

Temporal Grounding of Activities using Multimodal Large Language Models

Stanford CS224N Custom Project

Young Chol Song
Stanford University
youngcs@stanford.edu

Abstract

Temporal grounding of activities, the identification of specific time intervals of actions within a larger event context, is a critical task in video understanding. Recent advancements in multimodal large language models (LLMs) offer new opportunities for enhancing temporal reasoning capabilities. In this paper, we evaluate the effectiveness of combining image-based and text-based large language models (LLMs) in a two-stage approach for temporal activity localization. We demonstrate that our method outperforms existing video-based LLMs. Furthermore, we explore the impact of instruction-tuning on a smaller multimodal LLM, showing that refining its ability to process action queries leads to more expressive and informative outputs, thereby enhancing its performance in identifying specific time intervals of activities. Our experimental results on the Charades-STA dataset highlight the potential of this approach in advancing the field of temporal activity localization and video understanding.

1 Introduction

We address the problem of temporal grounding of activities, also known as temporal activity localization (Gao et al., 2017) or video moment retrieval (Cui et al., 2022). This task involves determining the interval of a specific target activity within a longer context of activities, where the context is provided as a video and the target activity is described in text. However, accurately grounding activities in time poses significant challenges due to the complex and dynamic nature of real-world events, variability in human behavior, and the need for fine-grained understanding of temporal relationships.

Recent advancements in multimodal large language models (LLMs) (Liu et al., 2023b; Zhang et al., 2024) have opened new avenues for addressing this task. Many methods for temporal grounding primarily rely on specialized architectures and task-specific training (Zhang et al., 2020; Cui et al., 2022), which can limit their generalizability and scalability. In contrast, multimodal LLMs, with their broad knowledge base and reasoning capabilities (Huang et al., 2023; Li et al., 2024), have the potential to overcome some of these limitations.

In this paper, we explore the use of pretrained and tuned multimodal LLMs for temporal grounding of fine-grained activities, such as ‘putting on shoes’ or ‘pouring a drink.’ We demonstrate that a combination of image-based and text-based LLMs achieves higher accuracy in temporal activity localization than a purely video-based LLM of a similar parameter size. Additionally, we show that instruction-tuning a smaller multimodal LLM on a dataset generated by GPT-4 Vision (OpenAI, 2023) not only qualitatively enhances the model’s visual description of activities but also improves its performance in accurately identifying the time intervals of these activities.

2 Related Work

Multimodal Large Language Models Historically, multimodal vision and language systems for activity or action recognition focused primarily on solving challenges in vision. While some models

person put on shoes.




Figure 1: Comparing different model outputs of temporal activity localization. Numbers to the right are normalized intervals of the activity “person put on shoes.”

emphasized aspects of language modeling (Li et al., 2017; Wang et al., 2020; Radford et al., 2021), the role of language was often reduced to providing simple captions or labels in many cases. With the advancement of pretrained large language models (LLMs), models have emerged that fully utilize the strength of LLMs (Li et al., 2023a; Alayrac et al., 2022), and their ability to query, interact, and reason using text within the context of the image (Liu et al., 2023a,b; Chen et al., 2023a,b). In the past year, models that can directly take videos as input, such as Video-ChatGPT (Maaz et al., 2023), have started to emerge (Zhang et al., 2023; Li et al., 2023b; Reid et al., 2024), enabling pretrained models that can conduct temporal grounding and reasoning out of the box without any customization or tuning.

Temporal Activity Localization Early approaches to temporal activity localization focused on computer vision approaches using convolutional neural networks and sentence embeddings (Gao et al., 2017) or LSTMs (Anne Hendricks et al., 2017; Zhang et al., 2020) on the language side. In recent years, methods using attention mechanisms, such as the Transformer model have started to emerge (Yuan et al., 2019; Cui et al., 2022), offering improved performance. With advances in video-based LLMs, models specifically tuned for temporal grounding (Huang et al., 2023; Qian et al., 2024), and more generally tuned for both spatial and temporal grounding, have started gaining traction (Li et al., 2024).

3 Approach

We introduce a two-stage approach for temporal activity grounding. The initial stage utilizes the image-based multimodal LLM to generate detailed descriptions of actions occurring in the individual video frames. In the second stage, we feed the responses from the multimodal LLM to a text-based LLM along with the activity query to pinpoint when the activity occurs within the sequence.

3.1 Multimodal LLM (Image-based)

The first stage consists of a multimodal LLM capable of image input, such as LLaMA-Adapter (Zhang et al., 2024), LLaVA (Liu et al., 2023b), or GPT-4 Vision (OpenAI, 2023). The role of this stage is to describe the activities occurring in each individual video frame. In our current approach, we use single frames from a video as input as it provides the best quality response. We have experimented with providing different types of inputs to the multimodal LLM and discuss this further in Section 5. Examples of such outputs are shown in Figure 2. We use two different strategies when providing prompts to the multimodal LLM, which we will refer to as the *activity prompt* and the *general prompt*.

Activity prompt This type of prompt provides the context of the action during the prompting of the multimodal LLM. This enables the model to search for the particular action in the frame(s) and provide information with regards to the particular action in question. However, the activity prompt

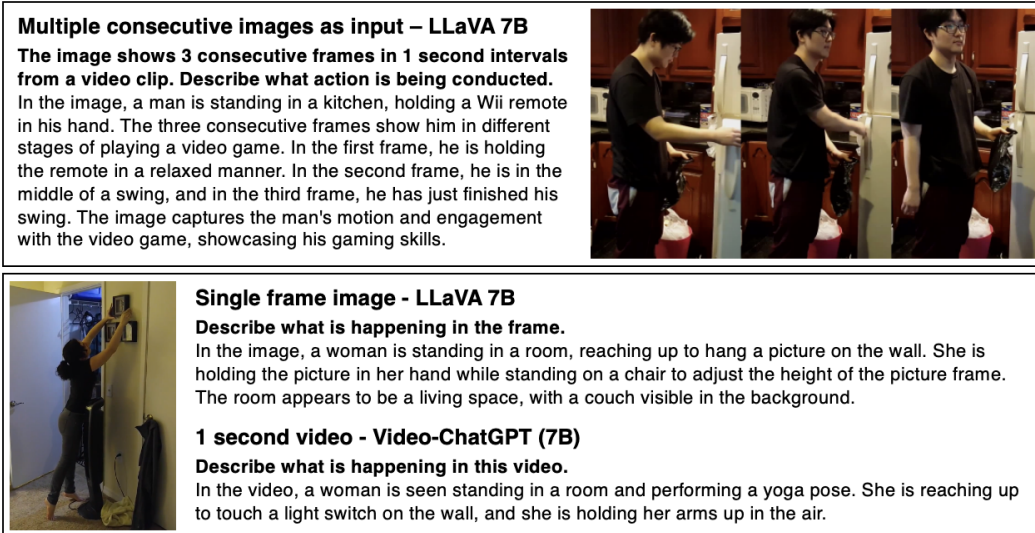


Figure 2: Responses from a multimodal LLM based on the type of input: consecutive frames (top), single frame (mid), short video (bottom)

comes with a complication where, if there are multiple activities to identify in a particular video sequence, separate prompts will be required for each frame image/activity combination.

[Image frame] This is one frame from a longer video clip. The video clip includes the action described as [activity query] however, the frame may or may not include this action. Briefly explain what action or actions the person is conducting in the frame.

Activity prompt

General prompt The second way to prompt the multimodal LLM is to provide the model with a generic prompt asking it to describe the scene. This saves on having to query multiple times for each frame image depending on the number of activities in the scene and may be useful when frames need to be processed beforehand. However, there is a tradeoff in accuracy as information about the particular activity in question may not be included in the response.

[Image frame] Describe what is happening in the frame.

General prompt

For our instruction-tuning experiments, we will evaluate models using the *general prompt* strategy to best illustrate how instruction-tuning a relatively smaller model (LLaVA 7B) with instruction-tuning data generated from a larger model (GPT-4 Vision) can provide better results due to enhanced descriptive outputs. For our model comparison experiments, we will use the *action prompt* strategy to prioritize model performance. In addition, we use the STA-Subset dataset (described in Section 4.1) for the model comparison experiments, as the image/activity combinations required for running these experiments using *action prompts* are not as prohibitive as the former.

3.2 Text-based LLM

The second stage uses a text-based LLM, which utilizes descriptions aggregated with each frame output from the first stage, and executes the task of predicting the most likely interval where the particular action would happen. The target action in question is necessarily included in this prompt.

```
The following are descriptions of actions for frames extracted 1
second apart from a video clip:
* Frame 1: [Description from frame 1]
* Frame 2: [Description from frame 2]
* ...
The action [activity query] has occurred in the video clip. What
interval is the action most likely to start and end? Provide your
best guess by providing the start and end frame numbers in json
format.
```

Template of a prompt used to query the text-based LLM

We explicitly request the responses in JSON format for parsing convenience; however, the text-based LLM is not bound to this convention and may generate non-compliant JSON responses. We employ a heuristic-based parser, which uses criteria such as the presence of specific keywords or patterns, to extract the start and end frames from the entire response text if the frame numbers from the JSON cannot be extracted.

In our experiments, we compare different text-based LLMs for our second stage processing, such as OpenAI’s GPT-4 (OpenAI, 2023), Qwen 14B, and Qwen 7B (Bai et al., 2023). As the aggregated descriptions from multiple video frames became quite lengthy, we were limited to models that could handle context lengths that covered text from our longest video. Models with context lengths of 4096 tokens or less generally had trouble processing our descriptions, which discouraged our use of LLaMA or LLaMA-based models such as Vicuna.

3.3 Video-based LLM

For comparison purposes, we also use Gemini 1.5 Pro, a multimodal LLM that can take videos directly as input. We use a prompt similar to the one used for the text-only LLM but without providing the image description context.

```
The action [activity query] has occurred in the video clip. In what
interval is the action most likely to occur? Please provide the
numbers for the start and end timestamps.
```

Template of a prompt used to query the video-based LLM

3.4 Constraints

A major limiting factor in our approach was the amount of compute resources and cost needed to process ~100k images via the multimodal LLM for each experiment. To mitigate this, while we use the Charades-STA full dataset for our instruction-tuning experiments on our 7B models, we employ the *general prompting* strategy as described above. For our model comparison experiments, we use the *activity prompting* strategy to focus on performance but select a smaller test set (STA-Subset, described in Section 4.1) during evaluations.

4 Experiments

We conduct experiments exploring the different aspects of temporal activity localization using multimodal LLMs. To demonstrate the benefits of instruction-tuning in multimodal LLMs, we create an instruction-tuning dataset from a larger multimodal LLM (e.g., GPT-4 Vision) and instruction-tune a smaller multimodal LLM (e.g., LLaVA 7B) based on this dataset to enhance its descriptive power in responding to activity-related prompts.

We also formally evaluate our two-stage approach to temporal activity localization using pretrained and instruction-tuned LLMs, where the first stage involves a multimodal LLM tasked with describing activities within frames from a video, and the second stage employing a language-only LLM for reasoning about the timing of the activity within the video clip.

We compare the performance metrics of our tuned model against similar parameter-sized video-based LLMs and evaluate our results by comparing the base multimodal LLM to our tuned model, showing our model performs better in the task of grounding activities temporally via our two-stage LLM approach. We also use LLMs of various types and sizes to compare our two-stage approach, showing how the multimodal and text-based stages perform differently when the underlying LLMs change.

4.1 Data

Instruction-tuning Dataset A set of ~19,000 activity-related queries generated by GPT-4 Vision using frames randomly extracted from the training set portion of the Charades-STA dataset (Gao et al., 2017). These queries are designed to help the model become more descriptive in responding to activity-related prompts. Each sample is split into two parts, the first prompt always asks for a description of what is happening in the frame, followed by an optional Q&A sequence asking for specific activity-related questions about the frame. The samples were created by using GPT-4 Vision to analyze the frame and provide a verbose description, and generate questions and their answers regarding activities occurring in the frame.

Charades-STA Activity Dataset We use an activity dataset with longer duration activities for our evaluations, which were also similarly used in Cui et al. (2022). The Charades (Sigurdsson et al., 2016) dataset contains 9848 videos of 157 action classes in total, with 27847 textual descriptions. The Charades-STA (Gao et al., 2017) dataset (abbreviated as *STA* in subsequent sections) is a selected portion of Charades that contains sentence temporal annotations, with a total of 6672 videos of 16128 sentence temporal annotations, where 3720 annotations from 1334 videos are used for testing. For image inputs, we extracted images from each video at a rate of 1 frame per second (fps), resulting in an average of 31 frames per video. We use the Charades-STA test set for our instruction-tuning evaluations for the purpose of comparing our results with existing video-based LLM evaluations.

Charades-STA Subset (STA-Subset) We defined a smaller subset of the Charades-STA activity dataset, denoted as *STA-Subset* in the subsequent sections, consisting of 128 videos each with one sentence temporal annotation. This results in a total of 4000 frames when extracted into images at 1 fps. The smaller subset is used due to the prohibitive cost in both compute and time in running multimodal LLMs over each frame/sentence temporal annotation combination in the original Charades-STA dataset. We use this subset for comparing between different LLMs in our two-stage approach comparison evaluations.

4.2 Evaluation method

We use the metric recall of threshold-bounded temporal intersection over union, denoted as “ $R@n$, $IoU=m$ ” (Gao et al., 2017). In all of our quantitative evaluations, we use $n = 1$, and abbreviate the metric as $R@IoU=m$. The variable m in $R@IoU$ can have fractional values, indicating the threshold of overlap of the temporal prediction of the action with the ground truth. As used in Gao et al. (2017), evaluations are conducted using $R@IoU$ thresholds of $\{0.3, 0.5, 0.7\}$, and a single value averaging the different $R@IoU$ values is noted as mean averaged IoU (mIoU).

4.3 Experimental details

Instruction-tuning For the baseline, we ran inference on a non-modified LLaVA 7B (LLaVA-1.5-7B) model using the *general prompt* strategy introduced in Section 3.1. We then instruction-tuned the model using the dataset described in Section 4.1 using the `finetune_task_lora.sh` script¹ provided by LLaVA over 1 epoch, with a learning rate of $2e^{-4}$, for a total of 1.4k steps. Total time for instruction-tuning took around 2 hours on an NVIDIA H100 GPU. Due to the longer context generated by describing each frame in the video, we chose Qwen 7B (Bai et al., 2023) for the text-based LLM as it has a context length of 8192 tokens. Using LLaMA 2 (Touvron et al., 2023), we ran into issues with context length, as the model only has half the context length of that of Qwen models.

¹https://github.com/haotian-liu/LLaVA/blob/main/scripts/v1_5/

Multimodal LLM	Text-based LLM	STA ($n=3720$) R@IoU				STA-Subset ($n=128$) R@IoU			
		0.3	0.5	0.7	mIoU	0.3	0.5	0.7	mIoU
VideoLLaMA (7B)		10.4	3.8	0.9	7.1				
VideoChat (7B)		9.0	3.3	1.3	6.5				
VideoChatGPT (7B)		20.0	7.7	1.7	13.7				
LLaVA 7B	Qwen 7B	14.1	5.6	1.8	12.0	7.0	3.1	0.0	8.6
Tuned 7B	Qwen 7B	20.1	7.9	2.5	14.9	15.6	7.0	0.8	12.7
GPT-4 Vision	GPT-4					28.9	10.9	6.3	20.7

Table 1: R@IoU for base LLaVA, instruction-tuned LLaVA using the *general prompting* strategy. For comparison, we also show video-based LLM (Huang et al., 2023), GPT-4 Vision metrics.

Two-stage LLM runs Next, we compare metrics of different combinations of multimodal LLM and text-based LLMs as listed in Table 2. We use the *activity prompting* strategy in our multimodal LLMs to optimize for performance. For comparison, we run models on the state-of-the-art task-specific model (ViGA) by Cui et al. (2022) and a state-of-the-art video-based LLM (Reid et al., 2024), Gemini.

For inference in all experiments, LLaVA 7B (LLaVA-1.5-7B²) and 16B (LLaVA-1.6-Vicuna-13B³) models were run on NVIDIA RTX A4500 and A6000 GPUs. Interactions with Qwen 14B (Qwen/Qwen1.5-14B-Chat), Qwen 7B (Qwen/Qwen1.5-7B-Chat) and Gemma 7B (google/gemma-7b-it) models were conducted via the Together AI API⁴. GPT-4 Vision (gpt-4-1106-vision-preview) and GPT-4 (gpt-4-0125-preview) models were accessed through the OpenAI API⁵. For GPT-4 Vision, we use low-res mode which resizes the images to 512px \times 512px. LLaVA models use CLIP-ViT-L-336px (Radford et al., 2021) for its vision encoders, which resize images to 336px \times 336px. We used pretrained ViGA models provided by the authors⁶ to generate metrics for ViGA using VGG and I3D features. Access to Gemini 1.5 Pro was available via Google AI Studio’s web interface⁷.

4.4 Results

Instruction-tuning metrics In Table 1, we report the R@IoU metrics for the base model and the instruction-tuned model for both the full Charades-STA dataset (STA) and the STA-Subset. For reference, we also include R@IoU metrics for video-based LLMs as described in Huang et al. (2023); Li et al. (2024), for VideoLLaMA (Zhang et al., 2023), VideoChat (Li et al., 2023b) and VideoChatGPT (Maaz et al., 2023). We observe that the instruction-tuned 7B model performs significantly better than the original LLaVA model and outperforms other video-based 7B LLM models. The instruction-tuned model is more action-focused in its descriptions and is more verbose, enabling the text-based LLM to work with more context compared with the base model. We cover qualitative analysis of instruction-tuning results with examples in Section 5.

Two-stage LLM comparisons Metrics for each combination of multimodal and text-based LLMs are shown in Table 2. Overall, the two-stage GPT-4 Vision with GPT-4 model performed best, mostly outperforming the current state-of-the-art video-based LLM, Gemini 1.5 Pro, except when R@IoU=0.7. An example of two-stage LLM output comparisons is illustrated in Figure 1. As seen in Table 2, both the multimodal LLM and text-based LLM contribute in different ways to the temporal activity localization accuracy. There are two interesting observations: in the multimodal LLM, LLaVA 7B seems to perform better than 16B. Also, there appear to be pairs of multimodal LLM and text-based LLMs that work well together and others that don’t. For example, given GPT-4 Vision’s output, Gemma 7B has a hard time providing proper time interval outputs, many of the outputs in this case are undefined, which significantly affects the resulting metric.

²<https://huggingface.co/liuhaotian/llava-v1.5-7b>

³<https://huggingface.co/liuhaotian/llava-v1.6-vicuna-13b>

⁴<https://api.together.xyz/>

⁵<https://openai.com/blog/openai-api>

⁶<https://github.com/r-cui/ViGA/tree/master>

⁷<https://aistudio.google.com/>

Multimodal LLM	Text-based LLM	STA-Subset ($n=128$) R@IoU			
		0.3	0.5	0.7	mIoU
LLaVA 7B	GPT-4	38.3	14.8	5.5	26.3
	Qwen 14B	36.7	17.2	5.5	24.8
	Qwen 7B	31.3	10.9	1.6	20.3
	Gemma 7B	35.2	11.7	3.1	23.1
LLaVA 16B	GPT-4	28.1	12.5	4.7	21.3
	Qwen 14B	24.2	9.4	1.6	18.2
	Qwen 7B	28.1	11.7	3.9	19.2
	Gemma 7B	29.7	16.4	3.1	21.0
GPT-4 Vision	GPT-4	44.5	18.0	7.0	29.8
	Qwen 14B	33.6	12.5	4.7	22.9
	Qwen 7B	40.6	16.4	5.5	25.1
	Gemma 7B	28.9	11.7	3.9	19.1
Gemini 1.5 Pro (Video)		34.4	16.4	7.8	25.0
ViGA (VGG)		56.2	35.2	14.8	35.8
ViGA (I3D)		72.7	45.3	15.6	43.9

Table 2: Comparing metrics (IoU) for different multimodal and text-based LLMs

5 Analysis

Instruction-tuning and multimodal LLM outputs We look at the qualitative effects of instruction-tuning a multimodal LLM to provide a more elaborate description of the scene. This will provide the text-based LLM with more information to better discern the interval in which the activity in question occurred. The following two text boxes show the description extracted from the scene in Figure 1 at timestep 00:19, where the individual is in the process of putting on their shoes.

In the image, a person is standing in a doorway, holding a pair of pants in their hand. They are likely in the process of putting on or taking off the pants. The person is positioned near the center of the frame, and the doorway appears to be a part of a hallway.

Base LLaVA 7B model output for Frame at 00:19 in Figure 1

In the frame provided, we see a person in the process of putting on a pair of shoes. The individual is bending forward slightly, with their right hand reaching into a shoe, likely to place their foot inside. The shoe appears to be a dark-colored, possibly black, and the person is wearing a green t-shirt and purple pants. The setting looks like a hallway, with a closed door to the left and a light-colored wall in the background. The floor is a warm-colored wood, and the overall lighting suggests an indoor setting. The person’s focus is on the task at hand, and their posture indicates they are engaged in the action of putting on their shoes.

Instruction-tuned LLaVA 7B model output for Frame at 00:19 in Figure 1

The example shows that the tuned model provides a more descriptive output regarding the scene. This includes detailed descriptions of the movements being conducted and the environment, all of which help in gathering context about when a particular activity starts or ends.

Text-based LLM outputs In the text-based LLM setting, each model has its own quirks with regards to outputs. GPT-4 LLM models, while most accurate, simply generate a strict JSON file, making for easy parsing. Gemma 7B models are similar, providing a simple start and end timestamp without explanation. The Qwen 14B and 7B models are more verbose, and the JSON strings that are generated sometimes cannot be parsed using a standard json parser. If this is the case, heuristics are

used to extract out time signatures embedded within the text, however, some time signatures may have been missed, negatively impacting performance.

Based on the descriptions provided, the action of "person putting on shoes" seems to have occurred in multiple instances, particularly from Frame 16 to Frame 28. However, the exact start and end frames are not explicitly stated. To provide a rough estimate, we can look for instances where the person is clearly seen bending over or interacting with their feet, which is a common action when putting on shoes. Here's a possible range for the start and end frames, considering the context and the actions described:

```
```json
{
 "start_frame": "16",
 "end_frame": "28"
}
```
```

Temporal activity localization results from Qwen 7B

Analyzing different input modes In Figure 2, we present prompting outputs for different types of visual input modes: a) using a single image, b) concatenating multiple consecutive images as a single image, and c) using a short 1-second video clip as input. We find that the multiple image and video modalities tend to give inconsistent descriptions, with the single image input mode providing the most detailed and accurate description. This may be due to the fact that multiple split images or 1-second clips are generally not as abundantly provided in the pretraining stage for these multimodal LLMs. Additionally, due to the fact that these models are relatively small (7B), it may even be a struggle to generate accurate descriptions for a single image, let alone multiple images or a video.

6 Conclusion

We demonstrate that instruction-tuning can be an effective way to generate elaborate and rich descriptions of scenes in smaller-sized multimodal LLMs. Our experiments show that a tuned multimodal LLM model can provide significant performance improvements in the task of temporal activity localization.

Our proposed two-stage approach, using only pretrained models, performs better than the state-of-the-art video-based LLMs, although it does not outperform state-of-the-art vision models specifically designed for temporal activity localization. The advantage over video-based LLMs is likely due to the higher accuracy and detail provided by image-based multimodal LLMs, as they are currently better trained and developed. Notably, the LLaVA 7B multimodal LLM performs well in the task of temporal grounding of activities, even exceeding the performance of Gemini 1.5 Pro when coupled with a strong text-based LLM.

While our current work focused on a single description of a scene, an obvious next step is to tap into the reasoning capabilities of the LLM. Our instruction-tuning dataset includes additional Q&A regarding actions within the scene, which, through further instruction-tuning, is expected to help the model gain additional insight into the actions occurring in the scene. Furthermore, exploring methods such as chain of thought (Lu et al., 2022) or tree of thought (Yao et al., 2024) prompting could be a valuable direction for tapping into the reasoning capabilities of the LLM and improving performance.

7 Key Information to include

- Mentor: Tathagat Verma

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736.
- Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. 2017. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. 2023a. Minigt-v2: large language model as a unified interface for vision-language multi-task learning.
- Wei-Ge Chen, Irina Spiridonova, Jianwei Yang, Jianfeng Gao, and Chunyuan Li. 2023b. Llava-interactive: An all-in-one demo for image chat, segmentation, generation and editing. *arXiv preprint arXiv:2311.00571*.
- Ran Cui, Tianwen Qian, Pai Peng, Elena Daskalaki, Jingjing Chen, Xiaowei Guo, Huyang Sun, and Yu-Gang Jiang. 2022. Video moment retrieval from text queries via single frame annotation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1033–1043.
- Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. 2017. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE international conference on computer vision*, pages 5267–5275.
- Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. 2023. Vtimellm: Empower llm to grasp video moments. *arXiv preprint arXiv:2311.18445*.
- Ang Li, Allan Jabri, Armand Joulin, and Laurens Van Der Maaten. 2017. Learning visual n-grams from web data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4183–4192.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*.
- KunChang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2023b. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*.
- Zhaowei Li, Qi Xu, Dong Zhang, Hang Song, Yiqing Cai, Qi Qi, Ran Zhou, Junting Pan, Zefeng Li, Van Tu Vu, Zhida Huang, and Tao Wang. 2024. Groundinggpt: language enhanced multi-modal grounding model.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. In *NeurIPS*.
- P. Lu, S. Mishra, T. Xia, L. Qiu, K.W. Chang, S.C. Zhu, O. Tafjord, P. Clark, and A. Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th Conference on Neural Information Processing Systems (NeurIPS)*.

- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. 2023. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*.
- OpenAI. 2023. Gpt-4 technical report.
- Long Qian, Juncheng Li, Yu Wu, Yaobo Ye, Hao Fei, Tat-Seng Chua, Yueting Zhuang, and Siliang Tang. 2024. Momentor: Advancing video large language model with fine-grained temporal reasoning.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy Lillicrap, Jean-baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Gunnar A. Sigurdsson, Gül Varol, Xiaolong Wang, Ivan Laptev, Ali Farhadi, and Abhinav Gupta. 2016. Hollywood in homes: Crowdsourcing data collection for activity understanding. *ArXiv e-prints*.
- H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample. 2023. Llama: Open and efficient foundation language models. <https://arxiv.org/abs/2302.13971>. ArXiv preprint arXiv:2302.13971.
- Jingwen Wang, Lin Ma, and Wenhao Jiang. 2020. Temporally grounding language queries in videos by contextual boundary-aware prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12168–12175.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2024. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems*, 36.
- Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. 2019. Semantic conditioned dynamic modulation for temporal sentence grounding in videos. *Advances in Neural Information Processing Systems*, 32.
- Hang Zhang, Xin Li, and Lidong Bing. 2023. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*.
- Renrui Zhang, Jiaming Han, Chris Liu, Aojun Zhou, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2024. LLaMA-adapter: Efficient fine-tuning of large language models with zero-initialized attention. In *The Twelfth International Conference on Learning Representations*.
- Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. 2020. Learning 2d temporal adjacent networks for moment localization with natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12870–12877.