




A Survey on Video Temporal Grounding With Multimodal Large Language Model

Jianlong Wu , Member, IEEE, Wei Liu , Ye Liu , Meng Liu , Member, IEEE, Liqiang Nie , Senior Member, IEEE, Zhouchen Lin , Fellow, IEEE, and Chang Wen Chen , Life Fellow, IEEE

(Survey Paper)

Abstract—The recent advancement in video temporal grounding (VTG) has significantly enhanced fine-grained video understanding, primarily driven by multimodal large language models (MLLMs). With superior multimodal comprehension and reasoning abilities, VTG approaches based on MLLMs (VTG-MLLMs) are gradually surpassing traditional fine-tuned methods. They not only achieve competitive performance but also excel in generalization across zero-shot, multi-task, and multi-domain settings. Despite extensive surveys on general video-language understanding, comprehensive reviews specifically addressing VTG-MLLMs remain scarce. To fill this gap, this survey systematically examines current research on VTG-MLLMs through a three-dimensional taxonomy: 1) the functional roles of MLLMs, highlighting their architectural significance; 2) training paradigms, analyzing strategies for temporal reasoning and task adaptation; and 3) video feature processing techniques, which determine spatiotemporal representation effectiveness. We further discuss benchmark datasets, evaluation protocols, and summarize empirical findings. Finally,

we identify existing limitations and propose promising research directions.

Index Terms—Video-language understanding, video temporal grounding, fine-grained temporal understanding, vision-language model, large language model, multimodal learning.

I. INTRODUCTION

THE proliferation of untrimmed video content across domains such as surveillance, entertainment, and autonomous systems has created an urgent need for systems capable of precise temporal understanding. Real-world applications, including moment retrieval, scene editing, and temporal question answering, demand accurate identification of not only what events occur but precisely when they occur. Existing video-language models primarily focus on global or coarse-level video comprehension [1], [2], [3], making them inadequate for tasks requiring fine-grained temporal grounding of events described by natural language. To address this capability gap, *Video Temporal Grounding* (VTG) has emerged as a pivotal research area. VTG involves localizing video segments that correspond specifically to given textual queries, enabling detailed interaction with video content. The core challenge of VTG lies in precisely aligning complex linguistic semantics with temporally distributed visual information, while simultaneously handling complex temporal relationships within the video.

As illustrated in Fig. 2, VTG encompasses several closely related but distinct tasks: (a) *Video Moment Retrieval* [4], [5], where the goal is to identify video segments matching natural language descriptions; (b) *Dense Video Captioning* [6], [7], which requires generating temporally aligned captions for multiple events; (c) *Video Highlight Detection* [8], [9], aims at selecting segments most relevant to a given query; and (d) *Temporally Grounded Video Question Answering* [10], [11], which involves pinpointing the temporal evidence needed to accurately answer questions. Collectively, these tasks define the contemporary scope of VTG research and highlight the necessity of sophisticated temporal reasoning.

Although substantial progress has been achieved, early VTG methods based on traditional deep learning architectures [12], [13] continue to face significant limitations. These include challenges in bridging semantic gaps between visual and linguistic modalities, inadequate temporal context modeling, and limited generalization capabilities. Previous methods often relied

Received 29 June 2025; revised 9 September 2025; accepted 22 September 2025. Date of publication 29 September 2025; date of current version 9 January 2026. This work was supported in part by the National Natural Science Foundation of China under Grant 62376069, Grant 62276004, and Grant 62376140, in part by the Beijing Natural Science Foundation under Grant L257007, in part by the Young Elite Scientists Sponsorship Program by CAST under Grant 2023QNRC001, in part by Guangdong Basic and Applied Basic Research Foundation under Grant 2024A1515012027, in part by Shenzhen Science and Technology Program under Grant KQTD20240729102207002 and Grant ZDSYS20230626091203008, in part by Jiangsu Science and Technology Major Program under Grant BG2024041, and in part by the Special Fund for Taishan Scholar Project of Shandong Province. Recommended for acceptance by G. Kim. (Corresponding authors: Liqiang Nie; Meng Liu.)

Jianlong Wu is with the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen 518055, China, also with Pengcheng Laboratory, Shenzhen 518108, China, and also with Shenzhen Loop Area Institute, Shenzhen 518038, China (e-mail: wujianlong@hit.edu.cn).

Wei Liu and Liqiang Nie are with the School of Computer Science and Technology, Harbin Institute of Technology, Shenzhen 518055, China (e-mail: liuwei030224@gmail.com; nieliqiang@gmail.com).

Ye Liu and Chang Wen Chen are with the Department of Computing, The Hong Kong Polytechnic University, Hong Kong SAR 999077, China (e-mail: coco.ye.liu@connect.polyu.hk; changwen.chen@polyu.edu.hk).

Meng Liu is with the School of Computer Science and Technology, Shandong Jianzhu University, Jinan 250101, China (e-mail: mengliu.sdu@gmail.com).

Zhouchen Lin is with the State Key Lab of General AI, School of Intelligence Science and Technology, Peking University, Beijing 100871, China, also with the Institute for Artificial Intelligence, Peking University, Beijing 100871, China, and also with the Pazhou Laboratory (Huangpu), Guangzhou 510335, China.

For additional resources and details, readers are encouraged to visit our repository at <https://github.com/ki-lw/Awesome-MLLMs-for-Video-Temporal-Grounding>.

This article has supplementary downloadable material available at <https://doi.org/10.1109/TPAMI.2025.3615586>, provided by the authors.

Digital Object Identifier 10.1109/TPAMI.2025.3615586

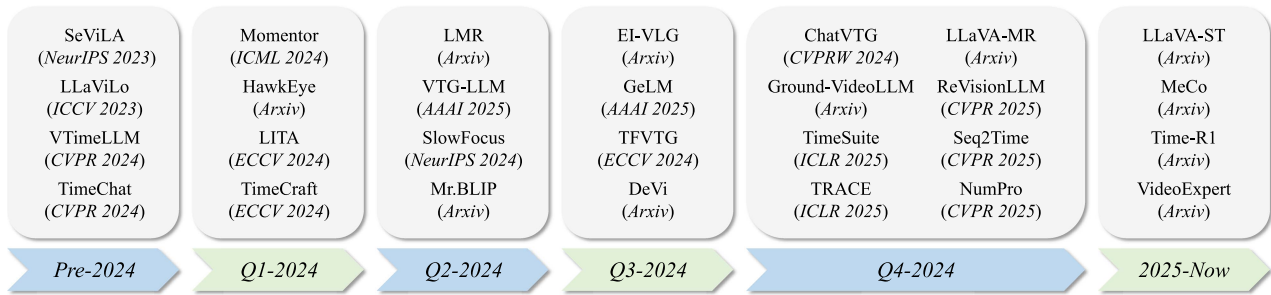


Fig. 1. An overview timeline of representative VTG-MLLMs. This timeline is organized according to the initial arXiv release dates of each work, with corresponding conference acceptance details included where applicable.

on manually designed proposal-generation mechanisms [14], [15] or simple temporal boundary regression [16], [17], which lacked flexibility and interpretability. Recently, the advent of Large Language Models (LLMs) [18], [19], [20] and their multimodal variants, i.e., Multimodal Large Language Models (MLLMs) [21], [22], [23], has dramatically reshaped the field of video-language understanding. These models provide powerful cross-modal reasoning, instruction-following capabilities, and robust zero-shot generalization, significantly enhancing the potential for effective and robust VTG.

Motivated by these advancements, a rapidly growing research direction termed *VTG-MLLMs* [24], [25], [26] has emerged. These new approaches address the limitations of their predecessors by fundamentally altering how the VTG task is framed. While often building upon powerful visual encoders [27], [28] from the pre-MLLM era, they diverge by replacing the paradigm of bespoke, task-specific modules with a unified, generalist reasoning engine. This represents a critical shift from *learning to localize* from scratch towards *reasoning to localize* based on vast pre-existing knowledge, a distinction that underpins the swift evolution of this subfield (depicted in Fig. 1). VTG-MLLMs adopt diverse architectural strategies, with some methods employing MLLMs as high-level facilitators for semantic grounding [29], [30], and others using them explicitly for boundary prediction [31], [32]. Consequently, the VTG-MLLM field now encompasses varied architectural innovations, training paradigms, and representation techniques [33], [34], [35]. However, the rapid evolution and complexity of VTG-MLLM research present challenges for navigating the current literature. Existing surveys predominantly focus on general video-language modeling [36], [37], [38], [39] or cover VTG from a pre-MLLM perspective [40], [41], [42], [43], [44], leaving a notable gap in systematic analyses of VTG in the MLLM era.

To fill this gap, we present the first comprehensive survey exclusively dedicated to VTG-MLLMs, covering literature up to May 2025. This survey systematically organizes recent developments, identifies emerging technical trends, and outlines future research opportunities. Specifically, we introduce a structured three-dimensional taxonomy:

- *Functional Roles of MLLMs*: Classifying models based on whether MLLMs act as *Facilitators* assisting downstream grounding tasks, or as *Executors* directly predicting temporal boundaries.
- *Training Paradigms*: Distinguishing between *pretraining*, *fine-tuning*, and *training-free* approaches, each with unique

trade-offs in generalization, task specialization, and supervision requirements.

- *Video Feature Processing Techniques*: Examining strategies for representing and integrating video inputs, including spatiotemporal tokenization and temporal modeling techniques.

Our taxonomy provides a progressive analytical framework, guiding readers from the high-level roles of MLLMs through training paradigms and down to video feature processing methods. By structuring our survey in this layered manner (see Fig. 3), we aim to offer clarity, advance comparative analyses, and identify underexplored avenues in VTG-MLLMs.

The remainder of this survey is organized as follows. Section II establishes the necessary background, covering an overview of VTG tasks, conventional pre-MLLM methods, and the foundational concepts of MLLMs. Section III presents a detailed taxonomy of recent VTG-MLLM research, categorizing methods along three key dimensions: the *functional roles* of MLLMs, *training paradigms*, and *video feature processing* strategies. These categories and their subtypes are illustrated in Fig. 3. Section IV provides an overview of benchmark datasets, evaluation protocols, and comparative analysis of empirical results across existing VTG-MLLMs. Section V discusses open challenges and future research directions. Finally, Section VI concludes the survey.

II. PRELIMINARIES


In this section, we provide an overview of core VTG tasks, review conventional pre-MLLM methodologies, and introduce the foundations of MLLMs.

A. Video Temporal Grounding

In this survey, we provide a comprehensive overview of four primary VTG tasks: video moment retrieval, dense video captioning, video highlight detection, and temporally grounded video question answering.

1) *Video Moment Retrieval*: Video moment retrieval (MR) [45], [46], also referred to as *temporal sentence grounding* [47], *video moment localization* [48], [49] or *temporal video grounding* [25], [50], aims to identify and localize temporal segments within untrimmed videos based on natural language queries (see Fig. 2(a)). This task serves as the most fundamental benchmark for assessing the temporal grounding capabilities of VTG-MLLMs, requiring both

Input: Locate the visual content mentioned in the text query “person tidy up the floor.” within the video.
Output: From 34.6 to 54.0 seconds.



(a) Video Moment Retrieval Task

Input: Determine the start and end times of various activity events in the video, accompanied by descriptions.

Output: From 62 to 65 seconds, cut the tofu into squares.

From 70 to 77 seconds, boil the tofu.

From 100 to 113 seconds, heat dashi soy sauce mirin and sake in a pan.


...



(b) Dense Video Captioning Task

Input: Find the highlight contents in the video corresponding to text query “A woman wearing glasses eating something at a street market.”, determining the highlight timestamps.

Output: The highlight timestamps are in the 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51, 52, 53, 54, 55, 56, 57, 58, 59, 60 seconds.



(c) Video Highlight Detection Task

Input: Why did the boy in green bend down on the sofa?

Output: Answer: Take closer look at baby. & Relevant video moment: From 24.5 to 27.9 seconds.



(d) Temporally Grounded Video Question Answering Task

Fig. 2. Illustration of four core tasks in VTG. These tasks encompass a range of temporal reasoning challenges, including precise segment localization, multi-event description generation, highlight identification, and time-sensitive question answering, each requiring fine-grained temporal understanding.

accurate alignment of textual descriptions with video segments and precise boundary localization.

2) *Dense Video Captioning*: Dense video captioning (DC) [51], [52] involves generating detailed, temporally grounded descriptions for all significant events or actions occurring in an untrimmed video, along with their corresponding start and end timestamps (see Fig. 2(b)). Unlike MR, which aims at localizing a single specific moment given a textual query, DC captures the complete narrative by identifying multiple events and their intricate temporal dependencies. This task assesses a model’s proficiency in comprehending extended temporal contexts and nuanced interactions within videos.

3) *Video Highlight Detection*: Video highlight detection (HD) [8], [53] aims to identify keyframes or short segments within an untrimmed video that best match a given textual query, typically by assigning an importance or relevance score to these moments (see Fig. 2(c)). Unlike MR and DC, which primarily operate at the event level, HD emphasizes frame-level precision. This task evaluates the ability of model to accurately pinpoint

salient video clips that closely correspond to textual prompts and to assess their contextual significance.

4) *Temporally Grounded Video Question Answering*: Temporally grounded video question answering (GQA) [50], [54], also known as *grounded video QA* [10], [55] or *temporal video grounding of questions* [32], extends video QA by requiring models to not only answer questions but also identify the precise temporal intervals containing relevant visual evidence (see Fig. 2(d)). GQA introduces the added complexity of integrating temporal localization with multimodal reasoning. This task is particularly critical for developing explainable video QA systems, as it demands explicit connections between textual answers and visual evidence.

B. Pre-MLLM VTG Methodologies

In the pre-MLLM era, VTG methodologies can be primarily categorized by their level of supervision: fully-supervised [45], [48] and weakly-supervised [14], [56], with the latter recently spawning a point-supervised sub-paradigm [57], [58]. Across

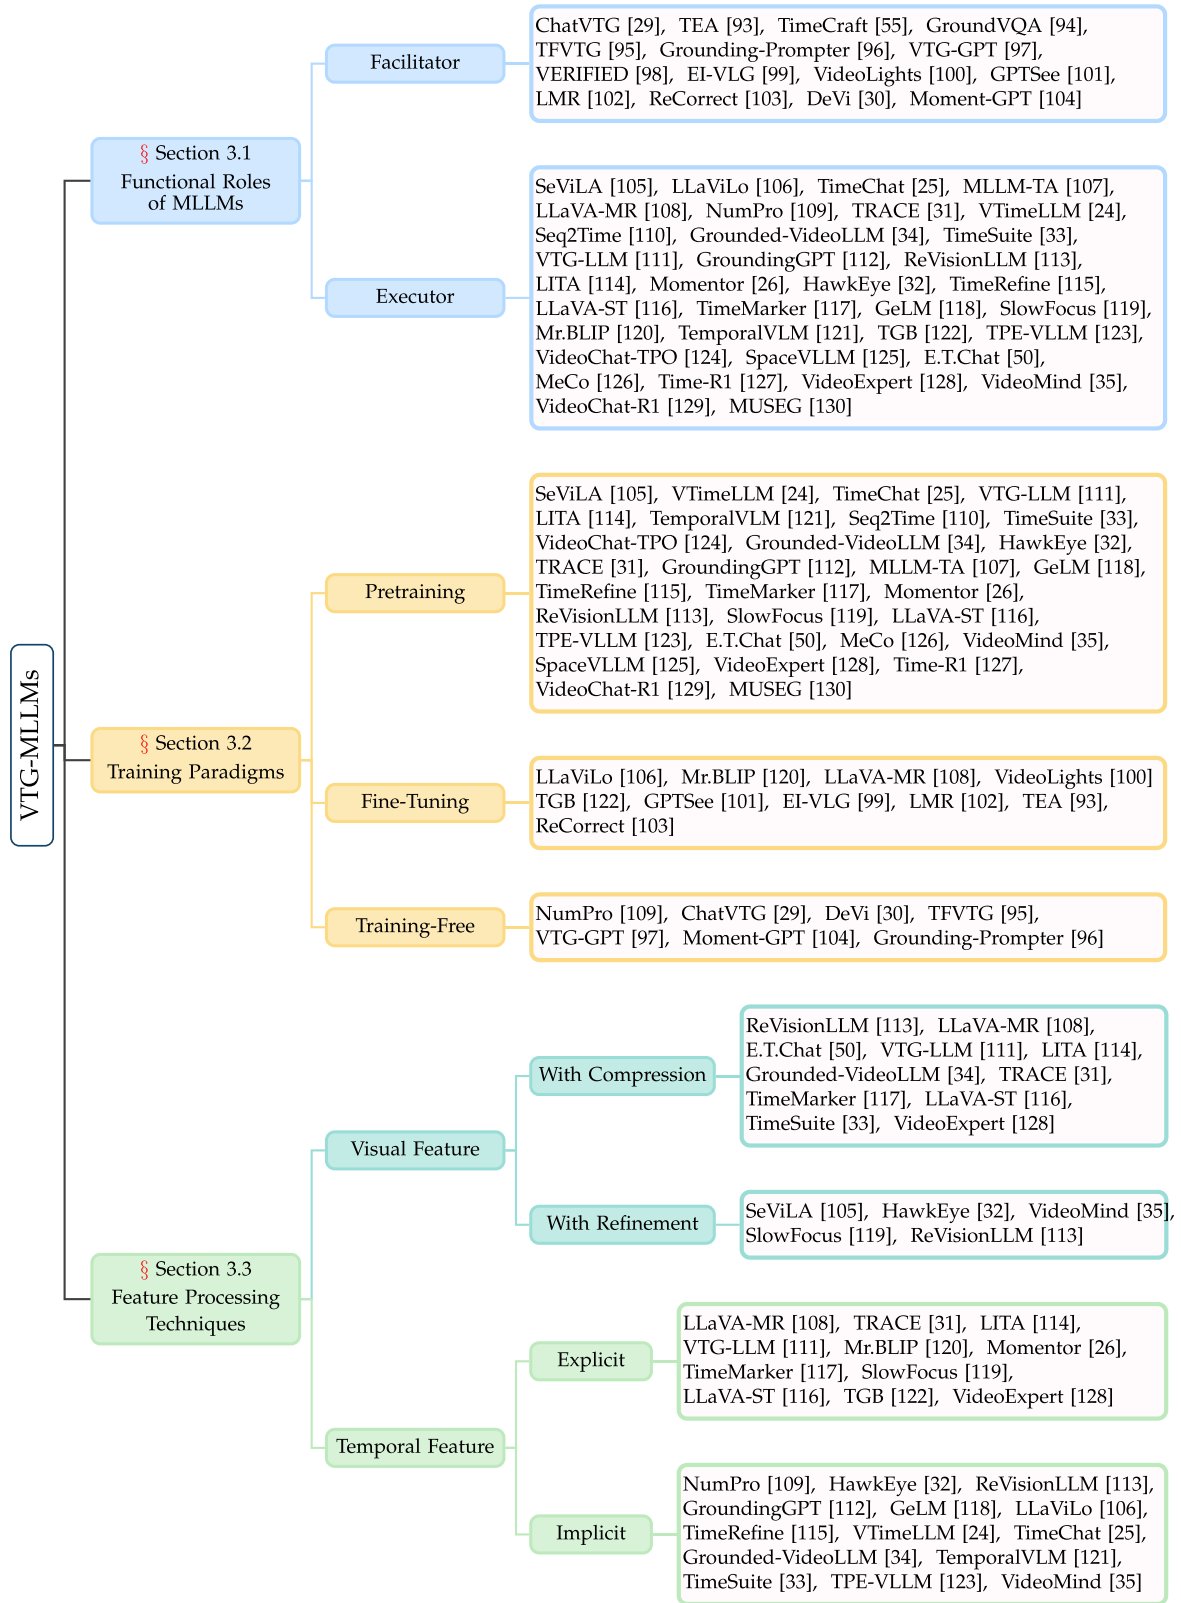


Fig. 3. Taxonomy of VTG-MLLMs, encompassing three primary dimensions: *Functional Roles of MLLMs* (Facilitators / Executors), *Training Paradigms* (Pretraining / Fine-Tuning / Training-Free), and *Feature Processing Techniques* (Visual Features / Temporal Features), each with distinct sub-categories reflecting the diverse strategies employed in this field.

these supervision schemes, architectural designs generally fall into three main categories: two-stage [48], [59], one-stage [12], [47], and reinforcement learning-based approaches [60]. A key differentiator among these architectures lies in their moment proposal strategy, which governs how candidate segments are generated and refined. Early methods generated moment candidates using sliding windows [14], matched by query similarity. This evolved into more efficient anchor-based [61], enumeration [62] techniques, and dedicated proposal networks [63] to handle diverse temporal scales. In contrast, proposal-free approaches relied on regression-based models [16], [64] to directly predict the temporal boundaries of the target moment. Despite achieving strong performance on specific benchmarks, these conventional approaches consistently struggled with poor generalization across diverse domains and complex queries. This fundamental limitation has catalyzed a significant shift in the field, paving the way for MLLMs to address these enduring challenges.

C. Multimodal Large Language Models

MLLMs [22], [65], [66] extend traditional LLMs by integrating multimodal encoders, such as image encoders [27], [28], [67], [68], video encoders [69], [70], [71], and specialized cross-modal adapters [72], [73]. Taking Video-LLM as an illustrative example, a video encoder processes a sequence of downsampled frames V , converting them into visual tokens $F^v = E_v(V)$. These visual tokens are then projected by a visual adapter to align with the embedding space of the language model, yielding aligned visual tokens $X^v = Q(F^v)$. Concurrently, an input textual query q , which may include instructions, prompts, or other textual elements, is encoded into linguistic tokens X^t via a textual encoder. The visual and textual tokens are concatenated into a unified input sequence $[X^v, X^t]$, which is subsequently processed by the LLM to generate the appropriate inference.

Current research efforts on MLLMs emphasize maximizing the efficiency of leveraging LLMs' advanced capabilities. Initial studies predominantly focus on designing cross-modality adapters aimed at mapping features from non-linguistic modalities into the semantic embedding space of language models. Flamingo [74], as a pioneering model, integrates visual and linguistic modalities through a gated cross-attention mechanism. Subsequently, various approaches, such as BLIP [75], mPLUG [76], and LanguageBind [77], adopt the Q-Former architecture to align visual representations, whereas models in the LLaVA series [21], [78], [79] introduce multilayer perceptrons (MLPs) as simpler yet effective connectors for modality integration. Additionally, more recent works have proposed lightweight and efficient alignment modules [80], [81], [82], continuing to enhance performance and model compactness.

Alongside architectural developments, training strategies constitute another critical research direction. Researchers have compiled large-scale multimodal pretraining datasets [83], [84], [85] to facilitate robust and diverse representation learning. Instruction-tuned datasets [86], [87] and specialized datasets tailored for chain-of-thought (CoT) reasoning [88], [89] have

been developed to improve task comprehension and generalization capabilities of MLLMs. Furthermore, parameter-efficient fine-tuning methods such as LoRA [90], LISA [91], and DoRA [92] have emerged, enabling efficient task-specific or domain-specific adaptation without extensive retraining.

III. A MULTI-FACETED TAXONOMY OF VTG-MLLMS

As established in the Introduction (Section I), we employ a three-dimensional taxonomy to deconstruct the complexities of VTG-MLLMs. This section delves into the specifics of this classification. Our taxonomy (visualized in Fig. 3), which progresses from high-level architectural considerations to fine-grained processing techniques, will be explored in detail through the following dimensions:

- *The Functional Roles of MLLMs (Section III-A)*: We will analyze how the architectural positioning of MLLMs—whether as *Facilitators* aiding downstream tasks or as *Executors* directly undertaking temporal prediction—shapes their overall design and impact on temporal perception.
- *The Training Paradigms (Section III-B)*: This subsection will differentiate among the *pretraining*, *fine-tuning*, and *training-free* paradigms. The analysis will center on the inherent trade-offs each strategy presents in terms of generalization capability, task-specific adaptation, and overall resource demands.
- *The Video Feature Processing Techniques (Section III-C)*: Here, we will systematically examine the diverse methodologies for representing and integrating video inputs. This includes a closer look at spatiotemporal tokenization mechanisms within the token budget, and various temporal modeling approaches that enable models to effectively process and reason about dynamic visual content.

This structured examination will provide the foundation for our detailed review of specific methodologies and trends within the VTG-MLLM field in the sections that follow.

A. Functional Roles of MLLMs in VTG-MLLMs

The functional role of MLLMs characterizes their architectural integration within VTG pipelines, determining whether they function primarily as auxiliary modules facilitating cross-modal understanding or as central reasoning engines directly conducting temporal grounding. Accordingly, existing VTG-MLLMs can be categorized into two paradigms: 1) *Facilitators*, where MLLMs generate structured textual representations from video content to support downstream modules; and 2) *Executors*, where MLLMs directly perform temporal boundary prediction via integrated multimodal reasoning.

1) *Facilitators*: In the Facilitator role, MLLMs act as intermediaries by transforming complex video data into structured textual forms, as depicted in Fig. 4(a). We formalize this process as a conditional generation problem:

$$T = \text{MLLM}_{\text{facilitator}}(V), \quad (1)$$

where V denotes the input video and T represents the generated textual descriptions. The generated textual outputs can either directly facilitate dataset construction or serve as semantic aids

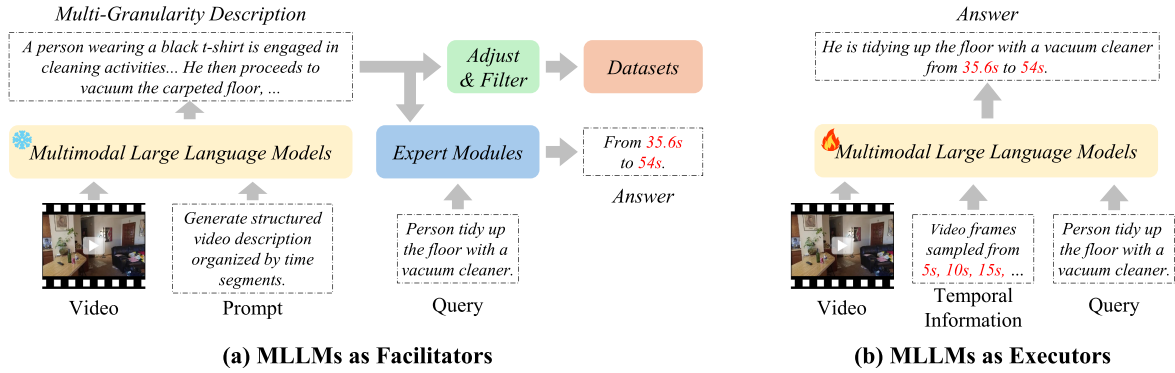


Fig. 4. Visualization of the distinct functional roles of MLLMs in VTG-MLLMs. The left panel depicts MLLMs as *Facilitators*, where they generate textual descriptions from video inputs, serving either as training data for downstream tasks or as auxiliary signals within expert modules. The right panel illustrates MLLMs as *Executors*, directly processing queries, video content, and temporal information (e.g., timestamps) to produce grounded outputs through a sequence-to-sequence prediction framework.

within dedicated downstream modules. Two primary application areas arise under this paradigm: *Dataset Construction* and *Expert Module Integration*.

Dataset Construction: MLLMs are extensively utilized to synthesize textual annotations, significantly enhancing the efficiency of dataset creation and expansion for model training and evaluation. For instance, Di and Xie [94] leverages Llama2 [131] to transform timestamped narrations from Ego4D [132] into temporally grounded QA pairs. Similarly, GPT-4o [133] was employed by Bao et al. [103] in Vid-Morp to automatically generate pseudo-labeled sentences aligned with video frames. Other works [98], [100] have similarly utilized advanced models like BLIP-2 [75], LLaVA [21], and Gemini-1.5 [134] to automate the annotation process and enrich datasets for VTG tasks.

Expert Module Integration: Beyond dataset generation, textual outputs from MLLMs can serve directly within VTG systems, either as semantic inputs in similarity-based grounding methods or as additional signals enhancing visual representations through cross-modal integration. For instance, Qu et al. [29] employs Video-ChatGPT [2] to generate multi-granularity clip captions, facilitating iterative query matching using Sentence-BERT [135]. Similarly, Xu et al. [97] utilizes MiniGPT-v2 [136] for caption generation and Baichuan2 [137] for query rewriting, reducing linguistic biases. Additionally, Cai et al. [93] leverages LLaVA-1.5 [78] to generate paragraph-level narrations, aligning them temporally with video features via cross-attention mechanisms [13], [16], thereby enhancing contextual understanding and robustness.

Summary: The Facilitator framework is advantageous due to its computational efficiency, ease of deployment, and inherent scalability, requiring minimal adaptation of pretrained MLLMs. However, reliance on static pretrained models carries limitations such as propagation of inherent biases [138] and constraints from original training data, potentially affecting the reliability of textual outputs and downstream performance. Furthermore, the fixed nature of off-the-shelf MLLMs inherently limits their capability for complex temporal reasoning, creating performance bottlenecks that sophisticated pipeline designs may not fully address. Nonetheless, as Facilitators, MLLMs remain valuable

for efficient dataset curation, providing abundant task-specific data crucial for advancing VTG research.

2) *Executors:* When functioning as Executors, MLLMs directly perform the core tasks of VTG, formulating the problem as an end-to-end sequence-to-sequence (seq-to-seq) prediction challenge. In this setting, illustrated in Fig. 4(b), the model jointly consumes raw video input and task-specific textual prompts to generate a temporally aligned output:

$$Y = \text{MLLM}_{\text{executor}}(V, Q, \tau), \quad (2)$$

where V denotes the input video stream, Q is the textual query, and τ optionally encapsulates temporal priors. The output Y represents predicted answers, timestamps, or task-specific tokens aligned with the video timeline.

This paradigm holds the potential to unify video understanding within a generative framework. However, it faces a significant hurdle: standard MLLMs [2], [66] often struggle to capture fine-grained temporal dependencies. This challenge largely stems from their vision encoders treating video as an unordered "bag-of-features," which discards the crucial sequential information essential for precise event localization. To address these challenges, research has focused on two primary directions: *Architectural Enhancement* and *Training Optimization*.

Architectural Enhancement: Architectural innovations are designed to improve the temporal perception and reasoning capabilities of MLLMs. These enhancements typically involve either modifying the input feature processing pipeline or integrating temporal awareness directly into the LLM's internal structure.

One line of research focuses on enhancing video feature representations before they are processed by the LLM. For instance, the method in Momentor [26] injects explicit temporal position encodings into frame-level features to improve temporal localization. Grounded-VideoLLM [34] adopts a dual-stream architecture to separately capture spatial and temporal dynamics, while LLaVA-MR [108] introduces components for reducing redundancy and emphasizing critical dynamic moments. These strategies aim to provide the LLM with richer temporal context, laying a stronger foundation for subsequent reasoning.

A complementary approach modifies the internal architecture or output mechanisms of LLMs to better handle temporal cues. For example, GeLM [118] incorporates flexible grounding tokens for temporal evidence retrieval, and TRACE [31] adds task-specific decoding heads for structured temporal output. VideoExpert [128] integrates parallel reasoning and generation modules for specialized processing. Similarly, VideoMind [35] decomposes complex tasks into specialized roles, like a Grounder with a timestamp-decoder and a Planner to coordinate them, using a Chain-of-LoRA strategy for seamless collaboration. While these approaches can significantly enhance temporal understanding, they may introduce trade-offs, such as increased computational overhead or reduced general-purpose flexibility.

Training Optimization: Effective optimization strategies are critical for equipping MLLMs with robust temporal understanding, transforming them into capable Executors. These strategies typically form a holistic pipeline combining novel training curricula with temporal tasks and datasets.

A prevalent approach involves multi-stage training frameworks, as seen in VTimeLLM [24] and SlowFocus [119], which progressively refine the model’s temporal localization abilities. Another emerging direction is optimization via reinforcement learning (RL). Time-R1 [127] adapts RL-based strategies specifically for temporal reasoning, and VideoChat-R1 [129] explores the effectiveness of Group Relative Policy Optimization (GRPO) [139].

To further strengthen temporal perception, many methods incorporate explicit reasoning tasks. For example, TPE-VLLM [123] introduces novel training objectives targeting boundary detection and duration reasoning, improving its handling of complex temporal relationships.

Underpinning all these strategies is the reliance on high-quality, time-annotated datasets. For instance, the TimeIT dataset from TimeChat [25] provides rich timestamp annotations essential for instruction tuning, while TimeSuite [33] offers a unified collection of diverse datasets to facilitate more comprehensive temporal learning.

Summary: The Executor paradigm represents a pivotal shift towards unified, end-to-end temporal grounding, allowing MLLMs to simultaneously process video content and textual queries within a tightly coupled seq-to-seq framework. This approach supports flexible input and output formats, capturing complex visual-textual correlations without relying on modular, cascaded architectures. However, this flexibility comes at a cost, often requiring extensive annotated datasets, significant computational resources, and complex training procedures. Despite these challenges, the Executor approach remains a promising direction for advancing fine-grained video understanding, with the potential to fundamentally reshape the field by integrating deeper temporal reasoning into multimodal models.

B. Training Paradigms of VTG-MLLMs

Building on the functional differentiation of MLLMs, this subsection examines the training paradigms used to adapt these models for effective video temporal grounding. The choice of

training approach reflects not only the system’s design goals, i.e., whether to build a domain-generalist or task-specific model, but also the trade-offs in supervision, resource efficiency, and scalability. We categorize current VTG-MLLM approaches into three main paradigms: *Pretraining*, *Fine-Tuning*, and *Training-Free* pipelines.

1) *Pretraining VTG-MLLMs:* Pretraining in VTG-MLLMs aims to equip models with robust temporal reasoning capabilities via large-scale supervised learning. Fundamentally, like most generative multimodal approaches, the pretraining process involves training the model to generate a target output T , typically encompassing temporal annotations such as event order, timestamps, or durations, conditioned on an input video V and textual prompt P . Formally, this objective is represented as minimizing the pretraining loss over a dataset \mathcal{D} :

$$\mathcal{L}_{\text{pretrain}} = \sum_{(V,T) \in \mathcal{D}} \mathcal{L}_{\text{gen}}(T | V, P). \quad (3)$$

However, unlike general-purpose MLLMs, the primary innovation in pretraining-based VTG-MLLMs lies not in architectural design but in the creation of sophisticated training strategies and specialized pretraining datasets that tailor the optimization of (3) for complex temporal understanding. We focus on two critical aspects of this approach: *Prevalent Pretraining Strategies* and *High-Quality Temporal Datasets*.

Prevalent Pretraining Strategies: The cornerstone of VTG pretraining is the multi-stage, progressive supervised learning pipeline. As introduced in the context of Executor models (Section III-A2), this strategy is founded on the principle of incremental learning, guiding the model from coarse-grained understanding to fine-grained localization. For instance, VTimeLLM [24] exemplifies this with its boundary-aware three-stage process, which sequentially tackles feature alignment, instruction tuning, and precise boundary optimization. Similarly, SlowFocus [119] integrates mixed-frequency sampling in its final training stages to enhance temporal resolution. This multi-stage philosophy has become a de facto standard, with numerous other models like TimeMarker [117], GroundingGPT [112], and LLaVA-ST [116] adopting similar hierarchical frameworks to progressively refine temporal perception.

While multi-stage supervision remains the mainstream, a new wave of research is exploring innovative refinements, aiming for greater precision and efficiency. A particularly prominent direction is the application of RL. Inspired by the success of RL and techniques like GRPO [139] in complex reasoning domains such as code generation and mathematics, researchers have begun adapting these methods for VTG. RL allows for the direct optimization of task-specific metrics like IoU, achieved by designing a composite reward function that encourages both a structured reasoning process and high prediction accuracy. Pioneering this direction, Time-R1 [127] introduces a reasoning-guided framework with a novel reward mechanism; VideoChat-R1 [129] further offers a systematic exploration with GRPO, while MUSEG [130] addresses the single-segment limitation by enabling reasoning over multiple distributed events.

Beyond these, other novel strategies are also emerging. For example, Seq2Time [110] adopts a data-centric strategy, synthesizing sequential training data with self-generated temporal cues, while TimeRefine [115] reformulates temporal grounding as an iterative refinement task, allowing model to self-improve its localization accuracy.

High-Quality Temporal Datasets: High-quality, temporally annotated multimodal datasets are crucial for pretraining and instruction-tuning VTG models, providing the diverse contexts required for robust generalization. Building upon early efforts like TimeIT [25] and VTimeLLM [24], subsequent work has evolved along several key directions. One major line of work aims to enhance data scale and diversity for more effective pretraining. For instance, VTG-IT-120K [111] expands on TimeIT by incorporating annotations from YT-Temporal-180M [140], InternVid-G [32] enriches InternVid10M-FLT [141] with segment-level captions and hard negative samples for more precise grounding, and Vid-Morp [103] leverages pseudo-labeling on real-world videos to scale data creation. A parallel direction develops specialized instruction-tuning datasets to align models with complex temporal reasoning. Moment-10M [26], sampled from YT-Temporal-1B [142], was designed for this purpose, while E.T. Instruct 164K [50] complements this by providing a curated dataset across nine distinct tasks, specifically tailored for multi-event and time-sensitive scenarios. More recently, frontier datasets have begun to broaden the scope of VTG by integrating spatial dimensions [125] or advancing spatiotemporal understanding [116]. Collectively, these datasets, along with numerous other contributions [33], [34], [127], underscore a clear trajectory toward building more comprehensive and fine-grained data resources to tackle the full spectrum of temporal reasoning challenges.

Summary: The pretraining paradigm empowers MLLMs with robust temporal grounding capabilities, supporting generalization across diverse downstream VTG tasks. However, this approach also presents significant challenges, including the high computational cost of training and the substantial effort required to construct large, high-quality temporal datasets. Effective pretraining strategies must carefully balance task complexity and learning progression to maximize temporal understanding, making this a critical area for ongoing research.

2) *Fine-Tuning VTG-MLLMs:* In contrast to the resource-intensive pretraining paradigm (Section III-B1), fine-tuning VTG-MLLMs provides a more computationally efficient approach, requiring smaller, task-specific datasets. Research within this paradigm can be broadly divided into two main directions: *Direct Fine-Tuning of MLLMs* and *Offline Textualization with MLLMs*, closely aligning with the functional roles discussed in Section III-A.

Direct Fine-Tuning of MLLMs: This approach directly fine-tunes general-purpose pre-trained MLLMs while maintaining their original architecture, adapting them for VTG tasks through task-specific training objectives. These methods typically reframe VTG as a seq-to-seq prediction problem, leveraging the contextual understanding capabilities already embedded in the models.

For instance, SeViLA [105] adapts BLIP-2 [75] into two interconnected components—a localizer and an answerer—strategically linking the outputs of the localizer to guide the answerer, thereby enhancing temporal precision. Similarly, LLaViLo [106] incorporates lightweight adapters to integrate video-text features, utilizing a multi-objective loss function for more refined temporal grounding. TGB [122] introduces a unique approach by leveraging CNN-extracted optical flow features as low-dimensional motion cues, enhancing temporal awareness without substantially increasing model complexity. Additionally, recent innovations have focused on optimizing BLIP-2 for fine-grained temporal understanding. Notably, Mr.BLIP [120] explores novel multimodal input sequences to improve temporal understanding of events, while LLaVA-MR [108] introduces a dynamic token compression strategy, reducing redundancy in spatiotemporal features and capturing more fine-grained event cues.

Despite their efficiency, these methods face the *catastrophic forgetting* challenge [143]. As the models adapt to fine-tuning data, they often lose the general-purpose capabilities acquired during pretraining, leading to performance degradation on broader video understanding tasks.

Offline Textualization with MLLMs: An alternative fine-tuning strategy employs MLLMs in a static capacity to convert raw video inputs into textual descriptions, which then guide downstream modules. This approach effectively bridges the gap between unstructured visual data and language-conditioned learning tasks, often integrating components from traditional VTG methods.

For example, GPTSee [101] generates detailed video descriptions, subsequently matched with textual queries to support moment localization. EI-VLG [99] incorporates these descriptions as environmental cues in a contrastive learning framework, refining the temporal precision of candidate segments. To address redundancy, LMR [102] uses cross-attention to highlight query-relevant segments, improving contextual alignment. Similarly, TEA [93] integrates these textual outputs with visual features to enhance semantic discriminability and temporal precision.

While this approach can significantly improve temporal grounding accuracy, it inherits certain limitations from traditional VTG methods. For instance, early two-stage matching approaches [59], [144] often rely on pre-defined segment boundaries, limiting global context modeling. Meanwhile, direct regression methods [64], [145] can misinterpret visually similar segments as semantically identical due to attention biases. In contrast, fine-tuned MLLMs provide richer, more context-aware embeddings, aligning visually similar events with distinct textual semantics and improving overall robustness.

Summary: Fine-tuning VTG-MLLMs offers a practical compromise between pretraining and fully training-free approaches, significantly reducing computational overhead while enhancing task-specific temporal alignment. However, these methods are inherently task-optimized, limiting their generalization across broader video understanding domains. As a result, fine-tuned models often excel in narrow, well-defined tasks but struggle with broader generalization—a critical challenge for future research.

3) *Training-Free VTG-MLLMs*: Training-free approaches represent a rapidly emerging paradigm in VTG-MLLMs, notable for their low computational overhead and zero-shot nature that eliminates the need for labeled supervision [146]. These methods bypass the need for end-to-end training by leveraging pre-trained foundation models (e.g., MLLMs and LLMs) and specialized expert tools [75], [135], [147], enabling temporal grounding through purely inference-based pipelines. While they share architectural similarities with fine-tuning approaches (Section III-B2), training-free methods distinguish themselves by substituting trainable components with off-the-shelf models, significantly reducing the need for task-specific parameter updates. Current training-free VTG-MLLMs generally adopt one of two principal strategies: *Feature Similarity Matching* and *LLM-Driven Reasoning*, distinguished by their approach to leveraging MLLM-generated textualizations for temporal localization.

Feature Similarity Matching: This strategy relies on extracting high-dimensional semantic representations from both natural language queries and textualized video content using pre-trained encoders. Temporal grounding is then achieved by identifying the video span s^* that maximizes a similarity score with the query \mathbf{q} , formalized as:

$$s^* = \arg \max_{s_i \in \mathcal{S}} \text{sim}(E_Q(\mathbf{q}), E_V(\mathbf{v}_{s_i})), \quad (4)$$

where $\mathcal{S} = \{s_1, s_2, \dots, s_N\}$ denotes the set of candidate video spans, $E_Q(\cdot)$ and $E_V(\cdot)$ are frozen encoders for the query and video, respectively, and $\text{sim}(\cdot, \cdot)$ denotes a similarity function such as cosine similarity.

For instance, Moment-GPT [104] extends the VTG-GPT [97] framework by combining frame-level captions from MiniGPT-v2 [136] with segment-level captions from Video-ChatGPT [2], matching them to textual queries using a similarity-based retrieval approach. TFVTG [95] further refines this paradigm by decomposing complex queries into sub-events using an LLM, followed by segment matching through the BLIP-2 Q-Former [75]. The final predictions are derived by integrating localized spans via temporally-aware filtering that accounts for the sub-event order and relations, improving the accuracy of temporal localization.

LLM-Driven Reasoning: This alternative strategy treats VTG as a high-level textual inference task, leveraging the reasoning capabilities of LLMs to comprehend and localize temporal segments based on enriched video descriptions.

For example, Grounding-prompter [96] reformulates VTG as a long-text comprehension task, aligning speech transcriptions and visual captions with timestamp annotations. It employs a four-step multiscale denoising chain-of-thought approach, progressively refining coarse temporal predictions through iterative prompts. In a similar vein, DeVl [30] uses Video-LLaVA [65] to perform hierarchical, multi-scale captioning, followed by query-driven refinement using GPT-4o [133]. This multi-stage reasoning process allows the model to better capture event dependencies and fine-grained temporal structures, leading to more accurate localization without additional training.

Beyond these dominant strategies, some emerging methods improve grounding by augmenting the input signals themselves, rather than altering the core model architecture. On the textual side, the pseudo-query generation and verification mechanism proposed by LoVe [146] enriches the query, enhancing zero-shot generalization for more robust and practical deployments. On the visual side, NumPro [109] embeds numerical identifiers directly onto video frames, providing temporal cues for LLMs to track sequences more effectively without requiring fine-tuning.

Summary: Training-free VTG-MLLMs provide a lightweight, modular alternative to conventional fine-tuning approaches, effectively decomposing VTG into manageable subtasks, i.e., captioning, matching, and reasoning, without the computational overhead of extensive training. By leveraging powerful pre-trained models and off-the-shelf components, these methods reduce the cost and complexity of domain-specific adaptation, making them a compelling choice for scenarios where data availability and computational resources are limited. However, their reliance on predefined embeddings and static representations can introduce challenges in capturing fine-grained temporal dependencies, presenting an ongoing area for innovation.

C. Video Feature Processing in VTG-MLLMs

At the most fine-grained level of our taxonomy, we examine the video feature extraction strategies that underpin VTG-MLLMs. As discussed in the taxonomy of functional roles, Facilitator-based methods often rely on pre-trained, frozen modules to provide high-level video embeddings. In contrast, Executor-oriented designs require more sophisticated mechanisms to handle raw video inputs, reflecting their more direct involvement in temporal reasoning and event localization. This subsection focuses on the critical strategies for extracting and processing visual and temporal features within Executor-based VTG-MLLMs.

1) *Efficient Visual Feature Handling*: Given the dense nature of frame-level information and the constrained input size of most LLMs, efficient visual feature handling is essential for capturing fine-grained temporal cues without overwhelming the model. These techniques can be broadly categorized into three main approaches, as illustrated in Fig. 5: *Learnable Token-Based Compression*, *Pooling-Based Compression*, and *Coarse-to-Fine Progressive Refinement*.

Learnable Token-Based Compression: This strategy employs learnable tokens to compress high-dimensional visual features into a concise and manageable representation, as depicted in Fig. 5(a). Formally, given a set of trainable tokens \mathbf{Q} and raw visual features \mathbf{V} , the compressed representation \mathbf{C} is obtained by a parameterized function f_{compress} :

$$\mathbf{C} = f_{\text{compress}}(\mathbf{Q}, \mathbf{V}; \theta) \quad (5)$$

where θ represents learnable parameters. For example, VTG-LLM [111] introduces Slot-Based Token Compression, where a set of learnable slot embeddings aggregate information from raw visual tokens based on similarity. TRACE [31] adopts a similar strategy, compressing dense visual inputs into compact

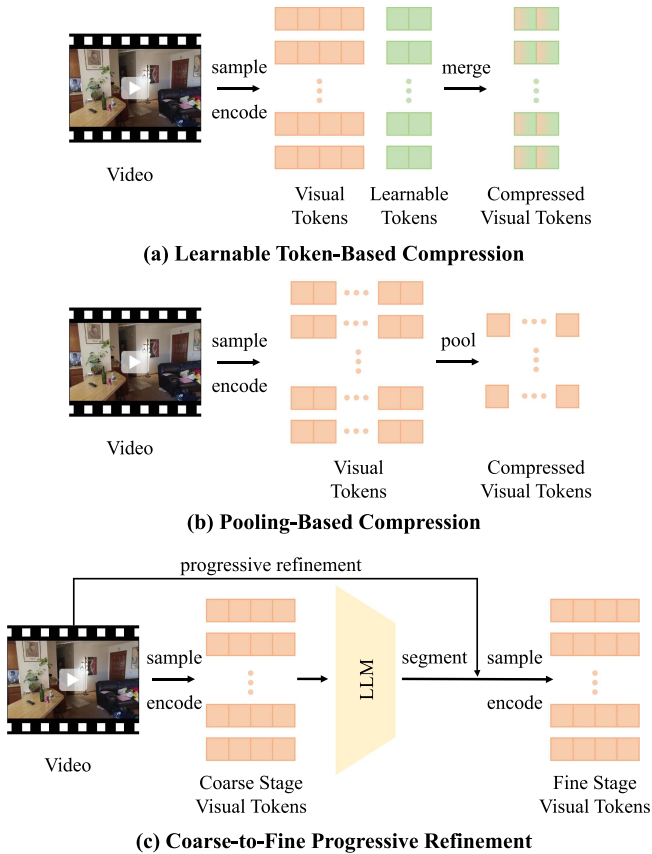


Fig. 5. Representative strategies for maximizing information utility within limited input token budgets: (a) Learnable Token-Based Compression, (b) Pooling-Based Compression, and (c) Coarse-to-Fine Progressive Refinement.

summaries using slot embeddings. ReVisionLLM [113] employs a [CLS]-like token [148] to aggregate segment features through self-attention, providing compact yet semantically rich representations, effectively acting as another instantiation of (5).

Pooling-Based Compression: Pooling techniques aggregate local or global visual features to reduce dimensionality yet retain key semantic information, as shown in Fig. 5(b). For instance, LITA [114] applies multi-granularity pooling across spatial and temporal dimensions, while Grounded-VideoLLM [34] and TimeMarker [117] use dynamically adjustable pooling kernels to capture hierarchical visual cues. LLaVA-MR [108] introduces token variance-based selection, prioritizing high-variance tokens to capture dynamic content more effectively. Alternatively, TimeSuite [33] reduces token overhead through token shuffling and projection, achieving compression without extra parameters.

Coarse-to-Fine Progressive Refinement: Rather than compressing dense frame-level features upfront, these methods progressively refine temporal predictions to improve efficiency under strict token constraints, as outlined in Fig. 5(c). SeViLA [105] is an early example that adopts a coarse-to-fine localization strategy by selecting language-aware keyframes before answer prediction. Similarly, HawkEye [32] adopts a recursive grounding approach, narrowing the temporal search space in iterative stages. Another prominent strategy involves

a multi-stage refinement—first identifying a coarse temporal segment and then adjusting its boundaries—a method effectively employed by ReVisionLLM [113], SlowFocus [119], and VideoMind [35] to reduce token overhead while preserving accuracy.

Summary: Efficient visual feature handling is paramount in VTG-MLLMs to reconcile the richness of dense video data with the input token limitations. The above-discussed strategies represent distinct philosophies for information reduction. Token-based and pooling methods achieve upfront compression. In contrast, progressive refinement adopts an iterative approach, selectively focusing computational resources on temporally relevant segments. Collectively, these techniques are crucial for enabling MLLMs to process detailed video sequences, striking a vital balance between capturing fine-grained temporal nuances necessary for precise grounding and maintaining computational tractability.

2) **Temporal Representation and Modeling:** Unlike global video understanding tasks, fine-grained temporal grounding demands precise reasoning about temporal relationships to align video frames with timestamp intervals. This requirement is critical for VTG tasks such as moment retrieval and dense video captioning, which rely on accurate temporal boundary predictions. To address these challenges, temporal feature modeling in VTG-MLLMs can be broadly categorized into *Explicit* and *Implicit* modeling strategies, distinguished by whether temporal cues are directly injected into the model’s input stream or contextually assimilated through its architectural design and reasoning capabilities.

Explicit Modeling: Explicit modeling strategies directly embed temporal information into the input or feature representations of MLLMs, as illustrated in Fig. 6(a)–(b). These approaches aim to provide precise temporal context by incorporating explicit time markers within the visual feature space, enhancing the model’s ability to align video frames with timestamps. Broadly, these methods can be categorized into two main approaches: *Temporal Embedding* and *Token Concatenation*, each with distinct mechanisms for integrating temporal cues.

Temporal Embedding: One common approach involves augmenting visual tokens with dedicated temporal embeddings, effectively integrating time information into the sequence of input tokens. If \mathbf{v}_i represents the i -th visual token and \mathbf{e}_{t_i} is its corresponding temporal embedding, the augmented visual token \mathbf{v}'_i can be formed as:

$$\mathbf{v}'_i = \mathbf{v}_i + \mathbf{e}_{t_i}. \quad (6)$$

For instance, VTG-LLM [111] introduces learnable absolute time embeddings initialized to zero, preserving the original semantic integrity of visual tokens generated by pre-trained encoders. In contrast, LITA [114] adopts a relative time representation, segmenting videos into T equal-length chunks and assigning unique temporal tokens (e.g., $\langle 1 \rangle$ to $\langle T \rangle$) to each segment, providing a coarse but computationally efficient temporal structure. Momentor [26] takes a more granular approach by defining N learnable anchor points, each representing a specific temporal position within the video. These anchors define a continuous temporal feature space, allowing for more precise

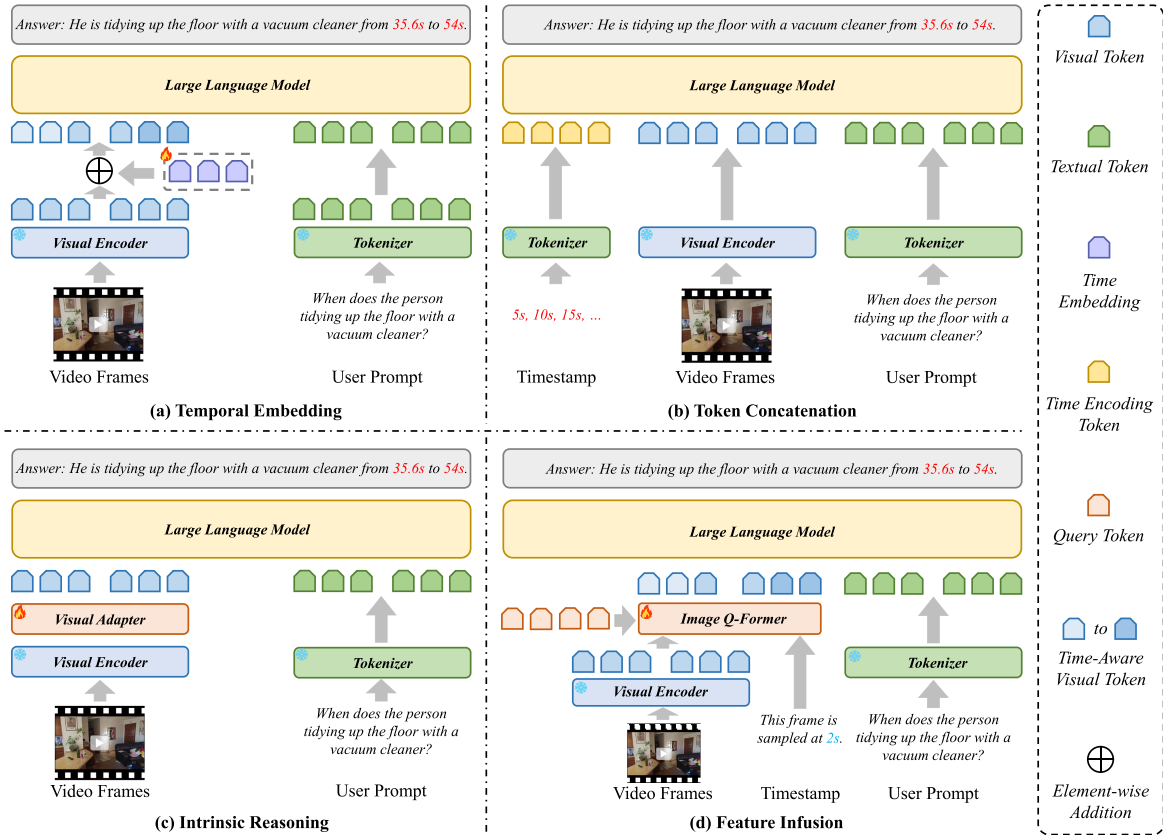


Fig. 6. Illustration of temporal feature processing strategies in VTG-MLLMs, categorized into (a)-(b) *Explicit Modeling*, which directly integrates timestamp information through methods like temporal embedding and token concatenation, and (c)-(d) *Implicit Modeling*, which relies on the reasoning capabilities of LLMs to infer temporal relationships through intrinsic reasoning and feature infusion.

localization through interpolation. Other methods also leverage explicit time encodings. For example, TGB [122], SlowFocus [119], and LLaVA-ST [116] incorporate temporal position embeddings to enhance temporal awareness during fine-tuning. These embeddings provide clear temporal context, improving temporal alignment without notably altering the underlying architecture of MLLMs.

Token Concatenation: An alternative approach involves tokenizing timestamps directly from sampled frames, integrating these temporal markers with visual and textual tokens to form a unified input sequence. If \mathbf{S}_P , \mathbf{S}_V , and \mathbf{S}_T represent sequences of prompt tokens, visual tokens, and tokenized timestamps, respectively, the final input sequence \mathbf{S}_{input} fed to the MLLM can be a concatenation:

$$\mathbf{S}_{input} = \text{Concat}(\mathbf{S}_P, \mathbf{S}_V, \mathbf{S}_T). \quad (7)$$

The specific order and interleaving strategy can vary. For instance, LLaVA-MR [108] dynamically selects relative frame indices or absolute timestamps based on the frame sampling rate, interleaving these markers with special tokens like $\langle time_begin \rangle$ and $\langle time_end \rangle$ to denote temporal boundaries. TimeMarker [117] adopts a similar strategy, inserting explicit temporal separators (e.g., “second2.0”) into the input sequence

to enhance temporal context. Additionally, Mr.BLIP [120] systematically explores various design choices for time representation, including relative versus absolute time, decimal versus integer formats, and different token ordering schemes, evaluating their impact on model performance. TRACE [31] further extends this approach by integrating temporal tokenization with visual feature embeddings, providing a tightly coupled representation of spatial and temporal information.

Implicit Modeling: Implicit modeling strategies aim to capture temporal relationships within video data through latent representations, leveraging the inherent reasoning and contextual understanding capabilities of large language models. Unlike explicit methods, which directly associate timestamps with visual inputs, implicit approaches integrate temporal cues more fluidly, embedding temporal knowledge without requiring explicit time markers. These strategies generally fall into two main categories: *Intrinsic Reasoning* and *Feature Infusion*, each employing distinct techniques to embed temporal context into visual representations, as illustrated in Fig. 6(c)–(d).

Intrinsic Reasoning: This approach relies on the LLM’s inherent ability to infer temporal relationships indirectly from the interplay between visual features and time-related language prompts. Rather than embedding explicit timestamps, these methods leverage numerical cues, iterative refinement, and boundary-aware reasoning to capture temporal dynamics. For

instance, NumPro [109] introduces numerical indices directly into video frames, allowing the LLM to infer sequence order through positional awareness. Grounded-VideoLLM [34] adopts a similar strategy, introducing specialized temporal tokens into the LLM’s vocabulary, enabling unified modeling of time and semantics. TimeRefine [115] reframes temporal grounding as a progressive refinement task, where the model first predicts coarse intervals (e.g., “15.0 s to 27.5 s”) and subsequently refines these estimates by predicting offset adjustments (e.g., “+4.0 s and -1.5 s”), achieving fine-grained localization through iterative reasoning. Other models, such as VTimeLLM [24] and TPE-VLLM [123], incorporate boundary-aware tasks during pretraining, explicitly teaching the model to reason about event durations and transitions, thereby enhancing temporal precision without the need for explicit time tokens.

Feature Infusion: Feature infusion techniques integrate temporal context directly into visual feature representations by conditioning the feature extraction process. This is often achieved using architectures like Q-Formers, designed to jointly learn spatiotemporal embeddings. Formally, given raw visual features \mathbf{V}_{raw} and a temporal descriptor \mathbf{T}_{desc} (e.g., “This frame is sampled at 2s”), the infused features $\mathbf{F}_{infused}$ can be generated as:

$$\mathbf{F}_{infused} = \text{Extractor}(\mathbf{V}_{raw}, \mathbf{T}_{desc}; \theta), \quad (8)$$

where Extractor (e.g., a Q-Former) processes \mathbf{V}_{raw} conditioned on \mathbf{T}_{desc} , allowing the model to capture subtle, context-dependent temporal cues without explicit time tokenization. For example, TimeChat [25] and TemporalVLM [121] leverage this by providing such temporal descriptors as conditional inputs to their Q-Former, guiding the model to incorporate temporal context. Similarly, TimeSuite [33] generates segment-level features that capture temporal dynamics across longer intervals, enabling more comprehensive temporal reasoning.

Summary: Temporal representation and modeling are foundational for endowing VTG-MLLMs with the capacity for precise temporal localization. Explicit modeling strategies directly furnish MLLMs with unambiguous temporal information, offering direct control and interpretability over temporal cues. Implicit modeling, on the other hand, leverages the inherent sequential processing and reasoning capabilities of LLMs or integrates temporal context more subtly during feature extraction. These approaches reflect ongoing exploration into how best to integrate the continuous nature of time with the discrete, symbolic processing of LLMs, ultimately shaping the model’s ability to perform nuanced temporal reasoning and accurate boundary prediction.

IV. EXPERIMENTAL EVALUATION

In this section, we provide a comprehensive analysis of the performance of various VTG-MLLMs across four core tasks. We begin by establishing the experimental settings, outlining the benchmark datasets and evaluation metrics. Subsequently, we present a detailed performance comparison, considering both zero-shot and fine-tuning scenarios to assess their generalization

capabilities and task-specific effectiveness, with further in-depth analyses of model behavior presented in Appendix B.

A. Experimental Settings

This subsection outlines the experimental settings. We first introduce the *Benchmark Datasets* used for evaluation, followed by a description of the *Evaluation Metrics* employed to assess model performance across different VTG tasks.

1) *Benchmark Datasets:* Benchmark datasets are foundational for the development and rigorous evaluation of VTG research, providing essential video content and corresponding natural language annotations necessary for model training and standardized comparisons. The current VTG dataset landscape can be broadly categorized into two groups: *Standard Benchmarks*, which are well-established and widely adopted by the research community, and *Emerging Benchmarks*, designed to address gaps in existing datasets by targeting complex scenarios such as long-term temporal reasoning and fine-grained event distinctions.

Given space limitations, we focus this subsection on key characteristics of select standard benchmarks that are utilized in our experimental evaluations (Sections IV-B and IV-C). Comparative statistics for these datasets are summarized in Table I. A comprehensive review, including additional standard and emerging benchmarks, is provided in the Appendix A.

Charades-STA [149], derived from the Charades dataset [151], is widely employed for MR. It includes approximately 6.7 k videos depicting indoor daily activities, annotated with about 16.1 k query-moment pairs. Videos average 30 seconds, moments average 8 seconds, and textual queries are concise (7.2 words on average), making it suitable for action-centric grounding tasks.

ActivityNet-Captions [6], an extension of the human activity dataset ActivityNet [152], is primarily utilized for DC, though also popular in MR tasks. It features around 15 k videos annotated with multiple temporally localized events per video (avg. 3.65 events per video). With an average video duration of 1.96 minutes, average moment length of 36.2 seconds, and longer descriptive queries (14.8 words), making it suitable for evaluating models on complex temporal structures and narratives.

YouCook2 [150] offers a large-scale instructional video dataset tailored for procedural activity understanding, primarily applied in DC. It comprises 2,000 long, untrimmed cooking videos, collectively spanning 176 hours and covering 89 distinct recipes, resulting in approximately 15.4 k query-moment pairs. Each video is densely annotated with precise temporal boundaries for step-level imperative sentences (avg. 8.8 words).

QVHighlights [53] is specifically designed for query-based HD and also finds application in MR. The dataset includes approximately 10.1 k YouTube videos, each cropped to a fixed duration of 150 seconds, and features around 10.3 k query-moment pairs. A distinctive aspect of QVHighlights is its allowance for multiple disjoint moments per query, reflecting the non-linear nature of salient video highlights, which average about 24.6

TABLE I

COMPARATIVE STATISTICS OF MAJOR DATASETS COMMONLY USED IN VTG. THESE DATASETS SPAN A RANGE OF TASKS, INCLUDING MR, DC, HD, AND GQA, AS DESCRIBED IN SECTION II-A. KEY DATASET ATTRIBUTES INCLUDE THE NUMBER OF VIDEOS, AVERAGE VIDEO DURATION, AVERAGE MOMENT DURATION, TOTAL NUMBER OF QUERIES, AND AVERAGE QUERY LENGTH, PROVIDING A COMPREHENSIVE OVERVIEW OF THEIR SCALE AND COMPLEXITY

Dataset	Year	Task	#Videos	Duration / Video	Duration / Moment	#Queries	#Words / Query
Charades-STA [149]	2017	MR	6,672	0.5min	8.1s	16,128	7.2
ActivityNet-Captions [6]	2017	MR & DC	14,926	1.96min	36.2s	71,957	14.8
YouCook2 [150]	2017	DC	2,000	5.26min	-	15,433	8.8
QVHighlights [53]	2021	MR & HD	10,148	2.5min	24.6s	10,310	11.3
NExT-GQA [10]	2023	GQA	1,557	0.70min	6.92s	8,911	-

seconds. Evaluation of its test set is managed via an official server¹.

NExT-GQA [10] extends the NExT-QA dataset [153], designed to support research on QA tasks that require temporal grounding as evidence. It adds precise temporal annotations for approximately 8.9 k QA pairs across 1.6 k videos, pinpointing critical segments (avg. 6.9 seconds) needed to answer questions. A notable feature is its weakly-supervised setup where only validation and test sets provide ground-truth temporal labels.

2) *Evaluation Metrics*: Effective evaluation of VTG-MLLMs relies on well-defined metrics that assess both the temporal localization accuracy and semantic alignment of predicted moments with ground-truth segments.

Video Moment Retrieval is quantified by two key metrics for localization accuracy: *Mean Intersection over Union (mIoU)* [149] quantifies the average overlap between predicted and ground-truth segments, indicating localization accuracy; and *Recall at Rank n with an IoU threshold of m ($R@n$ ($IoU=m$))* [154] quantifies the percentage of cases where the top- n retrieved segments have an IoU with the ground truth greater than or equal to the threshold m , under varying IoU thresholds (e.g., $R@1$ at $IoU=0.3/0.5/0.7$).

Dense Video Captioning hinges on both temporal precision and linguistic quality. Key metrics include: *SODA_c* [155] assesses structural alignment between captions and temporal segments, emphasizing event segmentation coherence; *ME-TEOR* [156] and *CIDEr* [157] evaluate linguistic similarity between generated and reference captions, considering language precision and recall; and *F1 Score* measures the harmonic mean of precision and recall for instances within a video, assessing coverage and relevance.

Video Highlight Detection commonly relies on metrics: *Mean Average Precision (mAP)* measures the area under the precision-recall curve at various IoU thresholds, such as 0.5 and 0.75, reflecting the model’s ability to accurately rank and localize key moments; and *HIT@1* [158] calculates the proportion of instances where the top-ranked highlight prediction matches a ground-truth highlight, providing a straightforward measure of first-choice accuracy.

Temporally Grounded Video Question Answering synthesizes metrics for both temporal localization and question-answering accuracy: *mIoU* follows the same setting as in MR; *Mean Intersection over Prediction (mIoP)* [10] evaluates the proportion of the predicted segment overlapping with the ground-truth

moment, as an alternative alignment metric; *Acc@QA* [159] measures the percentage of correctly answered questions, irrespective of temporal grounding; and *Acc@GQA* [10] extends *Acc@QA* by requiring correct answers with accurate temporal grounding, where the predicted segment must achieve an IoP score of at least 0.5.

B. Zero-Shot Performance Comparison

Zero-shot evaluation measures the ability of VTG-MLLMs to generalize to new datasets without any dataset-specific fine-tuning. This evaluation is crucial for understanding the inherent adaptability and robustness of different models, reflecting their capacity to transfer knowledge across diverse video grounding tasks. For clarity, models are categorized based on their training paradigms: *Pretraining (PT)* approaches (Section III-B1), which rely on large-scale multimodal datasets to build general-purpose grounding capabilities, and *Training-Free (TF)* approaches (Section III-B3), which leverage pre-trained foundation models without requiring additional task-specific training.

Video Moment Retrieval: Table II presents a comparison of zero-shot performance on the Charades-STA [149] and ActivityNet-Captions [6] benchmarks.

PT Approaches: Among pretraining-based methods, models like Time-R1 [127], VideoMind [35], and TimeMarker [117] demonstrate state-of-the-art zero-shot performance, particularly on the Charades-STA dataset. Time-R1, in particular, success in achieving the highest recall scores likely stems from its innovative use of reinforcement learning. Similarly, TRACE [31] and TPE-VLLM [123] achieve competitive results on ActivityNet-Captions, benefiting from their fine-grained temporal modeling and boundary-aware pertaining.

TF Approaches: In the training-free category, methods like TFVTG [95] and VTG-GPT [97] also exhibit competitive performance, particularly on Charades-STA. Notably, TFVTG outperforms several pretraining-based models in terms of mIoU and $R@0.5$. These models benefit from flexible, modular designs that leverage powerful pre-trained language and vision models without requiring task-specific fine-tuning.

Dense Video Captioning: Table III presents zero-shot performance comparisons on the ActivityNet-Captions [6] and YouCook2 [150] benchmarks.

PT Approaches: Within the pretraining category, methods like TRACE [31], Grounded-VideoLLM [34], and VTimeLLM [24] consistently achieve strong results across key metrics, including

¹Evaluation: <https://codalab.lisn.upsaclay.fr/competitions/6937>

TABLE II

ZERO-SHOT PERFORMANCE COMPARISON ON VIDEO MOMENT RETRIEVAL BENCHMARKS (CHARADES-STA [149] AND ACTIVITYNET-CAPTIONS [6]). METHODS ARE GROUPED BY TRAINING PARADIGM (PT: PRETRAINING, TF: TRAINING-FREE). THE BEST RESULTS ARE **BOLDFACED**, AND THE SECOND-BEST RESULTS ARE UNDERLINED. THIS FORMATTING CONVENTION IS UNIFORMLY ADOPTED ACROSS ALL SUBSEQUENT TABLES UNLESS EXPLICITLY NOTED OTHERWISE

Method	Paradigm	Charades-STA [149]				ActivityNet-Captions [6]			
		mIoU	R@0.3	R@0.5	R@0.7	mIoU	R@0.3	R@0.5	R@0.7
Time-R1 [127]	PT	-	78.1	60.8	35.3	-	-	-	-
VideoMind [35]	PT	50.2	<u>73.5</u>	<u>59.1</u>	<u>31.2</u>	33.3	48.4	30.3	15.7
TimeMarker [117]	PT	<u>48.4</u>	<u>73.5</u>	51.9	26.9	-	-	-	-
VideoChat-T [33]	PT	-	69.9	48.7	24.0	-	-	-	-
LLaVA-ST [116]	PT	42.4	63.1	44.8	23.4	-	-	-	-
MeCo [126]	PT	-	-	44.4	17.5	-	-	-	-
E.T.Chat [50]	PT	-	-	43.2	19.4	-	-	-	-
VideoExpert [128]	PT	41.1	61.5	40.3	20.9	-	-	-	-
TRACE [31]	PT	-	-	40.3	19.4	39.0	-	37.7	24.0
VideoChat-TPO [124]	PT	38.1	58.3	40.2	18.4	-	-	-	-
MLLM-TA [107]	PT	-	-	37.9	18.1	-	-	27.6	18.3
Ground-VideoLLM [34]	PT	36.8	54.2	36.4	19.7	<u>36.1</u>	46.2	30.3	19.0
VTG-LLM [111]	PT	-	-	33.8	15.7	-	-	-	-
TPE-VLLM [123]	PT	34.7	55.5	33.1	14.7	<u>36.1</u>	50.4	<u>35.4</u>	<u>19.2</u>
TimeChat [25]	PT	-	-	32.2	13.4	-	-	-	-
HawkEye [32]	PT	33.7	50.6	31.4	14.5	32.7	<u>49.1</u>	29.3	10.7
TemporalVLM [121]	PT	-	-	30.1	13.2	-	-	-	-
GroundingGPT [112]	PT	-	-	29.6	11.9	-	-	-	-
VTimeLLM [24]	PT	31.2	51.0	27.5	11.4	30.4	44.0	27.8	14.3
Momentor [26]	PT	28.5	42.6	26.6	11.6	29.3	42.9	23.0	12.4
TFVTG [95]	TF	44.5	67.0	50.0	<u>24.3</u>	34.1	<u>49.3</u>	27.0	<u>13.4</u>
VTG-GPT [97]	TF	<u>39.8</u>	<u>59.5</u>	<u>43.7</u>	25.9	30.5	47.1	<u>28.3</u>	12.8
Moment-GPT [104]	TF	36.5	58.2	38.4	21.6	<u>30.8</u>	48.1	31.1	14.9
ChatVTG [29]	TF	34.9	52.7	33.0	15.9	27.2	40.7	22.5	9.4

TABLE III

ZERO-SHOT PERFORMANCES ON DENSE VIDEO CAPTIONING BENCHMARKS (ACTIVITYNET-CAPTIONS [6] AND YOUCOOK2 [150])

Method	Paradigm	ActivityNet-Captions [6]				YouCook2 [150]		
		SODA_c	CIDEr	METEOR	F1 Score	SODA_c	CIDEr	F1 Score
TRACE [31]	PT	6.0	<u>25.9</u>	<u>6.4</u>	39.3	2.2	8.1	22.4
Ground-VideoLLM [34]	PT	6.0	-	6.8	-	-	-	-
VTimeLLM [24]	PT	<u>5.8</u>	27.6	6.8	-	0.9	3.4	-
VideoExpert [128]	PT	-	-	-	-	<u>2.1</u>	<u>6.0</u>	-
VTG-LLM [111]	PT	5.1	20.7	5.9	34.8	1.5	5.0	<u>17.5</u>
TimeChat [25]	PT	4.7	19.0	5.7	<u>36.9</u>	1.2	3.4	12.6
Momentor [26]	PT	2.3	14.9	4.7	-	-	-	-
TemporalVLM [121]	PT	-	-	-	-	1.2	3.7	13.1

TABLE IV

ZERO-SHOT PERFORMANCE COMPARISON ON VIDEO HIGHLIGHT DETECTION BENCHMARK QVHIGHLIGHTS [53]

Method	Paradigm	mAP	HIT@1
VideoExpert [128]	PT	35.8	<u>52.7</u>
TRACE [31]	PT	<u>26.8</u>	42.7
VideoChat-T [33]	PT	26.5	54.1
MLLM-TA [107]	PT	19.2	31.8
VTG-LLM [111]	PT	16.5	33.5
TemporalVLM [121]	PT	16.4	31.3
TimeChat [25]	PT	14.5	23.9
Momentor [26]	PT	7.6	-

SODA_c, CIDEr, and METEOR. TRACE stands out as a top performer on both ActivityNet-Captions and YouCook2, reflecting its effective integration of fine-grained temporal encoding and sequence generation.

TF Approaches: Training-free methods are generally not included in this evaluation, as they often rely on offline MLLMs for generating descriptions, which do not directly align with the standard online evaluation protocols typically used for dense captioning.

Video Highlight Detection: Table IV presents zero-shot performance comparisons on the QVHighlights [53] benchmark, a challenging dataset specifically designed for localizing query-relevant highlight moments within videos.

PT Approaches: Among pretraining-based methods, VideoExpert [128] stands out with the highest mAP scores, reflecting its strong ability to capture fine-grained temporal relationships and accurately rank highlight moments. In contrast, videoChat-T [33] leads in Hit@1, demonstrating superior single-shot localization accuracy, which highlights its effectiveness in identifying salient temporal regions.

TF Approaches: Training-free methods are typically excluded from direct comparison in this context, as they often

TABLE V
ZERO-SHOT PERFORMANCE COMPARISON ON TEMPORALLY GROUNDED VIDEO QA BENCHMARK NEXT-GQA [10]

Method	Paradigm	mIoP	IoP@0.3	IoP@0.5	mIoU	IoU@0.3	IoU@0.5	Acc@GQA	Acc@QA
DeVi [30]	TF	39.3	-	37.9	22.3	-	17.4	<u>28.0</u>	71.6
VideoMind [35]	PT	<u>39.0</u>	56.0	<u>35.3</u>	31.4	50.2	25.8	28.2	-
VideoChat-TPO [124]	PT	35.6	<u>47.5</u>	32.8	27.7	<u>41.2</u>	<u>23.4</u>	25.5	-
VideoExpert [128]	PT	34.6	45.3	29.3	<u>27.9</u>	41.0	22.4	21.6	<u>71.1</u>
HawkEye [32]	PT	-	-	-	25.7	37.0	19.5	-	-
Ground-VideoLLM [34]	PT	34.5	42.6	34.4	21.1	30.2	18.0	26.7	-
SeViLA [105]	PT	29.5	34.7	22.9	21.7	29.2	13.8	16.6	68.1
VTimeLLM [24]	PT	<u>27.9</u>	30.3	23.8	18.3	27.7	14.1	12.7	-

rely on frame-query similarity scoring, which inherently aligns well with highlight detection tasks. These approaches, while effective, lack the structured, end-to-end temporal grounding capabilities required for competitive performance in strictly defined zero-shot scenarios.

Temporally Grounded Video Question Answering: Table V presents zero-shot performance on the NEXT-GQA [10] benchmark, a relatively recent dataset designed to evaluate temporally grounded question answering in videos.

PT Approaches: Among pretraining-based methods, VideoMind [35] showcases exceptional strength in temporal localization, achieving the highest scores across all IoU-related metrics. This proficiency is complemented by its outstanding Acc@GQA performance, largely attributed to its methodical and step-by-step approach to solving GQA problems. In parallel, VideoChat-TPO [124] demonstrates robust and competitive results, underscoring the efficacy of its unique task preference optimization strategy.

TF Approaches: In the training-free category, DeVi [30] stands out, achieving the highest scores in mIoP, Acc@GQA, and Acc@QA. This indicates a strong capacity for fine-grained video temporal understanding, likely due to the carefully designed pipeline that effectively combines the perceptual capabilities of the pretrained MLLMs with the reasoning ability of LLMs.

Summary: Experimental results across the four tasks suggest that *PT Approaches* generally exhibit strong zero-shot performance due to their explicit temporal learning and large-scale data exposure, particularly excelling in tasks like moment retrieval. While *TF Approaches* are not directly comparable in dense captioning and highlight detection due to architectural differences, models like DeVi [30] show that with well-designed pipelines, *TF Approaches* can still achieve strong temporal reasoning in suitable tasks.

C. Fine-Tuning Performance Comparison

Fine-tuning evaluation assesses the performance of VTG-MLLMs after they are trained on the train-split of specific benchmark datasets. Unlike zero-shot evaluation, which tests the generalization ability of models without exposure to task-specific data, this setting allows models to adapt to domain-specific distributions and task requirements, often resulting in substantial performance gains. Notably, the NEXT-GQA [10] benchmark is excluded from this evaluation, as its dataset definition does not provide a dedicated training split, precluding fine-tuning.

TABLE VI
FINE-TUNING PERFORMANCE COMPARISON ON VIDEO MOMENT RETRIEVAL BENCHMARK CHARADES-STA [149]

Charades-STA [149]				
Method	Paradigm	mIoU	R@0.5	R@0.7
LLaVA-MR [108]	FT	59.8	70.7	49.6
Mr.BLIP [120]	FT	<u>58.6</u>	<u>69.3</u>	<u>49.3</u>
VideoLights [100]	FT	52.9	62.0	41.1
ReCorrect [103]	FT	48.7	54.4	31.1
LMR [102]	FT	-	55.9	35.2
LLaViLo [106]	FT	-	55.7	33.4
Time-R1 [127]	PT	-	72.2	50.1
VideoChat-R1 [129]	PT	60.8	<u>71.7</u>	50.2
VideoChat-T [33]	PT	-	67.1	43.0
VideoChat-TPO [124]	PT	<u>55.0</u>	65.0	40.7
SpaceVLLM [125]	PT	-	63.6	38.5
TRACE [31]	PT	-	61.7	41.4
VideoExpert [128]	PT	52.2	60.8	36.5
HawkEye [32]	PT	49.3	58.3	28.8
VTG-LLM [111]	PT	-	57.2	33.4
TemporalVLM [121]	PT	-	54.5	29.0
MLLM-TA [107]	PT	-	48.9	25.3
TimeChat [25]	PT	-	46.7	23.7

To better reflect differences in training paradigms, we categorize models into two types: *Pretraining (PT)* (Section III-B1) and *Fine-Tuning (FT)* (Section III-B2). The *PT* refers to models that undergo large-scale multimodal pretraining prior to task-specific fine-tuning on downstream training splits. In contrast, the *FT* is specifically designed to be trained directly on the training split of the downstream task.

Video Moment Retrieval: Table VI presents the fine-tuning performance on Charades-STA [149].

PT Approaches: Pretraining-based models exhibit significant performance enhancements when fine-tuned on the benchmark. Notably, reinforcement learning-based methods, Time-R1 [127] and VideoChat-R1 [129], emerge as the top performers, underscoring the profound potential of RL in the MLLM fine-tuning process. Closely following these leaders, models like VideoChat-T [33] and VideoChat-TPO [124] also deliver highly competitive outcomes. These results underscore the benefit of large-scale pretraining as a foundation for task-specific adaptation, allowing models to further refine their temporal grounding capabilities.

FT Approaches: Models designed for direct fine-tuning also exhibit remarkable performance, particularly on Charades-STA. LLaVA-MR [108] leads with the highest mIoU, R@0.5, and R@0.7, outperforming all PT methods on these metrics.

TABLE VII
FINE-TUNING PERFORMANCE COMPARISON ON DENSE VIDEO CAPTIONING
BENCHMARK YOU-COOK2 [150]

Method	Paradigm	SODA_c	CIDEr	F1 Score
TRACE [31]	PT	6.7	35.5	31.8
VideoExpert [128]	PT	4.2	18.7	-
VTG-LLM [111]	PT	3.6	13.4	20.6
TimeChat [25]	PT	3.4	11.0	19.5
TemporalVLM [121]	PT	3.4	13.2	20.0

TABLE VIII
FINE-TUNING PERFORMANCE COMPARISON ON VIDEO HIGHLIGHT DETECTION
BENCHMARK QVHIGHLIGHTS [53]

Method	Paradigm	mAP	HIT@1
VideoChat-TPO [124]	PT	38.8	66.2
VideoExpert [128]	PT	36.1	61.0
TRACE [31]	PT	31.8	51.5
VideoChat-T [33]	PT	27.0	55.3
TemporalVLM [121]	PT	25.1	43.0
VTG-LLM [111]	PT	24.1	41.3
MLLM-TA [107]	PT	23.9	40.1
TimeChat [25]	PT	21.7	37.9

Mr.BLIP [120] is a close second with mIoU and strong recall scores. On ActivityNet-Captions, Mr.BLIP [120] achieves the best R@0.5 and R@0.7 across all reported models. This demonstrates that well-optimized fine-tuning can effectively exploit the strong generalization capabilities of existing large vision-language models without requiring extensive pretraining related to temporal understanding.

Dense Video Captioning and Video Highlight Detection: The fine-tuning performance for DC on YouCook2 [150] and HD on QVHighlights [53] are presented in Tables VII and VIII, respectively.

PT Approaches: For DC on YouCook2, TRACE [31] stands out significantly among pretraining-based methods, achieving the highest SODA_c, CIDEr, and F1 Score. VideoExpert [128] also shows competitive results, particularly in SODA_c and CIDEr. Turning to HD on QVHighlights, VideoChat-TPO [124] demonstrates exceptional performance among PT approaches, leading with the highest mAP and Hit@1. VideoExpert [128] also achieves a competitive mAP, while VideoChat-T [33] shows a strong Hit@1 score. For both DC and HD, these results show that PT models can be effectively adapted to the specific nuances of these tasks through fine-tuning.

FT Approaches: For both DC and HD, results for *FT approaches* are absent in the presented tables under standard fine-tuning protocols. This is primarily due to two factors. Firstly, *FT approaches* of the “Offline Textualization with MLLMs” type (as described in Section III-B2) adopt architectures incompatible with the evaluation settings of DC and HD, which is analogous to the exclusion of *TF approaches* from zero-shot evaluations on these benchmarks. Secondly, “Direct Fine-Tuning of MLLMs” for these specific tasks, while theoretically possible, remains unreported in current public research efforts. The complexity of DC (generating multiple, temporally ordered descriptions) and the nuanced requirements of HD (identifying subtle, query-relevant highlights) may necessitate more specialized architectural

adaptations or fine-tuning strategies for direct application of MLLMs, which are less explored in current literature compared to their application in MR.

Summary: Fine-tuning enables VTG-MLLMs to better align with task-specific objectives and dataset characteristics, leading to notable performance gains. *PT Approaches* like TRACE [31] and VideoExpert [128] benefit significantly from this process, consistently ranking among top performers across tasks. Meanwhile, *FT Approaches* such as LLaVA-MR [108] and Mr.BLIP [120] achieve competitive or even superior results in MR, highlighting the effectiveness of direct supervision. The absence of *FT Approaches* in tasks like DC and HD reflects current architectural limitations, marking this as an underexplored yet promising direction.

V. LIMITATIONS AND FUTURE DIRECTIONS

While the integration of MLLMs has significantly advanced VTG, several fundamental limitations persist, defining the key challenges and future frontiers for the field. This section outlines these critical areas, spanning from immediate technical bottlenecks to broader considerations of reliability and forward-looking paradigms.

Efficient and Expressive Feature Representation: A fundamental challenge for VTG-MLLMs lies in reconciling the richness of video data with the limited input capacity of LLMs. Long-form or high-resolution videos require vast numbers of visual tokens, making fine-grained grounding computationally prohibitive. While existing feature compression and selection strategies (Section III-C1) provide partial solutions, they risk discarding subtle yet crucial temporal cues. Concurrently, approaches to temporal representation—ranging from explicit timestamp encoding to implicit knowledge learning (Section II-I-C2)—remain diverse, and their comparative trade-offs are not yet fully understood.

Addressing these challenges requires advances in both efficiency and expressiveness. On the efficiency side, future research should explore adaptive token selection strategies, such as saliency-based pruning and hierarchical aggregation, to dynamically filter redundant information while preserving critical details [160], [161]. On the expressiveness side, a unified temporal encoding framework is needed that goes beyond simple positions or timestamps and captures richer temporal attributes such as duration, order, and causality. This could be achieved by explicitly incorporating temporal reasoning objectives into pretraining, enabling models to develop a more native and robust understanding of time.

Expanding to Richer Modalities: While most VTG-MLLMs rely solely on visual and textual inputs, incorporating additional modalities—particularly audio—offers a promising avenue for improving temporal grounding. Auditory cues such as speech, sound effects, and ambient noise often provide complementary or even superior temporal signals, especially in visually ambiguous scenes [162], [163]. However, effective multimodal integration remains challenging due to difficulties in synchronizing heterogeneous data streams, the scarcity of large-scale annotated multimodal datasets, and increased computational

demands. Future work should focus on multimodal fusion frameworks with high temporal fidelity, leveraging advanced training paradigms, adaptive cross-modal attention, and synthetic data generation to reduce annotation costs. Integrating pre-trained audio-language models (e.g., Whisper [164]) within VTG pipelines further holds promise for enhancing temporal reasoning in complex, multi-channel video content.

Reliability and Trustworthiness: As VTG-MLLMs grow more capable, ensuring their reliability and trustworthiness becomes critical. Our error and failure mode analyses (Appendix B.3, B.4) reveal persistent challenges, including difficulty achieving high-precision temporal localization across VTG tasks—a key reliability gap with potentially severe consequences in safety-critical domains. A particularly concerning failure mode is the “decoupling of reasoning from grounding,” observed in GQA, where models produce correct answers while referencing irrelevant evidence, creating a misleading sense of explainability.

In addition, the generative nature of MLLMs raises risks of hallucination, bias, and malicious use. VTG-MLLMs may hallucinate events, misrepresent localized segments, or enable unethical applications such as large-scale surveillance and targeted misinformation. These risks underscore the need for robust safeguards, including uncertainty estimation, automated fact-checking, and clear ethical guidelines to ensure responsible deployment.

Forward-Looking Perspectives: The trajectory of VTG-MLLMs points toward reasoning over long-form videos, a challenge transcending mere token capacity to require innovations in hierarchical memory and efficient spatio-temporal indexing. As these capabilities mature, advanced VTG will become a cornerstone for general-purpose video agents. The ultimate vision, however, extends to enabling embodied AI to interact with the physical world. This demands a paradigm shift from merely localizing what happened to understanding why it happened and predicting what comes next—a feat requiring the integration of VTG with world models [165]. In this synthesis, VTG would anchor language to spatio-temporal evidence, while the world model supplies the intrinsic knowledge of causality and physics that current MLLMs lack. This positions temporal grounding not just as a retrieval task, but as a foundational component for building genuinely intelligent and interactive systems.

VI. CONCLUSION

In this survey, we provide a comprehensive review of VTG-MLLMs, a critical area in fine-grained video understanding. MLLMs have significantly transformed VTG by introducing advanced reasoning and cross-modal alignment capabilities that surpass traditional task-specific methods. We systematically categorize and analyze current approaches, distinguishing between the functional roles of MLLMs as facilitators and executors. Additionally, we examine diverse training paradigms, including pretraining, fine-tuning, and training-free methods, alongside cutting-edge video feature processing techniques, such as visual feature compression and explicit versus implicit temporal modeling. To support ongoing research, we provide a structured overview of widely used datasets and benchmark, highlighting

key performance trends across different VTG-MLLM architectures. Finally, we summarize current challenges and outline promising directions for advancing VTG research.

REFERENCES

- [1] Y. Wang et al., “InternVideo: General video foundation models via generative and discriminative learning,” 2022, *arXiv:2212.03191*.
- [2] M. Maaz, H. Rasheed, S. Khan, and F. S. Khan, “Video-chatGPT: Towards detailed video understanding via large vision and language models,” in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2024, pp. 12585–12602.
- [3] M. Wang, J. Xing, J. Mei, Y. Liu, and Y. Jiang, “ActionCLIP: Adapting language-image pretrained models for video action recognition,” *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 36, no. 1, pp. 625–637, Jan. 2025.
- [4] M. Regneri, M. Rohrbach, D. Wetzel, S. Thater, B. Schiele, and M. Pinkal, “Grounding action descriptions in videos,” in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, vol. 1, pp. 25–36, 2013.
- [5] L. Anne Hendricks, O. Wang, E. Shechtman, J. Sivic, T. Darrell, and B. Russell, “Localizing moments in video with natural language,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5803–5812.
- [6] R. Krishna, K. Hata, F. Ren, L. Fei-Fei, and J. C. Niebles, “Dense-captioning events in videos,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 706–715.
- [7] T. Wang, R. Zhang, Z. Lu, F. Zheng, R. Cheng, and P. Luo, “End-to-end dense video captioning with parallel decoding,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 6847–6857.
- [8] Y. Song, J. Vallmitjana, A. Stent, and A. Jaimes, “TVSum: Summarizing web videos using titles,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 5179–5187.
- [9] B. Xiong, Y. Kalantidis, D. Ghadiyaram, and K. Grauman, “Less is more: Learning highlight detection from video duration,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 1258–1267.
- [10] J. Xiao, A. Yao, Y. Li, and T.-S. Chua, “Can i trust your answer? Visually grounded video question answering,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 13204–13214.
- [11] J.-J. Chen, Y.-C. Liao, H.-C. Lin, Y.-C. Yu, Y.-C. Chen, and F. Wang, “Rextime: A benchmark suite for reasoning-across-time in videos,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2024, pp. 28662–28673.
- [12] D. Zhang, X. Dai, X. Wang, Y.-F. Wang, and L. S. Davis, “Man: Moment alignment network for natural language moment retrieval via iterative graph adjustment,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 1247–1257.
- [13] W. Moon, S. Hyun, S. Park, D. Park, and J.-P. Heo, “Query-dependent video representation for moment retrieval and highlight detection,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23023–23033.
- [14] M. Ma, S. Yoon, J. Kim, Y. Lee, S. Kang, and C. D. Yoo, “VLANet: Video-language alignment network for weakly-supervised video moment retrieval,” in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 156–171.
- [15] M. Zhang et al., “Multi-stage aggregated transformer network for temporal language localization in videos,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 12669–12678.
- [16] Y. Liu, S. Li, Y. Wu, C.-W. Chen, Y. Shan, and X. Qie, “UMT: Unified multi-modal transformers for joint video moment retrieval and highlight detection,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 3042–3051.
- [17] Y. Yuan, T. Mei, and W. Zhu, “To find where you talk: Temporal sentence localization in video with attention based location regression,” in *Proc. AAAI Conf. Artif. Intell.*, 2019, pp. 9159–9166.
- [18] S. Zhang et al., “OPT: Open pre-trained transformer language models,” 2022, *arXiv:2205.01068*.
- [19] H. W. Chung et al., “Scaling instruction-finetuned language models,” *J. Mach. Learn. Res.*, vol. 25, no. 70, pp. 1–53, 2024.
- [20] D. Guo et al., “Deepseek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning,” 2025, *arXiv:2501.12948*.
- [21] H. Liu, C. Li, Q. Wu, and Y. J. Lee, “Visual instruction tuning,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 34892–34916.
- [22] H. Zhang, X. Li, and L. Bing, “Video-LLaMA: An instruction-tuned audio-visual language model for video understanding,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2023, pp. 543–553.
- [23] L. Xu, Y. Zhao, D. Zhou, Z. Lin, S. K. Ng, and J. Feng, “PLLaVA: Parameter-free LLaVA extension from images to videos for video dense captioning,” 2024, *arXiv:2404.16994*.

- [24] B. Huang, X. Wang, H. Chen, Z. Song, and W. Zhu, "VTimeLLM: Empower LLM to grasp video moments," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 14271–14280.
- [25] S. Ren, L. Yao, S. Li, X. Sun, and L. Hou, "Timechat: A time-sensitive multimodal large language model for long video understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 14313–14323.
- [26] L. Qian et al., "Momentor: Advancing video large language model with fine-grained temporal reasoning," in *Proc. Int. Conf. Mach. Learn.*, 2024, pp. 41340–41356.
- [27] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Representations*, 2020, pp. 1–21.
- [28] Q. Sun, Y. Fang, L. Wu, X. Wang, and Y. Cao, "EVA-clip: Improved training techniques for clip at scale," 2023, *arXiv:2303.15389*.
- [29] M. Qu, X. Chen, W. Liu, A. Li, and Y. Zhao, "ChatVTG: Video temporal grounding via chat with video dialogue large language models," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 1847–1856.
- [30] H. Qin, J. Xiao, and A. Yao, "Question-answering dense video events," 2024, *arXiv:2409.04388*.
- [31] Y. Guo, J. Liu, M. Li, X. Tang, Q. Liu, and X. Chen, "Trace: Temporal grounding video LLM via causal event modeling," in *Proc. Int. Conf. Learn. Representations*, 2025, pp. 1–25.
- [32] Y. Wang, X. Meng, J. Liang, Y. Wang, Q. Liu, and D. Zhao, "Hawkeye: Training video-text LLMs for grounding text in videos," 2024, *arXiv:2403.10228*.
- [33] X. Zeng et al., "Timesuite: Improving MLLMs for long video understanding via grounded tuning," in *Proc. Int. Conf. Learn. Representations*, 2025, pp. 1–25.
- [34] H. Wang et al., "Grounded-videoLLM: Sharpening fine-grained temporal grounding in video large language models," 2024, *arXiv:2410.03290*.
- [35] Y. Liu, K. Q. Lin, C. W. Chen, and M. Z. Shou, "VideoMind: A chain-of-loRA agent for long video reasoning," 2025, *arXiv:2503.13444*.
- [36] C. Li et al., "Multimodal foundation models: From specialists to general-purpose assistants," 2023, *arXiv:2309.10020*.
- [37] Y. Zhu et al., "A comprehensive study of deep video action recognition," 2020, *arXiv:2012.06567*.
- [38] Z. Xing et al., "A survey on video diffusion models," *ACM Comput. Surv.*, vol. 57, no. 2, pp. 1–42, 2024.
- [39] J. Zhang, J. Huang, S. Jin, and S. Lu, "Vision-language models for vision tasks: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 46, no. 8, pp. 5625–5644, Aug. 2024.
- [40] X. Liu, X. Nie, Z. Tan, J. Guo, and Y. Yin, "A survey on natural language video localization," 2021, *arXiv:2104.00234*.
- [41] M. Liu, L. Nie, Y. Wang, M. Wang, and Y. Rui, "A survey on video moment localization," *ACM Comput. Surv.*, vol. 55, no. 9, pp. 1–37, 2023.
- [42] Y. Yang, Z. Li, and G. Zeng, "A survey of temporal activity localization via language in untrimmed videos," in *Proc. IEEE Int. Carnahan Conf. Secur. Technol.*, 2020, pp. 596–601.
- [43] X. Lan, Y. Yuan, X. Wang, Z. Wang, and W. Zhu, "A survey on temporal sentence grounding in videos," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 19, no. 2, pp. 1–33, 2023.
- [44] H. Zhang, A. Sun, W. Jing, and J. T. Zhou, "Temporal sentence grounding in videos: A survey and future directions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 8, pp. 10443–10465, Aug. 2023.
- [45] S. Zhang, H. Peng, J. Fu, and J. Luo, "Learning 2D temporal adjacent networks for moment localization with natural language," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 12870–12877.
- [46] K. Q. Lin et al., "Univtg: Towards unified video-language temporal grounding," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 2794–2804.
- [47] D. Liu et al., "Context-aware biaffine localizing network for temporal sentence grounding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 11235–11244.
- [48] S. Xiao et al., "Boundary proposal network for two-stage natural language video localization," in *Proc. AAAI Conf. Artif. Intell.*, 2021, pp. 2986–2994.
- [49] M. Liu, X. Wang, L. Nie, Q. Tian, B. Chen, and T.-S. Chua, "Cross-modal moment localization in videos," in *Proc. ACM Int. Conf. Multimedia*, 2018, pp. 843–851.
- [50] Y. Liu, Z. Ma, Z. Qi, Y. Wu, Y. Shan, and C. W. Chen, "ET bench: Towards open-ended event-level video-language understanding," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2024, pp. 32076–32110.
- [51] L. Zhou, Y. Zhou, J. J. Corso, R. Socher, and C. Xiong, "End-to-end dense video captioning with masked transformer," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 8739–8748.
- [52] V. Iashin and E. Rahtu, "Multi-modal dense video captioning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, 2020, pp. 958–959.
- [53] J. Lei, T. L. Berg, and M. Bansal, "Detecting moments and highlights in videos via natural language queries," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 11846–11858.
- [54] G. Chen et al., "CG-bench: Clue-grounded question answering benchmark for long video understanding," 2024, *arXiv:2412.12075*.
- [55] H. Liu, X. Ma, C. Zhong, Y. Zhang, and W. Lin, "Timecraft: Navigate weakly-supervised temporal grounded video question answering via BI-directional reasoning," in *Proc. Eur. Conf. Comput. Vis.*, 2025, pp. 92–107.
- [56] F. Ma, L. Zhu, and Y. Yang, "Weakly supervised moment localization with decoupled consistent concept prediction," *Int. J. Comput. Vis.*, vol. 130, no. 5, pp. 1244–1258, 2022.
- [57] K. Wang, H. Liu, L. Jie, Z. Li, Y. Hu, and L. Nie, "Explicit granularity and implicit scale correspondence learning for point-supervised video moment localization," in *Proc. ACM Int. Conf. Multimedia*, 2024, pp. 9214–9223.
- [58] H. Li et al., "D3G: Exploring gaussian prior for temporal sentence grounding with glance annotation," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 13734–13746.
- [59] S. Chen, W. Jiang, W. Liu, and Y.-G. Jiang, "Learning modality interaction for temporal sentence localization and event captioning in videos," in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 333–351.
- [60] D. Cao, Y. Zeng, X. Wei, L. Nie, R. Hong, and Z. Qin, "Adversarial video moment retrieval by jointly modeling ranking and localization," in *Proc. ACM Int. Conf. Multimedia*, 2020, pp. 898–906.
- [61] K. Ning, L. Xie, J. Liu, F. Wu, and Q. Tian, "Interaction-integrated network for natural language moment localization," *IEEE Trans. Image Process.*, vol. 30, pp. 2538–2548, 2021.
- [62] W. Yang, T. Zhang, Y. Zhang, and F. Wu, "Local correspondence network for weakly supervised temporal sentence grounding," *IEEE Trans. Image Process.*, vol. 30, pp. 3252–3262, 2021.
- [63] D. Liu, X. Qu, J. Dong, and P. Zhou, "Adaptive proposal generation network for temporal sentence localization in videos," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2021, pp. 9292–9301.
- [64] J. Wang, L. Ma, and W. Jiang, "Temporally grounding language queries in videos by contextual boundary-aware prediction," in *Proc. AAAI Conf. Artif. Intell.*, 2020, pp. 12168–12175.
- [65] B. Lin et al., "Video-LLaVA: Learning united visual representation by alignment before projection," 2023, *arXiv:2311.10122*.
- [66] K. Li et al., "VideoChat: Chat-centric video understanding," 2023, *arXiv:2305.06355*.
- [67] Y. Fang et al., "EVA: Exploring the limits of masked visual representation learning at scale," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 19358–19369.
- [68] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 8748–8763.
- [69] C. Feichtenhofer et al., "Masked autoencoders as spatiotemporal learners," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 35946–35958.
- [70] Z. Tong, Y. Song, J. Wang, and L. Wang, "Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 10078–10093.
- [71] L. Wang et al., "Videomae v2: Scaling video masked autoencoders with dual masking," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 14549–14560.
- [72] W. Dai et al., "InstructBLIP: Towards general-purpose vision-language models with instruction tuning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 49250–49267.
- [73] S. Azad, V. Vineet, and Y. S. Rawat, "HierarQ: Task-aware hierarchical Q-former for enhanced video understanding," 2025, *arXiv:2503.08585*.
- [74] J.-B. Alayrac et al., "Flamingo: A visual language model for few-shot learning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 23716–23736.
- [75] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 19730–19742.
- [76] Q. Ye et al., "mPLUG-OWL: Modularization empowers large language models with multimodality," 2023, *arXiv:2304.14178*.
- [77] S. Huang et al., "Language is not all you need: Aligning perception with language models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 72096–72109.

- [78] H. Liu, C. Li, Y. Li, and Y. J. Lee, "Improved baselines with visual instruction tuning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 26296–26306.
- [79] F. Li et al., "LlaVA-next-interleave: Tackling multi-image, video, and 3D in large multimodal models," 2024, *arXiv:2407.07895*.
- [80] J. Ye et al., "mPLUG-OWL3: Towards long image-sequence understanding in multi-modal large language models," 2024, *arXiv:2408.04840*.
- [81] K. Chen et al., "EVLm: An efficient vision-language model for visual understanding," 2024, *arXiv:2407.14177*.
- [82] M. Shi et al., "Slow-fast architecture for video multi-modal large language models," 2025, *arXiv:2504.01328*.
- [83] A. Miech, D. Zhukov, J.-B. Alayrac, M. Tapaswi, I. Laptev, and J. Sivic, "Howto100m: Learning a text-video embedding by watching hundred million narrated video clips," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2019, pp. 2630–2640.
- [84] M. Bain, A. Nagrani, G. Varol, and A. Zisserman, "Frozen in time: A joint video and image encoder for end-to-end retrieval," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2021, pp. 1728–1738.
- [85] C. Schuhmann et al., "Laion-5b: An open large-scale dataset for training next generation image-text models," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2022, pp. 25278–25294.
- [86] L. Li et al., "M³ it: A large-scale dataset towards multi-modal multilingual instruction tuning," 2023, *arXiv:2306.04387*.
- [87] Z. Yin et al., "LAMM: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 26650–26685.
- [88] T. Gao et al., "Cantor: Inspiring multimodal chain-of-thought of MLLM," in *Proc. ACM Int. Conf. Multimedia*, 2024, pp. 9096–9105.
- [89] J. Wu et al., "Visual prompting in multimodal large language models: A survey," 2024, *arXiv:2409.15310*.
- [90] E. J. Hu et al., "LoRA: Low-rank adaptation of large language models," in *Proc. Int. Conf. Learn. Representations*, 2022, pp. 1–13.
- [91] R. Pan et al., "LISA: Layerwise importance sampling for memory-efficient large language model fine-tuning," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2024, pp. 57018–57049.
- [92] S.-Y. Liu et al., "DoRA: Weight-decomposed low-rank adaptation," in *Proc. Int. Conf. Mach. Learn.*, 2024, pp. 32100–32121.
- [93] W. Cai, J. Huang, S. Gong, H. Jin, and Y. Liu, "MLLM as video narrator: Mitigating modality imbalance in video moment retrieval," 2024, *arXiv:2406.17880*.
- [94] S. Di and W. Xie, "Grounded question-answering in long egocentric videos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 12934–12943.
- [95] M. Zheng, X. Cai, Q. Chen, Y. Peng, and Y. Liu, "Training-free video temporal grounding using large-scale pre-trained models," in *Proc. Eur. Conf. Comput. Vis.*, 2025, pp. 20–37.
- [96] H. Chen, X. Wang, H. Chen, Z. Song, J. Jia, and W. Zhu, "Grounding-prompter: Prompting LLM with multimodal information for temporal sentence grounding in long videos," 2023, *arXiv:2312.17117*.
- [97] Y. Xu, Y. Sun, Z. Xie, B. Zhai, and S. Du, "VTG-GPT: Tuning-free zero-shot video temporal grounding with GPT," *Appl. Sci.*, vol. 14, no. 5, 2024, Art. no. 1894.
- [98] H. Chen et al., "Verified: A video corpus moment retrieval benchmark for fine-grained video understanding," 2024, *arXiv:2410.08593*.
- [99] H. Lee, S. Hong, M. Sung, and J. Choi, "Infusing environmental captions for long-form video language grounding," 2024, *arXiv:2408.02336*.
- [100] D. Paul, M. R. Parvez, N. Mohammed, and S. Rahman, "Videolights: Feature refinement and cross-task alignment transformer for joint video highlight detection and moment retrieval," 2024, *arXiv:2412.01558*.
- [101] Y. Sun, Y. Xu, Z. Xie, Y. Shu, and S. Du, "GPTSee: Enhancing moment retrieval and highlight detection via description-based similarity features," *IEEE Signal Process. Lett.*, vol. 31, pp. 521–525, 2024.
- [102] W. Liu et al., "Context-enhanced video moment retrieval with large language models," 2024, *arXiv:2405.12540*.
- [103] P. Bao, C. Kong, Z. Shao, B. P. Ng, M. H. Er, and A. C. Kot, "Vid-morp: Video moment retrieval pretraining from unlabeled videos in the wild," 2024, *arXiv:2412.00811*.
- [104] Y. Xu et al., "Zero-shot video moment retrieval via off-the-shelf multimodal large language models," in *Proc. AAAI Conf. Artif. Intell.*, 2025, pp. 8978–8986.
- [105] S. Yu, J. Cho, P. Yadav, and M. Bansal, "Self-chained image-language model for video localization and question answering," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2023, pp. 76749–76771.
- [106] K. Ma et al., "LLaVilo: Boosting video moment retrieval via adapter-based multimodal modeling," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2023, pp. 2798–2803.
- [107] Y. Liu, H. Hou, F. Ma, S. Ni, and F. R. Yu, "MLLM-TA: Leveraging multimodal large language models for precise temporal video grounding," *IEEE Signal Process. Lett.*, vol. 32, pp. 281–285, 2025.
- [108] W. Lu, J. Li, A. Yu, M.-C. Chang, S. Ji, and M. Xia, "LlaVA-MR: Large language-and-vision assistant for video moment retrieval," 2024, *arXiv:2411.14505*.
- [109] Y. Wu et al., "Number it: Temporal grounding videos like flipping manga," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2025, pp. 13754–13765.
- [110] A. Deng et al., "Seq2time: Sequential knowledge transfer for video LLM temporal grounding," 2024, *arXiv:2411.16932*.
- [111] Y. Guo et al., "VTG-LLM: Integrating timestamp knowledge into video LLMs for enhanced video temporal grounding," in *Proc. AAAI Conf. Artif. Intell.*, 2025, pp. 3302–3310.
- [112] Z. Li et al., "GroundingGPT: Language enhanced multi-modal grounding model," in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2024, pp. 6657–6678.
- [113] T. Hannan, M. M. Islam, J. Gu, T. Seidl, and G. Bertasius, "ReVisionLLM: Recursive vision-language model for temporal grounding in hour-long videos," 2024, *arXiv:2411.14901*.
- [114] D.-A. Huang et al., "LITA: Language instructed temporal-localization assistant," in *Proc. Eur. Conf. Comput. Vis.*, 2025, pp. 202–218.
- [115] X. Wang et al., "Temporal grounding with time refining video LLM," *arXiv:2412.09601*, 2024.
- [116] H. Li et al., "LlaVA-ST: A multimodal large language model for fine-grained spatial-temporal understanding," 2025, *arXiv:2501.08282*.
- [117] S. Chen, X. Lan, Y. Yuan, Z. Jie, and L. Ma, "Timemarker: A versatile video-LLM for long and short video understanding with superior temporal localization ability," 2024, *arXiv:2411.18211*.
- [118] Q. Chen, S. Di, and W. Xie, "Grounded multi-hop videoqa in long-form egocentric videos," in *Proc. AAAI Conf. Artif. Intell.*, 2025, pp. 2159–2167.
- [119] M. Nie et al., "Slowfocus: Enhancing fine-grained temporal understanding in video LLM," in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2024, pp. 81808–81835.
- [120] B. Meinardus, A. Batra, A. Rohrbach, and M. Rohrbach, "The surprising effectiveness of multimodal large language models for video moment retrieval," 2024, *arXiv:2406.18113*.
- [121] F. J. Fateh, U. Ahmed, H. Khan, M. Z. Zia, and Q.-H. Tran, "Video LLMs for temporal reasoning in long videos," 2024, *arXiv:2412.02930*.
- [122] Y. Wang et al., "Efficient temporal extrapolation of multimodal large language models with temporal grounding bridge," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2024, pp. 9972–9987.
- [123] X. Li, B. Wang, G. Shi, C. Feng, and J. Teng, "Mitigating the discrepancy between video and text temporal sequences: A time-perception enhanced video grounding method for LLM," in *Proc. Int. Conf. Comput. Linguistics*, 2025, pp. 9804–9813.
- [124] Z. Yan et al., "Task preference optimization: Improving multimodal large language models with vision task alignment," 2024, *arXiv:2412.19326*.
- [125] J. Wang et al., "SpacevLLM: Endowing multimodal large language model with spatio-temporal video grounding capability," 2025, *arXiv:2503.13983*.
- [126] Z. Pang, M. Otani, and Y. Nakashima, "Measure twice, cut once: Grasping video structures and event semantics with LLMs for video temporal localization," 2025, *arXiv:2503.09027*.
- [127] Y. Wang et al., "Time-R1: Post-training large vision language model for temporal video grounding," 2025, *arXiv:2503.13377*.
- [128] H. Zhao, G.-P. Ji, R. Yan, H. Xiong, and Z. Li, "Videoexpert: Augmented LLM for temporal-sensitive video understanding," 2025, *arXiv:2504.07519*.
- [129] X. Li et al., "Videochat-R1: Enhancing spatio-temporal perception via reinforcement fine-tuning," 2025, *arXiv:2504.06958*.
- [130] F. Luo et al., "MUSEG: Reinforcing video temporal understanding via timestamp-aware multi-segment grounding," 2025, *arXiv:2505.20715*.
- [131] H. Touvron et al., "LLAMA: Open and efficient foundation language models," 2023, *arXiv:2302.13971*.
- [132] K. Grauman et al., "Ego4D: Around the world in 3,000 hours of egocentric video," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 18995–19012.
- [133] OpenAI, "GPT-4O," 2024. [Online]. Available: <https://openai.com/index/hello-gpt-4o/>

- [134] G. Team et al., “Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context,” 2024, *arXiv:2403.05530*.
- [135] N. Reimers and I. Gurevych, “Sentence-BERT: Sentence embeddings using Siamese BERT-networks,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2019, pp. 3982–3992.
- [136] J. Chen et al., “MiniGPT-V2: Large language model as a unified interface for vision-language multi-task learning,” 2023, *arXiv:2310.09478*.
- [137] A. Yang et al., “Baichuan 2: Open large-scale language models,” 2023, *arXiv:2309.10305*.
- [138] K. C. Fraser and S. Kiritchenko, “Examining gender and racial bias in large vision-language models using a novel dataset of parallel images,” in *Proc. 13th Conf. Eur. Chapter Assoc. Comput. Linguistics*, 2024, pp. 690–713.
- [139] Z. Shao et al., “Deepseekmath: Pushing the limits of mathematical reasoning in open language models,” 2024, *arXiv:2402.03300*.
- [140] R. Zellers et al., “Merlot: Multimodal neural script knowledge models,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 23634–23651.
- [141] Y. Wang et al., “InternVid: A large-scale video-text dataset for multimodal understanding and generation,” 2023, *arXiv:2307.06942*.
- [142] R. Zellers et al., “Merlot reserve: Neural script knowledge through vision and language and sound,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 16375–16387.
- [143] Y. Zhai et al., “Investigating the catastrophic forgetting in multimodal large language models,” 2023, *arXiv:2309.10313*.
- [144] M. Liu, X. Wang, L. Nie, X. He, B. Chen, and T.-S. Chua, “Attentive moment retrieval in videos,” in *Proc. Int. ACM SIGIR Conf. Res. Develop. Inf. Retrieval*, 2018, pp. 15–24.
- [145] Y.-W. Chen, Y.-H. Tsai, and M.-H. Yang, “End-to-end multi-modal video temporal grounding,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, 2021, pp. 28442–28453.
- [146] Y. Lu, R. Quan, L. Zhu, and Y. Yang, “Zero-shot video grounding with pseudo query lookup and verification,” *IEEE Trans. Image Process.*, vol. 33, pp. 1643–1654, 2024.
- [147] S. Chen et al., “Beats: Audio pre-training with acoustic tokenizers,” 2022, *arXiv:2212.09058*.
- [148] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics*, 2019, pp. 4171–4186.
- [149] J. Gao, C. Sun, Z. Yang, and R. Nevatia, “Tall: Temporal activity localization via language query,” in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5267–5275.
- [150] L. Zhou, C. Xu, and J. Corso, “Towards automatic learning of procedures from web instructional videos,” in *Proc. AAAI Conf. Artif. Intell.*, 2018, pp. 7590–7598.
- [151] G. A. Sigurdsson, G. Varol, X. Wang, A. Farhadi, I. Laptev, and A. Gupta, “Hollywood in homes: Crowdsourcing data collection for activity understanding,” in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 510–526.
- [152] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Nieves, “ActivityNet: A large-scale video benchmark for human activity understanding,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 961–970.
- [153] J. Xiao, X. Shang, A. Yao, and T.-S. Chua, “NExT-QA: Next phase of question-answering to explaining temporal actions,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 9777–9786.
- [154] R. Hu, H. Xu, M. Rohrbach, J. Feng, K. Saenko, and T. Darrell, “Natural language object retrieval,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4555–4564.
- [155] S. Fujita, T. Hirao, H. Kamigaito, M. Okumura, and M. Nagata, “SODA: Story oriented dense video captioning evaluation framework,” in *Proc. Eur. Conf. Comput. Vis.*, 2020, pp. 517–531.
- [156] S. Banerjee and A. Lavie, “Meteor: An automatic metric for MT evaluation with improved correlation with human judgments,” in *Proc. Annu. Meeting Assoc. Comput. Linguistics Workshop*, 2005, pp. 65–72.
- [157] R. Vedantam, C. Lawrence Zitnick, and D. Parikh, “CIDEr: Consensus-based image description evaluation,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 4566–4575.
- [158] W. Liu, T. Mei, Y. Zhang, C. Che, and J. Luo, “Multi-task deep visual-semantic embedding for video thumbnail selection,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 3707–3715.
- [159] Y. Zhong, W. Ji, J. Xiao, Y. Li, W. Deng, and T.-S. Chua, “Video question answering: Datasets, algorithms and challenges,” in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2022, pp. 6439–6455.
- [160] X. Wang, Q. Si, S. Zhu, J. Wu, L. Cao, and L. Nie, “AdaReTaKe: Adaptive redundancy reduction to perceive longer for video-language understanding,” in *Proc. Annu. Meeting Assoc. Comput. Linguistics*, 2025, pp. 5417–5432.
- [161] Y. Li, C. Wang, and J. Jia, “LLaMA-VID: An image is worth 2 tokens in large language models,” in *Proc. Eur. Conf. Comput. Vis.*, 2024, pp. 323–340.
- [162] I. Viertola, V. Iashin, and E. Rahtu, “Temporally aligned audio for video with autoregression,” in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2025, pp. 1–5.
- [163] G. Yariv, I. Gat, S. Benaim, L. Wolf, I. Schwartz, and Y. Adi, “Diverse and aligned audio-to-video generation via text-to-video model adaptation,” in *Proc. AAAI Conf. Artif. Intell.*, 2024, pp. 6639–6647.
- [164] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 28492–28518.
- [165] H. Xu, P. Peng, G. Tan, Y. Chang, Y. Zhao, and Y. Tian, “Temporal triplane transformers as occupancy world models,” 2025, *arXiv:2503.07338*.
- [166] M. Soldan et al., “Mad: A scalable dataset for language grounding in videos from movie audio descriptions,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2022, pp. 5026–5035.
- [167] A. Zala et al., “Hierarchical video-moment retrieval and step-captioning,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2023, pp. 23056–23065.
- [168] C. Tan et al., “Synopground: A large-scale dataset for multi-paragraph video grounding from TV dramas and synopses,” in *Proc. ACM Int. Conf. Multimedia*, 2024, pp. 8383–8392.
- [169] W. Wang et al., “Lvbench: An extreme long video understanding benchmark,” 2024, *arXiv:2406.08035*.



Jianlong Wu (Member, IEEE) received the BEng degree from the Huazhong University of Science and Technology, China, in 2014, and the PhD degree from Peking University, China, in 2019. He is currently a professor with the Harbin Institute of Technology (Shenzhen), China. He has published more than 40 papers in top journals and conferences, such as *IEEE Transactions on Pattern Analysis and Machine Intelligence*, ICML, NeurIPS, and ICCV. His research interests include computer vision and multi-modal learning. He received many awards, such as the Outstanding Reviewer of ICML 2020 and the Best Student Paper of SIGIR 2021. He serves as the Area Chair for ICML, NeurIPS, CVPR and ACM MM, and reviewer for many top journals, such as *IEEE Transactions on Pattern Analysis and Machine Intelligence* and *International Journal of Computer Vision*.



Wei Liu received the BEng degree from NanKai University, Tianjin, China, in 2024. He is currently working toward the MS degree in computer science and technology with the Harbin Institute of Technology (Shenzhen). He won the champion of the 2025 CVPR VidLLMs Workshop Challenge Track (CoVR). His research interests include computer vision and video understanding.



Ye Liu received the BS degree in geographical information science and the BE degree in computer science from Wuhan University, in 2020. He is currently working toward the PhD degree with the Hong Kong Polytechnic University. He was a research assistant with the State University of New York at Buffalo and The Chinese University of Hong Kong, Shenzhen, in 2019. He served as a visiting scholar with National University of Singapore, an associate with Harvard John A. Paulson School of Engineering and Applied Sciences (SEAS), and a research intern with Tencent

ARC Lab. His research interests include computer vision, multimedia, and urban computing, with a particular focus on foundation models and vision-language understanding.



and INS.

Meng Liu (Member, IEEE) received the PhD degree from the Shandong University, China, in 2019. She is currently a professor with the School of Computer Science and Technology, Shandong Jianzhu University. Her research interests include multimedia computing and information retrieval. Various parts of her work have been published in top forums and journals, such as SIGIR, MM, and *IEEE Transactions on Image Processing*. She has served as a reviewer and a Subreviewer for various conferences and journals, such as MMM 2018, ACM MM 2018/2019, JVCI,



CVPR 2023. He is an associate editor of the *International Journal of Computer Vision and Optimization Methods and Software*. He is a fellow of IAPR, AAAI and CSIG.

Zhouchen Lin (Fellow, IEEE) received the PhD degree in applied mathematics from Peking University, in 2000. He is currently a Boya Special professor with the State Key Laboratory of General Artificial Intelligence, School of Intelligence Science and Technology, Peking University. His research interests include computer vision, image processing, machine learning, pattern recognition, and numerical optimization. He was the program co-chair of International Conference on Pattern Recognition (ICPR) 2022 and the senior area chair of ICML 2022, NeurIPS 2022 and



Scholar citations. He is an AE of *IEEE Transactions on Image Processing*, *IEEE Transactions on Knowledge and Data Engineering*, *IEEE Transactions on Multimedia*, *IEEE Transactions on Circuits and Systems for Video Technology*, ACM ToMM, and *Information Science*. Meanwhile, he serves as the chair of ICMR 2025, ICME 2025, and ACM MM 2027. He is a member of ICME steering committee. He has received many awards over the past four years, like SIGMM rising star in 2020, MIT TR35 China 2020, SIGIR best student paper in 2021, IEEE AI's 10 to Watch in 2022, ACM MM Best paper award in 2022, and the National Youth Science and Technology Award in 2024.

Liqiang Nie (Senior Member, IEEE) received the BEng degree from Xi'an Jiaotong University, and the PhD degree from the National University of Singapore (NUS), respectively. After Ph.D., he continued his research in NUS as a research fellow for three years. He is currently the dean with the Department of Computer Science and Technology, Harbin Institute of Technology (Shenzhen). His research interests include primarily in multimedia computing and information retrieval. He has co-authored more than 150 papers and 5 books, with 34,000 plus Google



Columbia, MO, USA, from 1996 to 2003. He was the Allen Henry Endow Chair Professor with the Florida Institute of Technology, Melbourne, FL, USA, from 2003 to 2007. He is currently a professor of Computer Science and Engineering with the University at Buffalo, The State University of New York, Buffalo, NY, USA.

Chang Wen Chen (Life Fellow, IEEE) received the BS degree from the University of Science and Technology of China, Hefei, China, in 1983, the MSEE degree from the University of Southern California, Los Angeles, CA, USA, in 1986, and the PhD degree from the University of Illinois at Urbana-Champaign, Champaign, IL, USA, in 1992. He was with the Faculty of Electrical and Computer Engineering, University of Rochester, Rochester, NY, USA, from 1992 to 1996, and the Faculty of Electrical and Computer Engineering, University of Missouri-Columbia,