GROUNDED-VIDEOLLM: SHARPENING FINE-GRAINED TEMPORAL GROUNDING IN VIDEO LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

Abstract

Video Large Language Models (Video-LLMs) have demonstrated remarkable capabilities in coarse-grained video understanding, however, they struggle with finegrained temporal grounding. In this paper, we introduce Grounded-VideoLLM, a novel Video-LLM adept at perceiving and reasoning over specific video moments in a fine-grained manner. We identify that current Video-LLMs have limitations for fine-grained video understanding since they lack effective temporal modeling and timestamp representation. In light of this, we sharpen our model by incorporating (1) an additional temporal stream to encode the relationships between frames and (2) discrete temporal tokens enriched with specific time knowledge to represent timestamps. To optimize the training of Grounded-VideoLLM, we employ a multi-stage training scheme, beginning with simple video-captioning tasks and progressively introducing video temporal grounding tasks of increasing complexity. To further enhance *Grounded-VideoLLM*'s temporal reasoning capability, we also curate a grounded VideoQA dataset by an automatic annotation pipeline. Extensive experiments demonstrate that *Grounded-VideoLLM* not only excels in fine-grained grounding tasks such as temporal sentence grounding, dense video captioning, and grounded VideoQA, but also shows great potential as a versatile video assistant for general video understanding.

029 030

006

008 009 010

011 012 013

014

015

016

017

018

019

021

023

025

026

027

028

031 032

033

1 INTRODUCTION

Multi-modal Large Language Models (MLLMs) have made remarkable progress in image-level understanding (Liu et al., 2023; Dai et al., 2023; Li et al., 2023a). However, extending their capabilities to the video domain poses distinct challenges. Unlike static images, the temporal nature of videos challenges models to process not only visual content but also the sequence and timing of events. While current Video-LLMs (Xu et al., 2024a; Li et al., 2024; Zhang et al., 2023b; Lin et al., 2023) are capable of capturing global visual semantics and generating coarse-grained captions for short clips, they struggle with fine-grained video understanding (Liu et al., 2024b; Wang et al., 2024d), which requires decomposing the video along the temporal axis to accurately perceive and reason over specific moments, such as subtle actions, transitions, and events that unfold over time.

Previous research efforts (Ren et al., 2024; Huang et al., 2024a; Qian et al., 2024a; Huang et al., 043 2024b; Guo et al., 2024) have explored *temporal grounding* to improve fine-grained video under-044 standing. However, two main challenges impede their potential for achieving effective *fine-grained* 045 temporal grounding: (1) Models like Video-ChatGPT (Maaz et al., 2024b), P-LLaVA (Xu et al., 046 2024a), and Video-LLAMA (Zhang et al., 2023b) typically sample multiple frames from a video 047 and encode each frame independently using an image encoder, followed by a feature projector 048 (e.g., sliding Q-former (Ren et al., 2024), slot-based token compression (Guo et al., 2024), or visual adapter (Huang et al., 2024a)). This approach focuses primarily on spatial details while potentially neglecting the temporal relationships between frames (Maaz et al., 2024a). (2) Cur-051 rent models also struggle with timestamp representation, which is crucial for pinpointing specific moments in time for fine-grained understanding. Models such as TimeChat (Ren et al., 2024) 052 and VTimeLLM (Huang et al., 2024a) represent timestamps as plain texts, for example, ["from 102.3 to 120.1 seconds"]. Despite being straightforward, this strategy needs to tokenize

060

061

062

063

064

065

067

068 069

071

073

074

<180> <276> <280> <12> <28> <34> <58> <60> <104> <254> <270> <288> Temporal Referring Temporal Localization Temporal Reasoning Why do the worker and assistants What is taking place between <12 When does 'the car left away approach the vehicle before it happen in the video? and <28>? leaves from <254> to <276 A large off-roading truck drives To help replace the tire as indicated ۲<mark>ش</mark> From <254> to <276>. by along a dirt track and comes between <58> and <180> to a stop. Can you compile a list of the Provide an overview of what What does the segment from <280> activities and their timestamps happens from <280> to <288> to <288> mean? featured in the video? A black screen appears with white words that read \"More Videos From <12> to <28>, a blue truck is coming and stops. From <34> to <60>, many 攟 It means the end of the video people approache the vehicle. From Coming, subscribe to my channel/

Figure 1: Grounded-VideoLLM enables Temporal Referring/Localizing/Reasoning for MLLM.

continuous floating-point values, which is inefficient for LLMs since their next-token prediction paradigm struggles with handling numerical data (Schwartz et al., 2024; Frieder et al., 2023).

To further improve video comprehension, we propose to sharpen the model with *fine-grained tem*-075 *poral grounding*, which allows the model to not only recognize *what* happens but also pinpoint *when* 076 it happens in a fine-grained manner. We identify three essential capabilities of *fine-grained temporal* 077 grounding, as illustrated in Figure 1, including Temporal Referring (Huang et al., 2024a), Localizing 078 (Gao et al., 2017; Caba Heilbron et al., 2015), and *Reasoning* (Xiao et al., 2024; Qian et al., 2024a). 079 *Referring* involves the model precisely describing events within user-specified time intervals. Lo*calizing* challenges the model to identify the timestamps of a given query and construct a coherent 081 storyline by locating a sequence of events within the video. *Reasoning* combines both *Referring* and 082 Localizing capabilities and further requires the ability to answer questions by understanding tem-083 poral dynamics and applying wrold knowledge. Achieving these tasks necessitates both effective temporal modeling and robust timestamp representation. 084

- 085 Targeting these goals, we introduce Grounded-VideoLLM, a novel Video-LLM that can perceive and reason over specific video moments with fine-grained precision. From the perspective of model ar-087 chitecture, *Grounded-VideoLLM* is built upon two key innovations: (1) **Two-Stream Encoding**: We decompose each segment of the video into spatial and temporal components, and encode each with an expert encoder respectively. The temporal stream extracts motion representations from dense frames and complements the spatial stream which captures appearance representations. This dual-090 stream approach forms comprehensive video representations enriched with both temporal and spatial 091 information. (2) Temporal Tokens: We extend the LLM's vocabulary by introducing discrete to-092 kens specifically crafted for timestamp representation. These temporal tokens denote relative time positions within a video and share a unified embedding space with the LLM. This integration al-094 lows Grounded-VideoLLM to avoid the inefficiency of tokenizing numerical text, enabling seamless prediction of both timestamps and textual outputs in a single sequence of discrete tokens. From the 096 perspective of training and datasets, given that such a unified model requires extensive pre-training, we start with an image-based MLLM (Microsoft, 2024) as the foundation and adopt a three-stage 098 training strategy to ensure efficiency. We meticulously select different tasks for each stage, and progressively refine the model in a "coarse-to-fine" manner, transitioning from image understanding to 099 video comprehension, and ultimately to fine-grained temporal grounding. Furthermore, we enhance 100 the model's temporal reasoning capability by curating 17K grounded VideoQA (Xiao et al., 2024) 101 samples with the assistance of GPT-4. Extensive experiments demonstrate that Grounded-VideoLLM 102 shows promising results over existing Video-LLMs not only in traditional video temporal grounding 103 tasks but also in general video understanding benchmarks. 104
- In summary, we make the following contributions: (1) We present key insights into fine-grained temporal grounding for Video-LLMs and introduce *Grounded-VideoLLM*, a Video-LLM adept at perceiving and reasoning over specific video moments. This is achieved through a two-stream architecture for effective temporal modeling and the temporal tokens for efficient timestamp repre-

sentation. (2) We propose a step-by-step training strategy that progressively adapts an image-based
 MLLM into a robust Video-LLM, while curating a grounded VideoQA dataset to further enhance the
 temporal reasoning capability. (3) We conduct extensive experiments across various tasks, including
 temporal sentence grounding, dense video caption, VideoQA, and general Video-LLM benchmarks,
 demonstrating the superiority of *Grounded-VideoLLM* for fine-grained video understanding.

114

2

113

115

RELATED WORK

116 Video Large Language Models have caught a growing interest with the advancements in image-117 based MLLMs (Zhang et al., 2023b; Lin et al., 2023; Maaz et al., 2024a; Luo et al., 2023). Despite 118 promising results, current Video-LLMs, such as P-LLaVA (Xu et al., 2024a) and Video-ChatGPT 119 (Maaz et al., 2024b), often struggle with temporal understanding (Liu et al., 2024b) and exhibit tem-120 poral hallucination (Wang et al., 2024d) when answering questions about specific video moments. These models encode each frame independently using a pre-trained image encoder, concatenating 121 the frame embeddings to create a video representation. This late fusion method can result in video 122 representations that lack inherent temporal information and heavily rely on the position embeddings 123 of LLM (Su et al., 2024) for temporal understanding, limiting the model's capability to perform fine-124 grained temporal grounding and understanding. In contrast to these studies, we employ a two-stream 125 architecture that integrates a video expert to extract motion features to complement the appearance 126 features during the early encoding process. Additionally, we employ a progressive training strategy 127 that gradually adapts an image-based MLLM for fine-grained video understanding. Unlike con-128 current studies such as SlowFast-LLaVA (Xu et al., 2024b) and VideoGPT+ (Maaz et al., 2024a), 129 which also utilize a two-stream architecture, we specifically targets fine-grained temporal grounding 130 through a unique encoding/pooling/training strategy for dense frames and grounding design.

131 Video Temporal Grounding (VTG) aims to associate specific video moments with their corre-132 sponding timestamps. Traditional VTG tasks include Temporal Sentence Grounding (Gao et al., 133 2017; Hendricks et al., 2018) and Dense Video Captioning (Caba Heilbron et al., 2015; Zhou et al., 134 2018). Other tasks, such as Grounded VideoQA (Xiao et al., 2024), emphasize reasoning over videos 135 with grounded information. Given the emerging reasoning capabilities of Video-LLMs, many stud-136 ies have investigated how to adapt them for VTG tasks. For example, TimeChat (Ren et al., 2024), 137 VTimeLLM (Huang et al., 2024a), and LITA (Huang et al., 2024b) perform temporal grounding us-138 ing a fully text-to-text approach through instruction-tuning datasets. Momentor (Qian et al., 2024a) introduces a temporal perception module to address the quantization errors associated with time to-139 kens, while VTG-LLM (Guo et al., 2024) incorporates a limited set of absolute-time tokens to han-140 dle timestamp knowledge. Compared to these studies, we avoid textual representation of timestamps 141 and instead introduce discrete and relative temporal tokens for more efficient timestamp encoding. 142 While Vid2Seq (Yang et al., 2023) also employs specialized tokens to indicate temporal positions, it 143 relies heavily on large-scale pre-training from scratch using noisy transcribed speech and is limited 144 to dense video captioning. Our two-stream architecture and progressive training strategy, however, 145 enable MLLMs to efficiently comprehend videos and effectively handle diverse tasks of fine-grained 146 temporal referring, localizing, and reasoning.

147 148

149

3 MODEL ARCHITECTURE

Given that current MLLMs already exhibit strong image-understanding capabilities, our architecture aims to sharpen temporal awareness by capturing motion dynamics across frames, which serve as a vital supplement to spatial content. As shown in Figure 2, we develop *Grounded-VideoLLM* upon a well-established MLLM for spatial comprehension and integrate an expert video encoder for temporal comprehension. Additionally, to avoid tokenizing numerical texts, we incorporate temporal tokens into the LLM's vocabulary for efficient and unified timestamp representation.

156 157 158

3.1 TWO-STREAM ENCODING

To effectively model the long-range temporal structure, given a video \mathcal{V} with T frames, we divide it into K segments and employ a segment-wise encoding strategy (Wang et al., 2016). Due to the inherent redundancy of consecutive frames in videos, each segment can be naturally represented from two perspectives: spatial and temporal. The spatial representation of each segment is derived

182

183

185

212



Figure 2: Overview of *Grounded-VideoLLM*. For temporal modeling, we employ a segment-wise encoding strategy by decomposing each segment into a spatial part and a temporal part and encoding each respectively. For timestamp representation, we introduce additional special temporal tokens sharing a unified embedding space with LLM.

from an individual keyframe, capturing the primary appearance semantics, while the temporal representation is learned from multiple frames depicting the motion evolution within the segment. We discuss the details of our two-stream segment encoding as follows.

Spatial Stream. We sample the middle frame from each segment as the keyframe and extract its spatial features using the original image encoder from the MLLM (Radford et al., 2021), resulting in spatial features $\mathbf{F}_S \in \mathbb{R}^{H_S \times W_S \times D_S}$, where H_S, W_S, D_S denote the height, width and dimension of the spatial features. Since dense frames are crucial for fine-grained temporal grounding, an appropriate pooling strategy is required to reduce token length for efficient computation. As indicated by Xu et al. (2024a) and Yao et al. (2024) that a 2D average pooling is both efficient and robust for spatial downsampling, we employ a 2D pooling kernel with a size $\sigma_S \times \sigma_S$ over the feature map and gets $\mathbf{F}_S \in \mathbb{R}^{N_S \times D_S}$ as the feature for spatial stream, where $N_S = \frac{H_S}{\sigma_S} \times \frac{W_S}{\sigma_S}$.

197 Temporal Stream. Traditional two-stream networks (Simonyan & Zisserman, 2014; Feichtenhofer et al., 2016) typically encode the optical flow as the temporal stream. However, given the scale of 199 data and parameters involved in MLLMs, extracting optical flow is computationally expensive and impractical. Consequently, we resort to a strong and well pre-trained video encoder, InternVideo2 200 (Wang et al., 2024b), to extract motion representations for each segment, using a lower resolution 201 but more frames. Specifically, we input each segment, containing $\frac{T}{K}$ frames, into the video encoder to obtain the segment-level features $\mathbf{F}_T \in \mathbb{R}^{\frac{T}{K} \times H_T \times W_T \times D_T}$, where H_T, W_T, D_T denote the height, 202 203 width and dimension of each frame feature. Similar to the spatial stream, we apply a 2D average 204 pooling strategy to downsample \mathbf{F}_T . However, as the temporal stream focuses primarily on temporal 205 modeling, we retain the complete temporal information by only pooling along the spatial dimen-206 sions. Specifically, we aggressively downsample \mathbf{F}_T using a kernel with a larger size of $\sigma_T \times \sigma_T$, 207 resulting in the compressed $\mathbf{F}_T \in \mathbb{R}^{\frac{T}{K} \times N_T \times D_T}$ for temporal stream, where $N_T = \frac{H_T}{\sigma_T} \times \frac{W_T}{\sigma_T}$. 208

To get the complete segment-level representation, we flatten the features of the spatial stream and temporal stream and concat them together :

$$\mathbf{F}_{Seq} = \text{Concat}\left[\text{Flatten}(f(\mathbf{F}_S)); \text{Flatten}(g(\mathbf{F}_T))\right], \mathbf{F}_{Seq} \in \mathbb{R}^{(N_S + \frac{T}{K} \cdot N_T) \times D}$$
(1)

where $f(\cdot)$ and $g(\cdot)$ are two MLPs that project the visual features to LLM's dimension D. The final video representation is formed by concatenating the K segment-level representations \mathbf{F}_{Seg} , resulting in $\mathbf{F}_{Vid} \in \mathbb{R}^{K \cdot (N_S + \frac{T}{K} \cdot N_T) \times D}$. This representation retains detailed spatial information across all 216 segments along with their global temporal contexts, while maintaining a manageable token length. 217 The combined video representation \mathbf{F}_{Vid} is then fed into the LLM serving as soft prompts, alongside 218 the word embeddings of the instruction text \mathbf{F}_{Text} to generate the target response \mathcal{A} . The model is 219 trained using the cross-entropy loss function with trainable parameters θ :

 $\mathcal{L} = -\sum_{t=1}^{L_a} log P_{\theta}(\mathcal{A}_t | \mathcal{A}_{< t}, \mathbf{F}_{Vid}, \mathbf{F}_{Text})$ ⁽²⁾

where A_t is predicted autoregressively at position t, and L_a is the sequence length of the ground truth answer text A.

3.2 UNIFIED TEMPORAL TOKENS

220

221

222

226 227

228

237

238 239

240

241

242 243 244

245

246 247 248

249

250

251

252

253

254 255 256

257

Given a text span depicting a particular video clip and its associated timestamps, we employ a 229 relative time representation that converts continuous timestamps into a sequence of discrete temporal 230 tokens. For a video \mathcal{V} with a duration of L seconds, we evenly divide \mathcal{V} into M equal-length chunks, 231 and then define M + 1 anchor points (ranging from <0> to <M>) across \mathcal{V} , representing relative 232 temporal positions. Each anchor point corresponds to a specific timestamp of \mathcal{V} and is encoded as a 233 temporal token. For instance, <0> denotes the very start of \mathcal{V} while <M> indicates the end. These 234 M + 1 tokens are added to the LLM's vocabulary to enable unified modeling alongside text. A 235 specific continuous timestamp τ can be easily converted to a temporal token $\langle t \rangle$ and vice verse: 236

$$t = \text{Round}(M \cdot \frac{\tau}{L}), \quad \tau = L \cdot \frac{t}{M}$$
 (3)

While this may introduce minor quantization errors, these can be minimized by selecting an appropriate M. We then organize the text span and its corresponding temporal tokens in a unified format. Both text tokens and temporal tokens are mapped to embeddings through the extended word embedding layer of LLM. For example, one input representation is as follows:

 $<s><video>F_{Vid}</video><grounded>From <0> to <6>, a baby is crying. From <7> to <16>, a man is coming and picking up the baby. From <20> to <25>, the baby is eating an apple. From <27> to <35>, the baby is smiling happily.</s>$

where <s> and </s> indicate start- and end-of-sequence, <video> and </video> represent the beginning and end of encoded video representations. <grounded> is a special token to tell the model should output the grounded timestamps. This strategy avoids the need to tokenize and process numerical values, which has been identified as a limitation of LLMs (Schwartz et al., 2024), and it greatly simplifies the representation of timestamps within the unified embedding space of LLMs. Consequently, text and timestamps can be jointly decoded as a single sequence, following the objective outlined in Eq. (2).

4 PROGRESSIVE TRAINING

Different from previous methods (Zhang et al., 2023b; Lin et al., 2023) that train models from scratch using mixed image and video datasets, we start with a pre-trained image-based MLLM (Microsoft, 2024) and progressively enhance its fine-grained temporal grounding capabilities. This strategy can be applied to any off-the-shelf MLLM and is more efficient. Table 1 enumerates the datasets used at different training stages, and Table 9 lists the prompts for different tasks.

Stage-1: Video-Caption Alignment. Feature alignment is widely used to improve training efficiency (Liu et al., 2023). In this stage, we leverage approximately 1.28 million video-text pairs sampled from diverse sources (Wang et al., 2024a; Bain et al., 2021; Chen et al., 2024b) to align the video encoder's features with the MLLM. This alignment allows the MLLM, which was pre-trained solely on images, to gain a foundational understanding of videos. Since 2D down-sampling has been applied to the visual features, only the two projectors $f(\cdot)$ and $g(\cdot)$ are set to be trainable, while the video encoder, image encoder, and LLM remain frozen. As this stage does not involve any video temporal grounding tasks, the temporal tokens described in Sec.3.2 are not yet incorporated.

| Training Stage | ng Stage Task | | Datasets |
|-------------------------------|-----------------------------|-------|---|
| Video-Caption Alignment | Video Captioning | 1.28M | WebVid-10M, Panda-70M, InternVid-10M |
| Temporal Token Alignment | Temporal Sentence Grounding | 149K | VTimeLLM-Stage2 |
| | Dense Video Captioning | 92K | VTimeLLM-Stage2, Moment-10M, InternVid-G |
| | Temporal Referring | 95K | VTimeLLM-Stage2, InternVid-G |
| | Grounded Conversation | 442K | ANet-RTL, Moment-10M |
| | Temporal Sentence Grounding | 84K | DiDeMo, HIREST, QuerYD, VTG-IT |
| | Dense Video Caption | 41K | COIN, ViTT, YouCook2, VTG-IT |
| | Grounded VideoQA | 17K | Self Collected |
| Multi-Task Instruction Tuning | Converstation | 233K | VCG-Plus-112K, Videochatgpt-100K, Videochat2-Conv |
| | VideoQA | 282K | EgoQA, NExT-QA, Intent-QA, STAR, CLEVRER, WebVid-QA |
| | Classification | 66K | SthSthV2, Kinetics |
| | Video Captioning | 136K | TextVR, YouCook2, WebVid, ShareGPT4Video |

Table 1: Datasets used at three training stages. Tasks with gray background consist of datasets regarding temporal grounding.

Stage-2: Temporal Token Alignment. While video-caption alignment effectively connects videos and the MLLM at a coarse semantic level, a gap persists between this alignment and fine-grained temporal grounding. To address this, we introduce the temporal tokens described in Sec.3.2 and continue pre-training the model on a diverse range of grounding datasets (Huang et al., 2024a; Qian et al., 2024a; Wang et al., 2024c), focusing on tasks such as Temporal Sentence Grounding, Dense Video Captioning, and Temporal Referring, which enables the model to refer to and localize temporal information effectively. Since new tokens are introduced, the trainable parameters in this stage include the two projectors, $f(\cdot)$ and $g(\cdot)$, the word embedding matrix, and the final classifier head of the LLM. This step enhances the model's ability to comprehend multiple events and aligns the temporal tokens with both the video timelines and the LLM's semantic space.

Stage-3: Multi-Task Instruction Tuning. Following the initial two stages of pre-training, the model has developed a basic understanding of video content and the ability to refer to and locate specific timestamps. In this stage, we will further enhance the model's fine-grained temporal ground-ing while improving its responsiveness to diverse user instructions. To achieve this, we gather two types of datasets: (1) We compile a wide range of public datasets for video temporal grounding tasks, similar to Time-IT (Ren et al., 2024) and VTG-IT (Guo et al., 2024), which include tasks of dense video captioning, temporal sentence grounding, and grounded VideoQA. (2) We incorporate a se-lection of instructional video-to-text datasets from VideoChat2 (Li et al., 2024), which feature tasks such as conversations, classification, question answering, and captioning. Additionally, we include ShareGPT-4Video (Chen et al., 2024a) to further enhance the model's ability to generate detailed video captions. By utilizing these diverse datasets, which encompass both temporal grounding and video instruction tasks, Grounded-VideoLLM excels in temporal referring, localization, reasoning, and general comprehension of video content. In this stage, the trainable parameters remain the same as in Stage 2, with the addition of LoRA parameters (Hu et al., 2022) for the LLM.

5 GROUNDED VIDEOQA DATASET GENERATION

Grounded VideoQA requires the model to not only answer questions but also identify relevant video moments that support predicted answers, thereby demonstrating the model's temporal reasoning abilities. The NExT-GQA dataset (Xiao et al., 2024) was manually developed by extending NExT-QA (Xiao et al., 2021) with temporal labels for start and end timestamps. However, annotating these temporal labels is labor-intensive and time-consuming, which has limited NExT-GQA only to QA pairs for the validation and test sets. To create a scalable training dataset, we utilized OpenAI GPT-4 (Achiam et al., 2023) to assist in constructing training sets for the grounded VideoQA task. These sets were built on public datasets that already contain temporal labels, such as ActivityNet-Caption (Caba Heilbron et al., 2015) and QVHighlights (Lei et al., 2021). We framed the task as a multiple-choice problem using a two-round conversational format, as depicted in Figure 3.

Specifically, we input event descriptions along with their timestamps into GPT-4 and prompted it to
 generate corresponding question-answer pairs, as shown in Table 8. To create distractor options for
 the multiple-choice questions, we retrieved the top 50 questions most similar to the generated ques tion, based on cosine similarity using an embedding model (Reimers, 2019). The answers to these
 for retrieved questions served as candidates for distractors. From this pool, we randomly sampled

339 340 341

342

343

344

345 346 347

348 349

350



Figure 3: Examples of annotation pipeline and generated data for Grounded VideoQA.

four distractors with cosine similarities to the correct answer ranging from 0.2 to 0.9, ensuring that the distractors were contextually relevant but not overly similar to the correct answer. The groundtruth timestamps for answering each question were derived from the timestamps of the associated event descriptions. The constructed dataset comprises 17K samples, which have been incorporated into the training sets for Stage 3, further enhancing the model's temporal reasoning performance.

EXPERIMENTS 6

6.1 EXPERIMENT SETTING

351 Implementation Details. We select the Phi3.5-Vision-Instruct-3.8B (Microsoft, 2024) as the base 352 MLLM for our Grounded-VideoLLM. For temporal stream encoding, we adopt InternVideo2-1B 353 (Wang et al., 2024b) as the video encoder. Each video is sampled as a sequence of T = 96 frames, which are evenly divided into K = 12 segments. We set the pooling size $\sigma_S = 2$ for the spatial 354 stream while $\sigma_T = 4$ for the temporal stream respectively, which results in $N_S = 144$ tokens per 355 frame for the spatial stream while $N_T = 16$ tokens per frame for the temporal stream. Moreover, 356 we introduce M = 300 temporal tokens into the LLM's vocabulary for timestamp representation. 357 More implementation details can be found in Appendix A.1. 358

Tasks and Benchmarks. To thoroughly evaluate Grounded-VideoLLM in fine-grained temporal 359 grounding, we assess it across three video temporal grounding tasks: Temporal Sentence Ground-360 ing, Dense Video Captioning, and Grounded VideoQA, utilizing datasets such as Charades-STA 361 (Gao et al., 2017), ActivityNet-Captions (Caba Heilbron et al., 2015), and NExT-GQA (Xiao et al., 362 2024). We also show its reasoning capability by the task of Open-Ended VideoQA with datasets including MSVD-QA, MSRVTT-QA (Xu et al., 2017), and ActicityNet-QA (Yu et al., 2019). Addi-364 tionally, to evaluate the model's general video understanding capabilities, we benchmark Grounded-365 VideoLLM against existing models using modern Video-LLM benchmarks including VCG-Bench 366 (Maaz et al., 2024b) and MVBench (Li et al., 2024).

367 **Evaluation Metrics.** For temporal sentence grounding, we report the metric of Intersection over 368 Union (IoU) (Gao et al., 2017) between the timestamps predicted by the model and the ground 369 truth, including **Recall at IoU** thresholds of $\{0.3, 0.5, 0.7\}$ and their **mean IoU**. For dense video 370 captioning, we use metrics including SODA₋c (Fujita et al., 2020) which is specifically tailored 371 for the video's storyline, and METEOR score (Banerjee & Lavie, 2005), which is the average of 372 traditional METEOR scores that are calculated based on matched pairs between generated events 373 and the ground truth across IoU thresholds of {0.3, 0.5, 0.7, 0.9}. For Visually-grounded VideoQA, 374 we calculate both the Intersection of Prediction (IoP) (Xiao et al., 2024) and Intersection of Union 375 (IoU), and use Acc@GQA (Xiao et al., 2024) to measure the percentage of questions that are both correctly answered and visually grounded with IoP ≥ 0.5 . The responses of Open-Ended VideoQA 376 and VCG-Bench are evaluated by GPT-3.5 with the prompts introduced by Video-ChatGPT (Maaz 377 et al., 2024b). More evaluation details are in Appendix A.3.

| Model | Model | | | Charade | | A | ctivityNet | ActivityN | et-Captions | | | |
|--------------------|---------------------|-------|-------------|---------|-------|-------------|------------|-------------|-------------|-------------|------------|--------|
| model | | Scale | R@0.3 | R@0.5 | R@0.7 | mIoU | R@0.3 | R@0.5 | R@0.7 | mIoU | SODA_c | METEOR |
| Video-LLaMA (Z | hang et al., 2023b) | 7B | 25.2 | 10.6 | 3.4 | 16.8 | 21.9 | 10.8 | 4.9 | 16.5 | 1.9 | 1.9 |
| SeViLA (Yu et al. | , 2023) | 3B | 27.0 | 15.0 | 5.8 | 18.3 | 31.6 | 19.0 | 10.1 | 23.0 | - | - |
| Video-ChatGPT (| Maaz et al., 2024b) | 7B | 27.2 | 6.2 | 1.9 | 19.7 | 19.5 | 10.6 | 4.8 | 14.2 | 1.9 | 2.1 |
| Valley (Luo et al. | , 2023) | 7B | 28.4 | 1.8 | 0.3 | 21.4 | 30.6 | 13.7 | 8.1 | 21.9 | 0.3 | 0.8 |
| VideoChat2 (Li e | t al., 2024) | 7B | 38.0 | 14.3 | 3.8 | 24.6 | 40.8 | 27.8 | 9.3 | 27.9 | - | - |
| VideoChat (Li et | al., 2023b) | 7B | 32.8 | 8.6 | 0.0 | 25.9 | 23.5 | 12.6 | 6.0 | 17.4 | 0.9 | 0.9 |
| Momenter (Qian | et al., 2024a) | 7B | 42.6 | 26.6 | 11.6 | 28.5 | 42.9 | 23.0 | 12.4 | 29.3 | 2.3 | 4.7 |
| VTimeLLM (Hua | ng et al., 2024a) | 7B | <u>51.0</u> | 27.5 | 11.4 | 31.2 | 44.0 | 27.8 | <u>14.3</u> | 30.4 | <u>5.8</u> | 6.8 |
| TimeChat (Ren et | al., 2024) | 7B | - | 32.2 | 13.4 | - | - | - | - | - | - | - |
| VTG-LLM (Guo | et al., 2024) | 7B | - | 33.8 | 15.7 | - | - | - | - | - | - | - |
| HawkEye (Wang | et al., 2024c) | 7B | 50.6 | 31.4 | 14.5 | <u>33.7</u> | 49.1 | <u>29.3</u> | 10.7 | <u>32.7</u> | - | - |
| Grounded-VideoL | LM | 4B | 54.2 | 36.4 | 19.7 | 36.8 | 46.2 | 30.3 | 19.0 | 36.1 | 6.0 | 6.8 |
| | | | | | | | | | | | | |

Table 2: Zero-shot results on temporal sentence grounding and dense video captioning tasks.

Table 3: Results on NExT-GQA. Acc@GQA is defined as the percentage of questions that are both correctly answered and visually grounded with IoP ≥ 0.5 .

| Model | Acc@GQA | mIoP | IoP@0.3 | IoP@0.5 | mIoU | IoU@0.3 | IoU@0.5 |
|--------------------------------------|-------------|-------------|---------|-------------|------|-------------|-------------|
| VIOLETv2 (Fu et al., 2023) | 12.8 | 23.6 | 25.1 | 23.3 | 3.1 | 4.3 | 1.3 |
| Temp[CLIP] NG+ (Xiao et al., 2024) | 16.0 | 25.7 | 31.4 | 25.5 | 12.1 | 17.5 | 8.9 |
| SeViLA (Yu et al., 2023) | 16.6 | 29.5 | 34.7 | 22.9 | 21.7 | 29.2 | 13.8 |
| HawkEye (Wang et al., 2024c) | - | - | - | - | 25.7 | 37.0 | 19.5 |
| LangRepo (Kahatapitiya et al., 2024) | 17.1 | 31.3 | - | 28.7 | 18.5 | - | 12.2 |
| FrozenBiLM NG+ (Yang et al., 2022) | 17.5 | 24.2 | 28.5 | 23.7 | 9.6 | 13.5 | 6.1 |
| VideoStreaming (Qian et al., 2024b) | 17.8 | 32.2 | - | 31.0 | 19.3 | - | 13.3 |
| LLoVi (Zhang et al., 2023a) | <u>24.3</u> | 37.3 | - | 36.9 | 20.0 | - | 15.3 |
| Grounded-VideoLLM | 26.7 | <u>34.5</u> | 42.6 | <u>34.4</u> | 21.1 | <u>30.2</u> | <u>18.0</u> |

401 402 403

404

378

379

390

391

392 393

397

399 400

6.2 MAIN RESULTS

405 **Temporal Sentence Grounding** requires the model to identify the precise time interval corresponding to a given query sentence. As shown in Table 2, Grounded-VideoLLM achieves performance 406 on "mIoU" with 36.8 for the Charades-STA (Gao et al., 2017) and 36.1 for ActivityNet-Grounding 407 (Caba Heilbron et al., 2015) respectively, surpassing previous SoTA end-to-end Video-LLMs, i.e., 408 HawkEye (Wang et al., 2024c), by a significant margin (+3.4). It is worth mentioned that the promis-409 ing performance of "mIoU" are largely attributed to the significant gains in terms of "R@0.7" com-410 pared with other thresholds, demonstrating that Grounded-VideoLLM is more advanced in localizing 411 specific moments within videos with finer granularity. 412

Dense Video Captioning involves generating descriptions for all events in a video, along with their corresponding start and end timestamps. We evaluated *Grounded-VideoLLM* on the ActivityNet-Captions (Caba Heilbron et al., 2015), and the results in Table 2 show that our method achieves the highest SODA_c score (6.0), which demonstrates that, thanks to the Temporal Token Alignment training stage, *Grounded-VideoLLM* is highly effective in identifying the multi-event structure of the video and capturing complete storylines. The highest METEOR score (6.8) also indicates that *Grounded-VideoLLM* provides more detailed event descriptions compared with other Video-LLMs.

NExT-GQA (Xiao et al., 2024) is quite challenging since it requires the model to not only correctly answer questions but also provide timestamps that support the answers, highlighting the temporal reasoning capability. According to Table 3, *Grounded-VideoLLM* achieves the highest Acc@GQA (26.7, +2.4) and delivers comparable IoU and IoP scores to models such as SeViLA (Yu et al., 2023) and LLoVi (Zhang et al., 2023a), which use specialized grounding modules or rely on proprietary large language models (Achiam et al., 2023). The highest Acc@GQA score further demonstrates *Grounded-VideoLLM*'s capability in both fine-grained temporal grounding and high-level reasoning.

Open-Ended VideoQA. As shown in Table 4, *Grounded-VideoLLM* achieves state-of-the-art or comparative performance across MSVD-QA, MSRVTT-QA (Xu et al., 2017), and ActivityNet-QA (Yu et al., 2019), highlighting its advancements in general video question answering.

General Video-LLM Benchmarks. While *Grounded-VideoLLM* excels in fine-grained temporal grounding, we aim to ensure that it maintains robust performance in general video understanding. Therefore, we conducted a comprehensive evaluation using VCG-Bench (Maaz et al., 2024b) and

449 450

458 459

471 472

473

| 432 | Table 4: Results on zero-shot Open-Ended VideoQA and VCG-Bench. VCG-Bench contains five |
|-----|---|
| 433 | aspects: Correctness of Information (CI), Detail Orientation (DO), Contextual Understanding (CU), |
| 434 | Temporal Understanding (TU), and Consistency (CO). |

| Model | MSV | MSVD-QA | | MSRVTT-QA | | ANet-QA | | | VCG | Bench | | |
|-----------------------------------|------------|---------|------|-----------|------|---------|------|------|------|-------|------|------|
| | Acc. | Score | Acc. | Score | Acc. | Score | CI | DO | CU | TU | CO | Avg |
| Video-LLMs w/o temporal groundin | g capabili | ty. | | | | | | | | | | |
| Video-LLaMA (Zhang et al., 2023b |) 51.6 | 2.5 | 29.6 | 1.8 | 12.4 | 1.1 | 1.96 | 2.18 | 2.16 | 1.82 | 1.79 | 1.98 |
| Video-ChatGPT (Maaz et al., 2024b |) 64.9 | 3.3 | 49.3 | 2.8 | 35.2 | 2.7 | 2.50 | 2.57 | 2.69 | 2.16 | 2.20 | 2.42 |
| Vista-LLaMA (Ma et al., 2024) | 65.3 | 3.6 | 60.5 | 3.3 | 48.3 | 3.3 | 2.44 | 2.64 | 3.18 | 2.26 | 2.31 | 2.57 |
| MovieChat (Song et al., 2024) | 75.2 | 3.8 | 52.7 | 2.6 | 45.7 | 3.4 | 2.76 | 2.93 | 3.01 | 2.24 | 2.42 | 2.67 |
| LongVLM (Weng et al., 2024) | 70.0 | 3.8 | 59.8 | 3.3 | 47.6 | 3.3 | 2.76 | 2.86 | 3.34 | 2.39 | 3.11 | 2.89 |
| VideoChat2 (Li et al., 2024) | 70.0 | 3.9 | 54.1 | 3.3 | 49.1 | 3.3 | 3.02 | 2.88 | 3.51 | 2.66 | 2.81 | 2.98 |
| Chat-UniVi (Jin et al., 2024) | 65.0 | 3.6 | 54.6 | 3.1 | 45.8 | 3.2 | 2.89 | 2.91 | 3.46 | 2.89 | 2.81 | 2.99 |
| P-LLaVA-7B (Xu et al., 2024a) | 76.6 | 4.1 | 62.0 | 3.5 | 56.3 | 3.5 | 3.21 | 2.86 | 3.62 | 2.33 | 2.93 | 3.12 |
| ST-LLM (Liu et al., 2024a) | 74.6 | 3.9 | 63.2 | 3.4 | 50.9 | 3.3 | 3.23 | 3.05 | 3.74 | 2.93 | 2.81 | 3.15 |
| VideoGPT+ (Maaz et al., 2024a) | - | - | - | - | - | - | 3.27 | 3.18 | 3.74 | 2.83 | 3.39 | 3.28 |
| Video-LLMs w/ temporal grounding | capability | V. | | | | | | | | | | |
| Momentor (Qian et al., 2024a) | 68.9 | 3.6 | 55.6 | 3.0 | 40.8 | 3.2 | - | - | - | - | - | - |
| VTimeLLM (Huang et al., 2024a) | - | - | - | - | - | - | 2.78 | 3.10 | 3.40 | 2.49 | 2.47 | 2.85 |
| LITA (Huang et al., 2024b) | - | - | - | - | - | - | 2.94 | 2.98 | 3.43 | 2.68 | 3.19 | 3.04 |
| Grounded-VideoLLM | 76.3 | 4.1 | 60.3 | 3.6 | 56.8 | 3.5 | 3.34 | 2.94 | 3.66 | 3.12 | 3.14 | 3.24 |

| TT 1 1 / | - D | 1. | | 1 1 | 1.1 | 1 . | . • | • |
|----------|--|--------|---------|---------|--------|--------|----------|-----------|
| Table | $\mathbf{v} \cdot \mathbf{R} \mathbf{e}$ | otilte | n N/l | V Rench | multi_ | choice | duestion | answering |
| Table . |). IX | Jourto | 011 111 | | munu- | CHOICE | question | answering |

| 451 | | | | | | | | | | | | | | | | | | | | | | |
|------|-------------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|------|-------------|------|-------------|------|------|-------------|------|------|------|-------------|-------------|------|-------------|-------------|
| 450 | Model | Avg. | AS | AP | AA | FA | UA | OE | OI | OS | MD | AL | ST | AC | MC | MA | SC | FP | CO | EN | ER | CI |
| 452 | VideoChatGPT (Maaz et al., 2024b) | 32.7 | 23.5 | 26.0 | 62.0 | 22.5 | 26.5 | 54.0 | 28.0 | 40.0 | 23.0 | 20.0 | 31.0 | 30.5 | 25.5 | 39.5 | 48.5 | 29.0 | 33.0 | 29.5 | 26.0 | 35.5 |
| 450 | VideoLLaMA (Zhang et al., 2023b) | 34.1 | 27.5 | 25.5 | 51.0 | 29.0 | 39.0 | 48.0 | 40.5 | 38.0 | 22.5 | 22.5 | 43.0 | 34.0 | 22.5 | 32.5 | 45.5 | 32.5 | 40.0 | 30.0 | 21.0 | 37.0 |
| 453 | VideoChat (Li et al., 2023b) | 35.5 | 33.5 | 26.5 | 56.0 | 33.5 | 40.5 | 53.0 | 40.5 | 30.0 | 25.5 | 27.0 | 48.5 | 35.0 | 20.5 | 42.5 | 46.0 | 26.5 | 41.0 | 23.5 | 23.5 | 36.0 |
| | TimeChat (Ren et al., 2024) | 38.5 | 40.5 | 36.0 | 61.0 | 32.5 | 53.0 | 53.5 | 41.5 | 29.0 | 19.5 | 26.5 | 66.5 | 34.0 | 20.0 | 43.5 | 42.0 | 36.5 | 36.0 | 29.0 | 35.0 | 35.0 |
| 454 | Video-LLaVA (Lin et al., 2023) | 43.0 | 46.0 | 42.5 | 56.5 | 39.0 | 53.5 | 53.0 | 48.0 | 41.0 | 29.0 | 31.5 | 82.5 | 45.0 | 26.0 | 53.0 | 41.5 | 33.5 | 41.5 | 27.5 | 38.5 | 31.5 |
| | P-LLaVA-7B (Xu et al., 2024a) | 46.6 | 58.0 | 49.0 | 55.5 | 41.0 | 61.0 | 56.0 | 61.0 | 36.0 | 23.5 | 26.0 | 82.0 | 39.5 | 42.0 | 52.0 | 45.0 | 42.0 | 53.5 | 30.5 | 48.0 | 31.0 |
| 455 | VideoChat2 (Li et al., 2024) | 51.1 | 66.0 | 47.5 | 83.5 | 49.5 | 60.0 | 58.0 | 71.5 | 42.5 | 23.0 | 23.0 | 88.5 | 39.0 | 42.0 | 58.5 | 44.0 | <u>49.0</u> | 36.5 | 35.0 | 40.5 | 65.5 |
| | ShareGPT4Video (Chen et al., 2024a) | 51.2 | 49.5 | 39.5 | 79.5 | 40.0 | 54.5 | 82.5 | 54.5 | 32.5 | 50.5 | 41.5 | 84.5 | 35.5 | 62.5 | 75.0 | 51.0 | 25.5 | 46.5 | 28.5 | 39.0 | 51.5 |
| 456 | ST-LLM (Liu et al., 2024a) | 54.9 | 66.0 | 53.5 | 84.0 | 44.0 | 58.5 | 80.5 | 73.5 | 38.5 | 42.5 | 31.0 | 86.5 | 36.5 | 56.5 | 78.5 | 43.0 | 44.5 | 46.5 | 34.5 | 41.5 | 58.5 |
| -100 | VideoGPT+ (Maaz et al., 2024a) | <u>58.7</u> | <u>69.0</u> | <u>60.0</u> | <u>83.0</u> | 48.5 | <u>66.5</u> | 85.5 | <u>75.5</u> | 36.0 | <u>44.0</u> | 34.0 | 89.5 | 39.5 | 71.0 | 90.5 | 45.0 | 53.0 | <u>50.0</u> | 29.5 | <u>44.0</u> | <u>60.0</u> |
| 457 | Grounded-VideoLLM | 59.4 | 76.0 | 75.5 | 77.0 | <u>48.0</u> | 67.5 | 85.5 | 77.0 | 34.5 | 39.5 | 59.5 | 86.5 | <u>44.5</u> | 60.5 | 79.0 | 51.5 | <u>49.0</u> | 46.0 | 35.0 | 42.5 | 54.0 |
| | | | | | | | | | | | | | | | | | | | | | | |

MVBench (Li et al., 2024). As shown in Table 4, Grounded-VideoLLM achieves promising results 460 in VCG-Bench, with an average score of 3.24, outperforming other Video-LLMs with temporal 461 grounding capabilities (e.g., LITA, VTimeLLM). Notably, Grounded-VideoLLM surpasses all other 462 Video-LLMs on the TU (Temporal Understanding) task, with a score of 3.12 (+7%), demonstrating 463 its superior temporal understanding, which can be attributed to the two-stream architecture that can 464 capture motion dynamics. For MVBench which provides 4,000 QA pairs spanning a wide range of 465 scenes categorized into 20 fine-grained tasks, the results, presented in Table 5, show that Grounded-466 VideoLLM achieves an average score of 59.4, surpassing other Video-LLMs. Notably, it achieves 467 top performance in several critical tasks requiring perceiving and reasoning over specific video moments, including Action Sequence (AS), Action Prediction (AP), Action Localization (AL), and 468 State Change (SC), with scores of 76.0 (+10%), 75.0 (+26%), 59.5 (+43%), and 51.5 (+14%), re-469 spectively, demonstrating significant advancements on fine-grained video understanding. 470

6.3 IN-DEPTH ANALYSIS

Two-stream Encoding. 474

We conduct ablations to 475 our two-stream encoding. 476 Specifically, we set two 477 variants by removing the 478 temporal stream while 479 only retaining the spa-480 tial stream, where all 481 frame embeddings are 482 concatenated as the video

Table 6: Impact of two-stream encoding and alignment stage.

| Model | C-STA | ANet-G | ANe | -Cap |
|---|---|---|---|--|
| | mIoU | mIoU | SODA_c | METEOR |
| Grounded-VideoLLM | 36.8 | 36.1 | 6.0 | 6.8 |
| w/o temporal-stream (sparse) w/o temporal-stream (dense) w/o temporal token alignment | $\begin{vmatrix} 30.4 (\downarrow 6.4) \\ 34.3 (\downarrow 2.5) \\ 27.5 (\downarrow 9.3) \end{vmatrix}$ | $\begin{array}{c} 28.0 (\downarrow 8.1) \\ 29.2 (\downarrow 6.9) \\ 23.1 (\downarrow 13.0) \end{array}$ | $\begin{array}{c} 4.9 \ (\downarrow 1.1) \\ 5.4 \ (\downarrow 0.6) \\ 4.7 \ (\downarrow 1.3) \end{array}$ | $5.5 (\downarrow 1.3) \\ 6.2 (\downarrow 0.6) \\ 6.4 (\downarrow 0.4)$ |
| | | | | |

483 representation: (1) w/o temporal-stream (dense) feeds T = 96 frames with a pooling size $\sigma_S = 4$ (36 tokens per frame), resulting in a total of $36 \times 96 = 3456$ tokens. (2) w/o temporal-stream 484 (sparse) feeds T = 24 frames with a pooling size $\sigma_S = 2$ (144 tokens per frame), also resulting 485 in a total of $144 \times 24 = 3456$ tokens. Both variants have a close number of tokens for video



Figure 4: We visualize the attention weights of the LLM when generating the temporal tokens.

498 representation compared to our two-stream encoding $(12 \times (144 + \frac{96}{12} \times 16)) = 3264$ tokens for 499 $\mathbf{F}_{Vid} \in \mathbb{R}^{K \cdot (N_S + \frac{T}{K} \cdot N_T) \times D}$). Table 6 shows that both variants result in a significant performance drop on temporal grounding tasks. Interestingly, the dense frame variant performs slightly better than the sparse frame variant, suggesting that dense temporal information is more critical for 502 grounding tasks than spatial details. Our two-stream architecture strikes a balance by effectively 503 capturing dense motion dynamics while maintaining essential appearance details.

504 **Temporal Token Alignment.** We further investigate the role of temporal tokens by ablation of the 505 second training stage, Temporal Token Alignment. Quantitative results in Table 6 reveal a sharp per-506 formance drop across all tasks, particularly in temporal sentence grounding tasks. Upon analyzing 507 the outputs, we found that the model often produces time intervals spanning nearly the entire video 508 (e.g., from <0> to <300>), neglecting the association between specific moments and temporal 509 tokens, which leads to misalignment. Qualitatively, we visualize the attention weights of the last layer in the LLM to demonstrate how temporal tokens attend to corresponding video moments. De-510 tails of the visualizations are provided in Appendix A.4. As shown in Figure 4 (a), when generating 511 the temporal token, e.g. <241> or <271>, the attention weights are higher and more focused on 512 their corresponding video moments. Conversely, in Figure 4 (b), when the model is trained without 513 the Temporal Token Alignment stage, the attention weights of temporal tokens become significantly 514 dispersed across irrelevant moments. This illustrates that our multi-stage training strategy is es-515 sential for achieving proper alignment between temporal tokens and the video timeline. We also 516 visualize the embedding distribution of temporal tokens in Appendix A.5. 517

Grounded VideoQA Dataset. We vali-518

date the role of our constructed grounded-519 VideoQA dataset by removing it from the 520 training sets of stage-3. Since the model 521 without training on our dataset usually 522 generates free-form texts when asked to 523 output the timestamps supporting the an-524 swer, we reformulate it as a temporal sen-525 tence grounding task, where we combine

| Model | NExT-GQA | | | | | | | | | |
|----------------------|--------------|---------------|--------------|--|--|--|--|--|--|--|
| | Acc@GQA | mIoP | mIoU | | | | | | | |
| Grounded-VideoLLM | 26.7 | 34.5 | 21.1 | | | | | | | |
| w/o grounded VideoQA | 18.1 (↓ 8.6) | 22.2 (↓ 12.3) | 12.9 (↓ 8.2) | | | | | | | |

526 the predicted answer and question into a single sentence and ask the model to localize its timestamps. Table 7 suggests that there is a significant performance decrease with regard to Acc@GQA 527 $(\downarrow 6.6)$, mIoP $(\downarrow 12.3)$, and mIoU $(\downarrow 8.2)$, from which we can conclude that our Grounded VideoQA 528 dataset is essential to further enhance the model's temporal reasoning capability. 529

530 531

532

496 497

500

501

7 CONCLUSION

533 We present Grounded-VideoLLM, a Video-LLM capable of fine-grained perception and reasoning 534 over specific video moments. This is achieved through a novel model architecture that incorporates two-stream encoding for effective temporal modeling, along with the temporal tokens for efficient 536 timestamp representation. We employ a multi-stage training scheme, starting with an image-based 537 MLLM and progressively equipping it with *fine-grained temporal grounding* capabilities. Additionally, we curate a grounded-VideoQA dataset to further enhance the model's temporal reasoning 538 ability. Extensive experiments demonstrate that Grounded-VideoLLM not only excels in video temporal grounding tasks but also performs strongly on general video understanding benchmarks.

540 REFERENCES

565

566

567

568

577

586

592

Hervé Abdi and Lynne J Williams. Principal component analysis. *Wiley interdisciplinary reviews: computational statistics*, 2(4):433–459, 2010.

- 544 Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, 546 Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dong-547 dong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang 548 Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit 549 Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin 550 Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, 551 Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, 552 Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong 553 Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro 554 Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-555 Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacroce, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, 558 Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua 559 Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Jilong Xue, Sonali Yadav, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, 561 Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan 562 Zhang, and Xiren Zhou. Phi-3 technical report: A highly capable language model locally on your 563 phone, 2024. URL https://arxiv.org/abs/2404.14219.
 - 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 preprint arXiv:2303.08774*, 2023.
- Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, pp. 1728–1738, 2021.
- Satanjeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pp. 65–72, 2005.
- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A
 large-scale video benchmark for human activity understanding. In *CVPR*, pp. 961–970, 2015.
- 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 preprint arXiv:2406.04325*, 2024a.
- Tsai-Shien Chen, Aliaksandr Siarohin, Willi Menapace, Ekaterina Deyneka, Hsiang-wei Chao,
 Byung Eun Jeon, Yuwei Fang, Hsin-Ying Lee, Jian Ren, Ming-Hsuan Yang, et al. Panda70m: Captioning 70m videos with multiple cross-modality teachers. In *CVPR*, pp. 13320–13331,
 2024b.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang,
 Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language
 models with instruction tuning. In *NeurIPS*, 2023.
- ⁵⁹⁰ Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network
 ⁵⁹¹ fusion for video action recognition. In *CVPR*, 2016.
- 593 Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Petersen, and Julius Berner. Mathematical capabilities of chatgpt. In *NeurIPS*, 2023.

| 594 595 596 | Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. An empirical study of end-to-end video-language transformers with masked visual modeling. In CVPR 2023 |
|-------------------|--|
| 597 | C VI R, 2025. |
| 598 | Soichiro Fujita, Tsutomu Hirao, Hidetaka Kamigaito, Manabu Okumura, and Masaaki Nagata. Soda: Story oriented dense video captioning evaluation framework. In <i>ECCV</i> , 2020. |
| 299 | |
| 600 601 | Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In <i>ICCV</i> , pp. 5267–5275, 2017. |
| 602 | |
| 603 | Yongxin Guo, Jingyu Liu, Mingda Li, Xiaoying Tang, Xi Chen, and Bo Zhao. Vtg-Ilm: Integrating |
| 604 605 | arXiv:2405.13382, 2024. |
| 606 607 | Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with temporal language. In <i>EMNLP</i> , 2018. |
| 608 609 | Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In <i>ICLR</i> , 2022. |
| 614 | |
| 612 | Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. Vtimellm: Empower llm to grasp video moments. In <i>CVPR</i> , 2024a. |
| 613 | De An Huene Chille Lice Subbachnes Dedhelmishner Henery Vir Deuls Malaham 71' 1' |
| 614 | De-An Huang, Snijia Liao, Subnashree Kadnakrishnan, Hongxu Yin, Pavio Molchanov, Zhiding |
| 615 | arVivi2403 10046 2024b |
| 616 | <i>urxiv</i> .2403.19040, 20240. |
| 617 | Peng Jin, Ryuichi Takanobu, Wancai Zhang, Xiaochun Cao, and Li Yuan. Chat-univi: Unified visual |
| 618 619 | representation empowers large language models with image and video understanding. In CVPR, 2024. |
| 620 | |
| 621 622 | Kumara Kahatapitiya, Kanchana Ranasinghe, Jongwoo Park, and Michael S Ryoo. Language repos- itory for long video understanding. <i>arXiv preprint arXiv:2403.14622</i> , 2024. |
| 623 624 | Jie Lei, Tamara L Berg, and Mohit Bansal. Detecting moments and highlights in videos via natural language queries. In <i>NeurIPS</i> , pp. 11846–11858, 2021. |
| 625 626 627 | 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 <i>ICML</i> , 2023a. |
| 628 629 630 | KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. <i>arXiv preprint arXiv:2305.06355</i> , 2023b. |
| 631 | |
| 632 633 | Runchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In |
| 63/ | <i>CVPR</i> , 2024. |
| 625 | Rin Lin Rin Zhu, Yang Ve, Munan Ning, Peng Jin, and Li Yuan, Video-Ilava: Learning united |
| 636 | visual representation by alignment before projection. <i>arXiv preprint arXiv:2311.10122</i> , 2023. |
| 637 | Haotian Liu, Chunyuan Li, Oingyang Wu, and Yong Jae Lee. Visual instruction tuning. In <i>NeurIPS</i> . |
| 638 639 | 2023. |
| 640 | Ruyang Liu, Chen Li, Haoran Tang, Yixiao Ge, Ying Shan, and Ge Li. St-llm: Large language |
| 641 | models are effective temporal learners. In ECCV, 2024a. |
| 042 | Yuanxin Liu, Shicheng Li, Yi Liu, Yuxiang Wang, Shuhuai Ren, Lei Li, Sishuo Chen, Xu Sun, |
| 043 | and Lu Hou. Tempcompass: Do video llms really understand videos? arXiv preprint |
| 644 | <i>arXiv:2403.00476</i> , 2024b. |
| 645 | Deim Lug Zimeng Zhao Min Verge Lug 'Dear De L' De la LUT W. L' 'M |
| 646 647 | Minghui Qiu, and Zhongyu Wei. Valley: Video assistant with large language model enhanced ability. <i>arXiv preprint arXiv:2306.07207</i> , 2023. |

662

683

684

691

692

- 648 Fan Ma, Xiaojie Jin, Heng Wang, Yuchen Xian, Jiashi Feng, and Yi Yang. Vista-Ilama: Reducing 649 hallucination in video language models via equal distance to visual tokens. In CVPR, 2024. 650
- Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Khan. Videogpt+: Integrating 651 image and video encoders for enhanced video understanding. arXiv preprint arXiv:2406.09418, 652 2024a. 653
- 654 Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: 655 Towards detailed video understanding via large vision and language models. In ACL, 2024b.
- Microsoft. microsoft/phi-3-vision-128k-instruct, 2024. URL https://huggingface.co/ 657 microsoft/Phi-3.5-vision-instruct. 658
- 659 Long Qian, Juncheng Li, Yu Wu, Yaobo Ye, Hao Fei, Tat-Seng Chua, Yueting Zhuang, and Siliang 660 Tang. Momentor: Advancing video large language model with fine-grained temporal reasoning. 661 In ICML, 2024a.
- Rui Qian, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Shuangrui Ding, Dahua Lin, and Jiaqi 663 Wang. Streaming long video understanding with large language models. arXiv preprint 664 arXiv:2405.16009, 2024b. 665
- 666 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, 667 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In ICML, 2021. 668
- 669 N Reimers. Sentence-bert: Sentence embeddings using siamese bert-networks. arXiv preprint 670 arXiv:1908.10084, 2019. 671
- Shuhuai Ren, Linli Yao, Shicheng Li, Xu Sun, and Lu Hou. Timechat: A time-sensitive multimodal 672 large language model for long video understanding. In CVPR, 2024. 673
- 674 Eli Schwartz, Leshem Choshen, Joseph Shtok, Sivan Doveh, Leonid Karlinsky, and Assaf Ar-675 belle. Numerologic: Number encoding for enhanced llms' numerical reasoning. arXiv preprint 676 arXiv:2404.00459, 2024. 677
- Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition 678 in videos. In NeurIPS, 2014. 679
- 680 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 682 for long video understanding. In CVPR, 2024.
- Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 685
- 686 Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. 687 Temporal segment networks: Towards good practices for deep action recognition. In ECCV, 688 2016.
- 689 Yi Wang, Yinan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinhao Li, Guo Chen, Xinyuan 690 Chen, Yaohui Wang, et al. Internvid: A large-scale video-text dataset for multimodal understanding and generation. In ICLR, 2024a.
- 693 Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yinan He, Guo Chen, Baoqi Pei, Rongkun Zheng, Jilan Xu, Zun Wang, et al. Internvideo2: Scaling video foundation models for multimodal video 694 understanding. arXiv preprint arXiv:2403.15377, 2024b.
- 696 Yueqian Wang, Xiaojun Meng, Jianxin Liang, Yuxuan Wang, Qun Liu, and Dongyan Zhao. Hawk-697 eye: Training video-text llms for grounding text in videos. arXiv preprint arXiv:2403.10228, 2024c. 699
- Yuxuan Wang, Yueqian Wang, Dongyan Zhao, Cihang Xie, and Zilong Zheng. Videohallucer: 700 Evaluating intrinsic and extrinsic hallucinations in large video-language models. arXiv preprint 701 arXiv:2406.16338, 2024d.

- 702 Yuetian Weng, Mingfei Han, Haoyu He, Xiaojun Chang, and Bohan Zhuang. Longvlm: Efficient 703 long video understanding via large language models. In ECCV, 2024. 704
- 705 Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of questionanswering to explaining temporal actions. In CVPR, 2021. 706
- 707 Junbin Xiao, Angela Yao, Yicong Li, and Tat-Seng Chua. Can i trust your answer? visually grounded 708 video question answering. In CVPR, 2024. 709
- Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. 710 Video question answering via gradually refined attention over appearance and motion. In ACM 711 *MM*, 2017. 712
- 713 Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. Pllava: 714 Parameter-free llava extension from images to videos for video dense captioning. arXiv preprint 715 arXiv:2404.16994, 2024a.
- 716 Mingze Xu, Mingfei Gao, Zhe Gan, Hong-You Chen, Zhengfeng Lai, Haiming Gang, Kai Kang, 717 and Afshin Dehghan. Slowfast-llava: A strong training-free baseline for video large language 718 models. arXiv preprint arXiv:2407.15841, 2024b. 719
- 720 Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Zero-shot video 721 question answering via frozen bidirectional language models. In NeurIPS, 2022.
- 722 Antoine Yang, Arsha Nagrani, Paul Hongsuck Seo, Antoine Miech, Jordi Pont-Tuset, Ivan Laptev, 723 Josef Sivic, and Cordelia Schmid. Vid2seq: Large-scale pretraining of a visual language model 724 for dense video captioning. In CVPR, pp. 10714–10726, 2023. 725
- 726 Linli Yao, Lei Li, Shuhuai Ren, Lean Wang, Yuanxin Liu, Xu Sun, and Lu Hou. Deco: Decoupling token compression from semantic abstraction in multimodal large language models. arXiv 727 preprint arXiv:2405.20985, 2024. 728
- 729 Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model 730 for video localization and question answering. In NeurIPS, 2023. 731
- Zhou Yu, Dejing Xu, Jun Yu, Ting Yu, Zhou Zhao, Yueting Zhuang, and Dacheng Tao. Activitynet-732 qa: A dataset for understanding complex web videos via question answering. In AAAI, 2019. 733
- 734 Ce Zhang, Taixi Lu, Md Mohaiminul Islam, Ziyang Wang, Shoubin Yu, Mohit Bansal, and Gedas 735 Bertasius. A simple llm framework for long-range video question-answering. arXiv preprint 736 arXiv:2312.17235, 2023a. 737
- Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. arXiv preprint arXiv:2306.02858, 2023b. 739
 - Luowei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from web instructional videos. In AAAI, volume 32, 2018.
- 741 742 743
- 744

740

- 745
- 746 747
- 748

749

- 750
- 751
- 752

754

14

756 A APPENDIX

758 A.1 MORE IMPLEMENTATION DETAILS

Phi3.5-Vison-Instruct (Microsoft, 2024) consists of a CLIP style ViT image encoder (Radford et al., 2021), an MLP projector $f(\cdot)$, and the large language model Phi3.5-mini-3.8B (Abdin et al., 2024). Each video is sampled as a sequence of T = 96 frames, which are evenly divided into K = 12segments. For the spatial stream encoded by the ViT in Phi3.5-V, we adopt a higher resolution 336×336 , but a lower resolution 224×224 for the temporal stream encoded by InternVideo-2.

We set the pooling size σ_S to be 2 while σ_T to be 4 respectively. For the spatial stream, each frame takes up $24 \times 24 = 576$ tokens before while $12 \times 12 = 144$ tokens after pooling. For the temporal stream, each frame takes up $16 \times 16 = 256$ tokens before while $4 \times 4 = 16$ tokens after pooling. Therefore, we have an overall of $K \times (144 + \frac{T}{K} \times 16) = 3264$ tokens in total.

⁷⁶⁹ During training, we use the AdamW optimizer with a cosine learning rate decay and set the learning rate as 2e-5 and 1e-3 for projector $f(\cdot)$ and $g(\cdot)$ in stage-1. During stage-2 and stage-3, we set the learning rate for both projectors and word embeddings as 2e-5, while 2e-4 for LoRA parameters (r= 128 and α = 256). All experiments are conducted on 8 NVIDIA A100/H800 GPUs.

773 774

775 A.2 INSTRUCTIONS FOR EACH TASK

The quality and diversity of instructions are essential in the training process. We manually write well-designed instructions for some tasks, combined with some templates in Time-IT (Ren et al., 2024). Table 9 lists the prompts for different tasks.

779

A.3 EVALUATION PROCESS

781 782

For the evaluation of the temporal sentence grounding task, we directly input the prompt ["At which time interval in the video can we see $\langle query \rangle$ occurring?"] in Table 9 to get the response ["From $\langle start \rangle$ to $\langle end \rangle$ "], and then calculate the predicted timestamps with the Equation (3) to get the IoU metrics.

For the evaluation of the dense video captioning task, we directly input the prompt ["Detect and report the start and end timestamps of activity events in

the video, along with descriptions."] in Table 9 to get the response ["From $< start_1 > to < end_1 >, < caption_1 >$. From $< start_2 > to < end_2 >,$ $< caption_2 >. ...$ "], and then calculate the SODA_c (Fujita et al., 2020) and Meteor scores (Banerjee & Lavie, 2005).

For the evaluation of the visually-grounded VideoQA task, we adopt a two-round conversation evaluation as follows:

Round-1:USER: Question: < question >. Options: < options >.ASSISTANT: Answer: < answer >.Round-2:USER: Provide the timestamps that correspond to your answer.ASSISTANT: From < start > to < end >.

801 802

796

797

798

799

800

In the first round, we input the question and options into the model and get the answer. In the second round, we input the question, options, and predicted answer as the contexts, together with the prompt ["Provide the timestamps that correspond to your answer."], into the model to get the predicted timestamps. With both the predicted answers and timestamps, we can calculate the metrics including IoU, IoP, and Acc@GQA (Xiao et al., 2024).

For the evaluation of the Open-ended VideoQA and VCG-Bench, we employed GPT-3.5 turbo to juxtapose model outputs with ground truth data as introduced by Video-ChatGPT (Maaz et al., 2024b), subsequently computing both accuracy and a score. To ensure a fair and consistent comparison with the results presented in Video-ChatGPT, we adopted the same prompt for our evaluations. For MVBench, we directly compute the accuracy of multiple-choice questions.

A.4 VISUALIZATION PROCESS

We visualize the attention weights from the last layer of the LLM during the generation of a new temporal token. Since the full video representation consists of a total of $K \times (N_S + \frac{T}{K} \times N_T)$ tokens—where T, K, N_S , and N_T denote the number of frames, number of segments, number of tokens per frame for the spatial stream, and number of tokens per frame for the temporal stream, respectively—we obtain an attention weight vector with the shape $[K \times (N_S + \frac{T}{K} \times \hat{N}_T), 1]$. First, we discard the spatial stream portion, focusing only on the temporal information, which results in a new vector with the shape $[K \times \frac{T}{K} \times N_T, 1]$. We then reshape this vector to the form $[T, N_T, 1]$ and average it along the spatial dimension, yielding [T, 1], which represents the final attention weights corresponding to each frame when generating a new token.

A.5 DISTRIBUTION OF TEMPORAL TOKENS

We visualize the embeddings of the M = 300 temporal tokens to investigate their distribution in embedding space. We employ PCA (Abdi & Williams, 2010) to reduce the dimensionality of the temporal tokens to 1D, 2D, and 3D representations. For all reductions, we use the reduced values as coordinates, incorporating a gradient color scheme in which the color of the data points changes progressively with the token index, as illustrated in Figure 5. Our observations reveal that temporal tokens with similar indices tend to cluster together, exhibiting a continuous transition from tokens with smaller indices (light colors) to those with larger indices (darker colors).



Figure 5: Visualization of temporal tokens with PCA.

| s used to generate visually-grounded VideoQA samples with GPT-4. |
|---|
| |
| n generator. I need your help in generating question-answer pairs pertaining |
| criptions. I have a video and I will provide you with descriptions of certain |
| esponding timestamps within this video. You need to consider these segments |
| I on the given description and timestamps, and select one segment which |
| a HIGH-QUALITY QUESTION. Based on the description of that segment, |
| to that segment, as well as one correct answer. Both the proposed answer |
| consistent with the content of the give description. BE CAREFUL! Your |
| answers should follow these rules: |
| segment spanning across the whole video. |
| alsed should include causal and temporal relationships as much as possible. |
| whing that is not covered in the given descriptions |
| NEVER appear in your question. |
| be a phrase no more than 7 words. Keep your answers concise and accurate. |
| 1 1 2 |
| |
| |
| |
| econds |
| 5] A young woman is seen standing in a room and leads into her dancing. |
| org the girl dances around the room while the camera captures lief move- |
| 421 She continues dancing around the room and ends by laying on the floor |
| |
| ent-3 |
| 56.26, 79.42] |
| girl do after she ended dancing? |
| r |
| |
| notration a) |
| 15U 4U0115) |
| nstrations) |
| 56.26, 79.42] girl do after she ended dancing? r |
| r |
| |
| netrations) |
| 1511 au(1115) |
| |

| 8 | |
|----|---|
| 9 | |
| 20 | |
| 1 | |
| 2 | |
| 3 | |
| 4 | Table 9: Prompts used for different tasks. |
| 25 | Dromata for Video Contion Alignment |
| 6 | (1) "Describe the following yideo concisely" |
| 7 | (1) Describe the following video concisely. , (2) "Provide a brief description of the given video clip." |
| } | (2) "Offer a succinct explanation of the footage presented" |
| | (4) "Summarize the visual content of the following video.". |
| | (5) "Give a short and clear explanation of the subsequent video clip.", |
| | (6) "Share a concise interpretation of the video provided.", |
| | (7) "Present a compact description of the clip's key features.", |
| | (8) "Relay a brief, clear account of the video shown.", |
| | (9) "Render a clear and concise summary of the video below.", |
| | (10) "Write a terse but informative summary of the following video clip." |
| | Prompts for Temporal Sentence Grounding: |
| | (1) "When does $\langle query \rangle$ happen in the video?", |
| | (2) "At what time does the occurrence $\langle query \rangle$ take place in the video?", |
| | (3) "During which part of the video does $\langle query \rangle$ occur?", |
| | (4) "When in the video does the $\langle query \rangle$ incident occur?", (5) "At which moment does $\langle query \rangle$ take place in the video?" |
| | (5) At which moment does $\langle query \rangle$ take place in the video?, |
| | (b) During which phase of the video does $\langle quer y \rangle$ happen? |
| | (8) "At what time does $\langle query \rangle$ occur in the video sequence?". |
| | (9) "When does the $\langle query \rangle$ situation take place in the video?", |
| | (10) "At which time interval in the video can we see $\langle query \rangle$ occurring?" |
| | |
| | Prompts for Dense Video Captioning: |
| | (1) "Localize a series of activity events in the video, output the start and end timestamp for each |
| | (2) "Detect and report the start and and timestamps of activity events in the video along with |
| | (2) Detect and report the start and end timestamps of activity events in the video, along with descriptions" |
| | (3) "Pinpoint the time intervals of activity events in the video, and provide descriptions for each |
| | event.", |
| | (4) "Can you compile a list of the activities and their timestamps featured in the video?". |
| | (5) "I need you to scrutinize the video and catalog every event it contains, along with the times- |
| | tamps." |
| | |
| | Prompts for Temporal Referring: |
| | (1) "What is happening from $< start > to < end >$?", |
| | (2) "What is taking place between $\langle start \rangle$ and $\langle end \rangle$?", |
| | (5) what events unfold between $\langle start \rangle$ and $\langle end \rangle$?", (4) "What is homomorphic during the pariod from $\langle start \rangle$ to $\langle scard \rangle$?" |
| | (4) what is happenning unning the period from $\langle start \rangle > 0 \langle ena \rangle$? |
| | (5) what is going on from $\leq start > to \leq end > ??$ |
| | (7) "How do things progress from $< start > to < end >$ " |
| | (8) "Can you describe what happens from $< start > to < end >?".$ |
| | (9) "Describe the events occurring between $\langle start \rangle$ and $\langle end \rangle$." |
| | (10) "Narrate the actions that unfold from $\langle start \rangle$ to $\langle end \rangle$." |
| | |
| | |
| | |
| | |
| | |