1: Describe the video in detail.

SlowFast-LLaVA-1.5: The video showcases a child's imaginative play in a toy kitchen setup. A person arranges colorful toy fruits and vegetables on a wooden countertop, simulating a cooking process. They wash the toys under running water, place them on a cutting board, and cut them with a toy knife. The person then arranges the cut pieces on a white plate, adding more toy fruits and vegetables, and uses a wooden spatula to mix them. After placing the plate back on the countertop, they wash their hands and continue arranging the toy fruits and vegetables. They use a black spatula to mix the contents of the pot, ensuring the toy fruits and vegetables are well combined. The video concludes with the person hanging the black spatula back on the hook and stepping away from the sink area, leaving the toy kitchen setup neatly arranged with the toy fruits and vegetables on the countertop.

Question 2: What happens after the person washes the toy fruits?

SlowFast-LLaVA-1.5: After washing the toy fruits, person arranges them on a white plate and then cuts them with a toy knife, simulating cooking.

Question 3: How many toy fruits are there in the video?

SlowFast-LLaVA-1.5: The fruits include a red strawberry, an orange carrot, a yellow banana, and a purple bunch of grapes.

Figure 5: SF-LLaVA-1.5 understands the relative sequence of different activities.