

Prior Knowledge Integration via LLM Encoding and Pseudo Event Regulation for Video Moment Retrieval

Anonymous Author(s)*

ABSTRACT

In this paper, we investigate the feasibility of leveraging large language models (LLMs) for integrating general knowledge and incorporating pseudo-events as priors for temporal content distribution in video moment retrieval (VMR) models. The motivation behind this study arises from the limitations of using LLMs as decoders for generating discrete textual descriptions, which hinders their direct application to continuous outputs like salience scores and inter-frame embeddings that capture inter-frame relations. To overcome these limitations, we propose utilizing LLM encoders instead of decoders. Through a feasibility study, we demonstrate that LLM encoders effectively refine inter-concept relations in multimodal embeddings, even without being trained on textual embeddings. We also show that the refinement capability of LLM encoders can be transferred to other embeddings, such as BLIP and T5, as long as these embeddings exhibit similar inter-concept similarity patterns to CLIP embeddings. We present a general framework for integrating LLM encoders into existing VMR architectures, specifically within the fusion module. The LLM encoder’s ability to refine concept relation can help the model to achieve a balanced understanding of the foreground concepts (e.g., persons, faces) and background concepts (e.g., street, mountains) rather focusing only on the visually dominant foreground concepts. Additionally, we introduce the concept of pseudo-events, obtained through event detection techniques, to guide the prediction of moments within event boundaries instead of crossing them, which can effectively avoid the distractions from adjacent moments. The integration of semantic refinement using LLM encoders and pseudo-event regulation is designed as plug-in components that can be incorporated into existing VMR methods within the general framework. Through experimental validation, we demonstrate the effectiveness of our proposed methods by achieving state-of-the-art performance in VMR. The source code can be accessed at https://github.com/open_upon_acceptance.

CCS CONCEPTS

• Computing methodologies → Visual content-based indexing and retrieval.

KEYWORDS

Video Moment Retrieval, Highlight Detection, LLMs

1 INTRODUCTION

The rapid expansion of video content, driven by advancements in digital platforms and devices, has elevated video to become one of the most captivating and information-rich media formats today. However, this surge has also presented a significant challenge in efficiently navigating through vast amounts of video content to locate specific user-requested moments or highlights. In response to this challenge, research efforts have progressed from traditional approaches of moment retrieval (MR) and highlight detection (HD) to

more advanced methods, including Moment-DETR [29], UMT [43], and UniVTG [36]. These cutting-edge techniques have pushed the boundaries in the field of video understanding, enabling more effective analysis and extraction of meaningful insights from videos.

An emerging trend observed in these advancements is the incorporation of prior knowledge into the learning process to enhance the semantic context and improve representation learning. This effort includes the expansion of datasets from a scale of hundreds [61] to tens of thousands [29], which improves the generality of resulting representations across broader domain. Another endeavor is the introduction of self-supervised pretraining such as UniVTG [36] which has been conducted by integrating existing MR/HD datasets for in-domain knowledge incorporation and QVHighlights [29] which has been pretrained on Youtube subtitles for cross-domain knowledge embedding. While these have broken new ground for prior integration, the question of whether we can innovate beyond still remains open.

In a post-ChatGPT era, the inclination to utilize large language models (LLMs) arises naturally due to their success across various tasks. However, employing LLMs in MR/HD tasks proves to be challenging since our focus lies primarily on capturing fine-grained inter-frame salience, whereas LLMs excel in high-level semantic comprehension. In simpler terms, LLMs perform well in tasks such as captioning [56] or grounding [27], effectively describing the content of a video or image as a whole, but they lack the ability to compare the degree of semantic salience among individual frames. From a technical standpoint, most applications utilize LLMs as decoders, where visual representations are first converted into LLM-compatible tokens using models like Q-Former [32]. These transformed representations are then integrated into the context for generating outputs. Unfortunately, neither the Q-Former nor the LLM models have been trained with frame-level salience information during this process. Furthermore, the LLM decoders output textual tokens which are discrete and determinate and thus less compatible to the continuous and comparative salience scores. This also makes the customization of the decoders or integrating MLPs for fine-grained decision making less feasible.

In this paper, we tackle these limitations by utilizing LLM encoders instead of decoders. Through a feasibility study, we demonstrate that LLM encoders effectively refine inter-concept relations in multimodal embeddings, even when not trained on textual embeddings. We confirm that the refinement ability of LLM encoders can be transferred to other embeddings, such as BLIP [33] and T5 [57], as long as the inter-concept similarities exhibited by these embeddings demonstrate similar patterns to CLIP embeddings (which serve as the original input for most LLMs). Based on these findings, we propose a general framework for applying LLM encoders in VMR. We showcase that LLM encoders can be inserted into the fusion module of existing VMR architectures. By doing this, the LLM encoder’s ability for inter-concept refinement can help the model

to have a comprehensive understanding of foreground concepts (e.g., *persons, faces*) and background concepts (e.g., *street, buildings*), to avoid the learner to be misled by the visually dominant foreground concepts. This is especially important when those concepts are scattered over consecutive frames and the combination of their semantics cannot be identified without an inter-frame relation modeling. Furthermore, we introduce the use of pseudo-events, obtained through event detection techniques (e.g., [63]), as priors for the content distribution of videos. This approach guides the predicted moments to align with events rather than crossing unreasonable event boundaries, which can effectively avoid distractions from adjacent moments. The integrations of semantic refinement using LLM encoders and pseudo-event regulation are designed as plug-in components that can be incorporated into existing VMR methods constructed within the general framework. Through our experiments, we validate that these components can enhance the performance of five VMR frameworks: Moment-Detr [29], Uni-VTG [36], QD-DETR [50], CG-DETR [49], and EaTR [23].

2 RELATED WORK

2.1 Moment Retrieval and Highlight Detection

Moment Retrieval (MR) and Highlight Detection (HD) have become pivotal in navigating the growing expanse of video content. MR focuses on localizing video moments pertinent to textual descriptions, employing cross-modal interactions [45, 49, 50, 85, 87, 93] and temporal relation context [20, 89]. These tasks focus on localizing user-desired moments and scoring clip-wise correspondence to queries, respectively. Moment retrieval, aimed at retrieving user-specific video segments, has evolved significantly [10, 17, 20, 37, 65]. Traditional methods in this domain are categorized into proposal-based and proposal-free approaches. Proposal-Based Methods utilize predefined proposals like sliding windows [18, 20, 41, 85, 90] or temporal anchors [8, 39, 85, 86, 91], and in some cases, generate proposals [38, 62, 76, 79, 89]. The essence of these methods lies in matching these candidates with the text query. In contrast, proposal-free methods bypass the use of predefined candidates. Instead, they focus on encoding multimodal knowledge and directly predict temporal spans using regression heads, making the process more streamlined and potentially more accurate.

Highlight detection, on the other hand, concentrates on scoring the importance of each video clip, whether based solely on visual [3, 60, 77, 80] or combined visual-audio inputs [3, 19, 21, 77]. The methods here vary in their approach to label granularity, with supervised [19, 68, 80], weakly supervised [5, 53, 77][6, 49, 68], and unsupervised methods [3, 25, 47, 60] all contributing to the field. The introduction of the QVHighlights dataset [29] marked a significant shift in this domain, prompting a combined consideration of these problems. Emerging approaches post-QVHighlights adopt either DETR or regression-based frameworks. For instance, UMT [43] explores additional audio modalities, while QD-DETR [50] and CG-DETR [49] innovate upon the DETR architecture. Other studies [36, 82] underscore the importance of pretraining.

2.2 Vision-Language Models

Inspired by the success of ChatGPT, significant efforts have been dedicated to developing visual-language models (VLMs) based on

LLMs [11, 12, 51, 52, 70–72]. One commonly employed approach involves encoding images into tokens that are compatible with LLM inputs. The pretrained LLMs then serve as decoders for generating textual descriptions. A well-known framework is BLIP-2 [32], which employs QFormer as the encoder and LLMs like FlanT5 [13] as the decoder. Several successful examples have adopted this framework. For instance, MiniGPT4 [94] replaces the LLM with a more advanced model called LLaMA [72]. InstructBLIP [14] fine-tunes the model using high-quality instructions, while LLaV [40] explores CLIP’s open-set visual encoder connected to Vicuna’s linguistic decoder and performs end-to-end fine-tuning on the generated instructional vision-language data. Similar frameworks have been extended to video understanding, resulting in models like VideoChat [34], VideoChatGPT [46], and LLaMA-VID [35].

Other variants that may not strictly stick to the framework include GILL [26], Emu [69], mPLUG [31], CogVLM [75] and MiniGPT-5 [92] for visual question answering, and Vision-LLM [74], Kosmos-2 [55], Qwen-VL [4], MiniGPT-v2 [9], GPT4-Vision [1], and GeminiProVision [70] that support multimodal inputs and outputs. However, the role of LLMs in these frameworks is still a decoder for discrete outputs (mainly textual descriptions/answers), which makes them limited to continuous outputs such as the salience scores or inter-frame correlations. In this paper, we explore the feasibility of employing LLMs as encoders for semantic relation refinement, based on which the outputs are still continuous embeddings and thus open the opportunity for fine-grained information processing or decoding. This particular direction remains under-explored, with only one existing work [54] in the literature that has attempted a similar approach for image understanding tasks. This paper distinguishes from [54] in two aspects: we conduct a feasibility study to investigate the rationale behind this approach, and we focus on the integration of prior knowledge in the LLM.

3 LEVERAGING LLM ENCODERS FOR INTER-CONCEPT RELATION REFINEMENT

Before introducing our VMR framework, let us examine the viability of utilizing LLM encoders to refine inter-concept relations. LLM encoders serve as natural inter-concept relation refiners since they operate as Transformers, taking concepts (represented as token embeddings) as input and generating refined embeddings as output. This becomes particularly evident when contrastive losses are employed, as they encourage similar concepts to move closer in the embedding space while pushing dissimilar ones further apart. However, this approach appears to be effective only for textual embeddings and is tightly coupled with the specific embeddings used by the LLM (e.g., CLIP [56], BLIP [33], T5 [57]). In this feasibility study, our aim is to investigate whether this approach can be applied to multimodal embeddings and delve into the underlying reasoning. Additionally, we will explore the possibility of reducing computational costs by utilizing only a subset of internal layers from an LLM encoder instead of loading the entire model.

3.1 LLM Encoders are Relation Refiners

We conducted an experiment to validate the hypothesis that an LLM Encoder can be utilized for relation refinement. The experiment commenced by requiring GPT-4 to recommend 1,000 triplets,

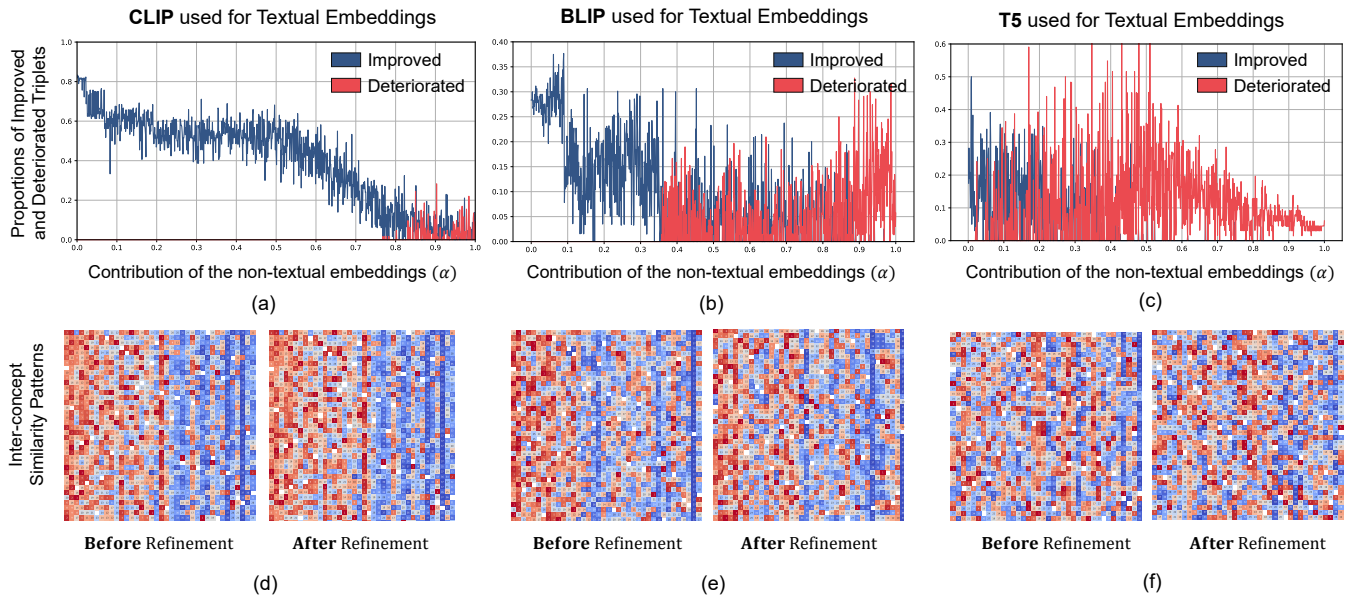


Figure 1: The proportions of improved and deteriorated triplets after the refinement (a–c), and the inter-concept similarity matrices of the concept embeddings before and after the refinement (d–f): in (a) and (d), CLIP is used as the textual embeddings, while BLIP is used for (b) and (e), and T5 is used for (c) and (f).

wherein each triplet (e.g., *sock, shoe, galaxy*) comprises two “paired” concepts (e.g., *sock* and *shoe*) that possess a strong semantic relation, while the third concept (e.g., *galaxy*) is significantly less related to the other two. We label a triplet as “unreasonable” if the similarity between the paired concepts is lower than that of any unpaired concepts; otherwise, it is deemed “reasonable.” We input CLIP embeddings of these concepts into the LLaMA-2 [72] encoder and obtain refined embeddings. To assess the extent of refinement in the embeddings, we measure the percentage of unreasonable triplets in the CLIP space that become reasonable in the refined space. The result shows that 83.33% (300/360) of the initially unreasonable triplets have been refined to a reasonable state, while none of the reasonable triplets were transformed into unreasonable ones.

3.2 LLM Encoders for Multimodal Embeddings

While the LLM encoder’s capability to refine semantic relations is expected, our primary focus lies in investigating whether this ability can be transferred to multimodal embeddings. In most applications, multimodal embeddings are generated by combining non-textual embeddings with textual ones using methods such as weighted summation or more advanced cross-attention mechanisms. These fusion approaches aim to align the non-textual embeddings with the textual ones, forming the foundation for stacking and fine-tuning task-specific decision-making layers. In our subsequent experiment, we simulate the fusion and alignment process by fusing the textual embeddings with randomly generated embeddings.

Dose the LLM encoders work for fused embeddings? Let \mathbf{c} represent a textual concept embedding obtained through CLIP, and let \mathbf{c}' denote a concept-dependent randomly generated vector used to simulate a non-textual embedding. We combine these two

embeddings using a weighted summation, as

$$\mathbf{c}^* = (1 - \alpha)\mathbf{c} + \alpha\mathbf{c}'. \quad (1)$$

It should be noted that the simulated embedding vector \mathbf{c}' is specific to each concept. This implies that each concept is associated with its own fixed vector, and the vectors corresponding to different concepts are distinct from one another. Due to the random nature of the vectors \mathbf{c}' , the expected similarities between such non-textual embeddings are zero. This allows \mathbf{c}' to simulate real-world scenarios where non-textual embeddings are extracted using specific encoders. In these scenarios, the embeddings are aligned to concepts, but the inter-concept similarities are not necessarily regulated in a reasonable manner. Consequently, the parameter α in Eq. (1) controls the extent to which the textual embeddings are influenced by the randomness and misalignment introduced by the non-textual embeddings.

The results in Fig. 1(a) indicate that when the textual embeddings are predominant (i.e., $\alpha \leq 0.5$), the LLaMA encoder effectively refines the majority of unreasonable pairs without causing any deterioration. Fortunately, this aligns perfectly with the requirements of many multimodal applications, where, during the representation learning phase, the goal is to align the non-textual modalities with the textual one.

What will happen when non-textual embeddings are distorted? To expand the experiment, we introduce a distortion factor by randomly setting elements of a non-textual embedding vector \mathbf{c}' to zero with a probability of p , which results in a distorted vector \mathbf{c}'_p . This simulation reflects real-world multimodal scenarios where the alignment of non-textual embeddings to concepts is not consistently reliable. We can substitute the \mathbf{c}' with \mathbf{c}'_p in Eq. (1) and repeat the experiment. The results shown in Fig. 2 demonstrate

that the distortion of non-textual embeddings would not degrade the refinement when the textual embeddings are dominant.

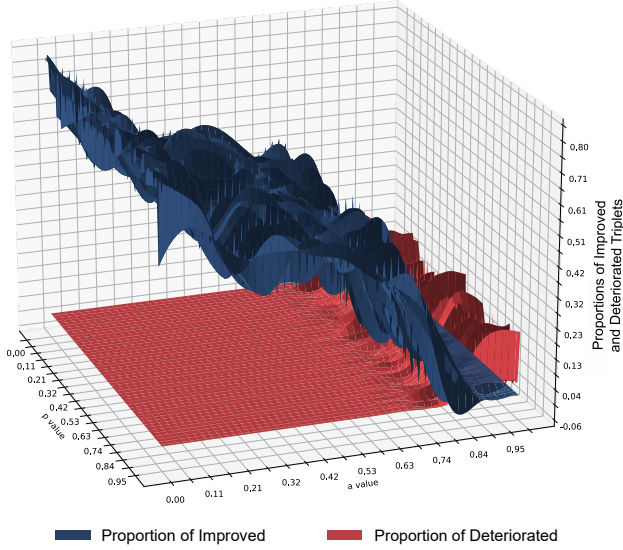


Figure 2: Proportions of Improved and Deteriorated Triplets over the contribution of non-textual embeddings controlled by the α and the degree of alignment between the textual and non-textual embeddings controlled by the distortion probability p .

Through the conducted experiments, we have confirmed the effectiveness of LLM encoders as reliable relation refiners. This can be attributed to the fact that LLM encoders are based on the Transformer architecture. Transforms heavily rely on self-attention mechanisms which are driven by the inter-concept (token) similarities. In other words, as long as the input concepts maintain inter-concept similarities similar to those of the CLIP embeddings, the LLM encoder can refine the concept relations in a manner comparable to using CLIP embeddings as input. In the experiment above, when there are two CLIP concept embeddings c_1 and c_2 , the similarity of their fused embeddings is written

$$\begin{aligned} \text{sim}(c_1^*, c_2^*) &= c_1^* (c_2^*)^\top \\ &= ((1-\alpha)c_1 + \alpha c_1')((1-\alpha)c_2 + \alpha c_2')^\top \\ &= (1-\alpha)^2 c_1 c_2^\top + (1-\alpha)\alpha(c_1 c_2'^\top + c_1' c_2^\top) + \alpha c_1' c_2'^\top. \end{aligned} \quad (2)$$

Obviously, the expectations of concept similarities before and after fusion have a relation of $\mathbb{E}[c_1^* (c_2^*)^\top] \approx (1-\alpha)^2 \mathbb{E}[c_1 (c_2)^\top]$. This is because the expectations of the last two terms in Eq. (2) are zero due to the randomness of the simulated non-textual embeddings. This indicates that the fused concept embedding still maintains similar (even when scaled) inter-concept similarities compared to the original embeddings, which explains why the relations of the fused concepts can also be refined using the LLM encoder.

To validate the discovery, we propose replacing the CLIP embeddings with its variant, BLIP, and introducing a more distant textual embedding model, T5. The results depicted in Fig. 1(b) and Fig. 1(c) indicate improvements when either BLIP or T5 embeddings are

utilized. It is not surprising that by using BLIP embeddings which are more similar to CLIP embeddings, the results demonstrate a better balance between the improved and deteriorated, and the pattern closely resembles that of the CLIP. However, T5 still works because it is also a well-trained model in the text domain, in which we expect a significant overlap in inter-concept similarities with CLIP. In Fig. 1(d–f), we visualize the inter-concept similarities of these three types of embeddings before and after the refinement. It is evident that the patterns exhibit similarity, and the refinement process further enhances this similarity.

3.3 Using A Subset of Layers as the Encoder

Our last experiment for the feasibility is to use a subset of LLaMA encoder as a relation refiner. The motivation is that the layers of a Transformer process data in a similar manner, suggesting that a subset of layers may possess a similar ability to refine relations. If this hypothesis holds true, it would enable a significant reduction in computational requirements, thereby enhancing the feasibility of applying this approach to a wider range of applications. The results in Fig. 3 validate the hypothesis by demonstrating that a subset of the LLaMA encoder can also effectively refine relations.

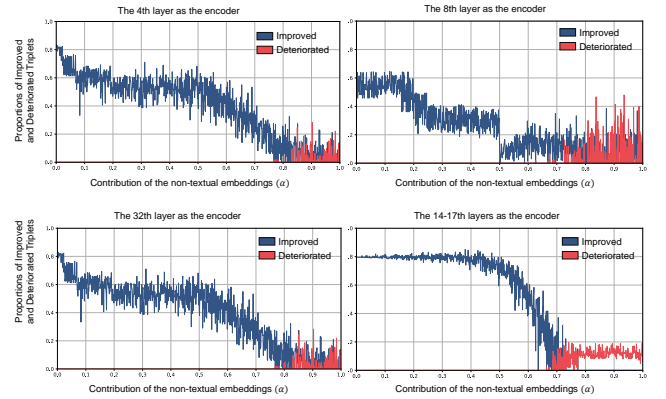


Figure 3: The impact of utilizing specific layers from the LLM encoder for relation refinement. The performance of individual layers (4^{th} , 8^{th} , and 32^{nd}) as well as combined layers (14^{th} to 17^{th}) have been studied.

4 METHOD

4.1 A General Framework for VMR

Before introducing the proposed method, we summarize existing VMR approaches into a general framework to ease the description (see Fig. 4). The VMR task takes a video V and a textual query Q as input and predicts a set of candidate video segments $\{m_k\}$ that are relevant to the query Q . This process is accomplished by an inference process of a neural network f_θ with parameters θ as

$$\{m_k\} = f_\theta(Q, V). \quad (3)$$

The process begins by encoding the query and video into embeddings $\{q_i\}$ and $\{v_j\}$, respectively, as

$$\{q_i\} = \text{CLIP}_t(Q), \quad \{v_j\} = \text{CLIP}_v(V) \oplus \text{SlowFast}(V), \quad (4)$$

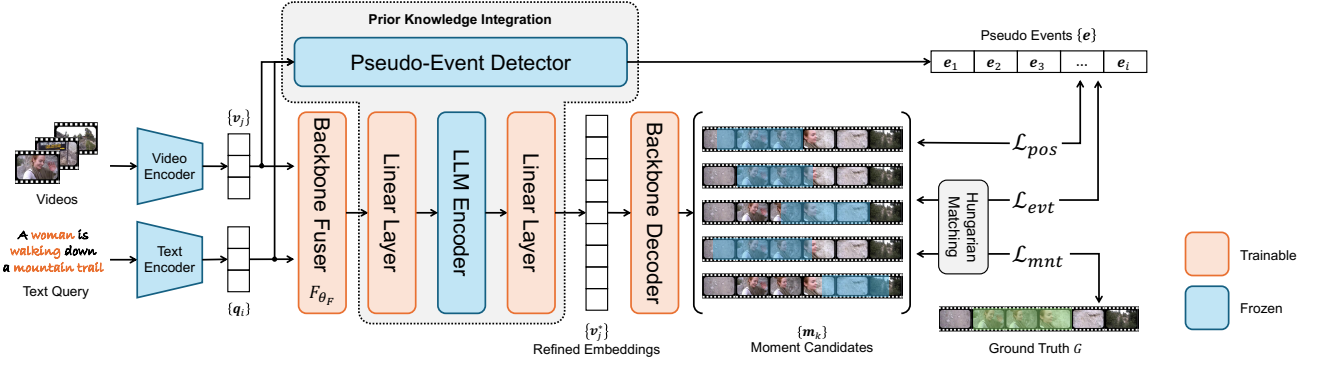


Figure 4: The proposed general framework for VMR with the proposed prior knowledge integration components.

where the $CLIP_t$ and $CLIP_v$ are the textual and visual encoders of CLIP [56], respectively, and the SlowFast [16] is a commonly used video encoder. The operator \oplus denotes concatenation. The embeddings are then fused by a network F with parameters θ_F as

$$\{\mathbf{v}_j^*\} = F_{\theta_F}(\{\mathbf{q}_i\}, \{\mathbf{v}_j\}). \quad (5)$$

A decoder D with parameters θ_D then predicts the moments as

$$\{\mathbf{m}_k\} = D_{\theta_D}(\{\mathbf{v}_j^*\}). \quad (6)$$

A loss function is defined with the ground truth G as

$$\mathcal{L}_{mnt} = \text{Dist}(\{\mathbf{m}_k\}, G), \quad (7)$$

where Dist is a distance metric and often implemented using Hungarian Match [6]. The learning goal is to find an optimal set of parameters $\theta = \{\theta_F, \theta_D\}$.

A large portion of VMR methods can be structured within this general framework, with their unique characteristics stemming from proposed variations in the fusion network Fu or adjustments to the loss function \mathcal{L} . In our paper, we introduce two plug-in elements to this framework, which implement prior integration of semantic refinement and pseudo-event regulation, respectively. These plug-in elements are designed to be compatible with a range of existing methods, emphasizing their versatility and adaptability.

4.2 Integration of Semantic Refinement

To incorporate semantic refinement within the general framework, we can insert an LLM encoder between the fusion and decoding processes. The feasibility is that the fusion process exhibits similarities to the process described in Eq. (1), where video embeddings are merged with textual query embeddings. Furthermore, CLIP embeddings play a significant role in both the encoding and fusion processes, making them suitable as inputs of the LLM encoder, as observed in our feasibility study. Considering the rationale behind it, the utilization of the LLM encoder for inter-concept refinement can enhance the model's overall comprehension of both foreground concepts (e.g., *persons*, *faces*) and background concepts. This is crucial in preventing the model from being misled by visually dominant foreground concepts. This becomes particularly significant when these concepts are dispersed across consecutive frames, and their combined semantics cannot be identified without proper modeling

of inter-frame relations. However, to adapt to the input dimension of the LLM encoder, we need to include a linear layer before and after the LLM encoder. The integration is then written by replacing Eq. (5) as

$$\{\mathbf{v}_j^*\} = L2_{\theta_{L2}} \left[\text{LLM} \left(L1_{\theta_{L1}} \left[F_{\theta_F}(\{\mathbf{q}_i\}, \{\mathbf{v}_j\}) \right] \right) \right], \quad (8)$$

where the $L1$ and $L2$ are the two linear adapter layers with parameters θ_{L1} and θ_{L2} , respectively. In order to improve computational efficiency, we can select a subset of layers from the LLM to serve as the refiner.

4.3 Integration of Pseudo-Event Regulation

The motivation for this regulation arises from the recognition that valid moments should remain within the boundaries of events instead of crossing them. By guiding the predicted moments to align with the content distribution indicated by the event boundaries, we can effectively eliminate distractions from adjacent irrelevant moments. To this end, we utilize event detectors such as the recursive TSM parsing mechanism in UBoCo [24] to generate pseudo events for a given video. These pseudo events serve as a prior for the distribution of event boundaries. This prior can be used to guide the predicted moments into positions between event boundaries by penalizing those that extend beyond the boundaries. Let us denote the position and width of a predicted moment \mathbf{m}_k as a vector $\mathbf{p}_k = [pos_k, wid_k]$ and the detected pseudo events as a set $\mathbf{e} = [pos_e, wid_e]$. This can be implemented by introducing a pseudo-event regulated loss as

$$\mathcal{L}_{evt} = \sum_{\mathbf{e} \in \{\mathbf{e}\}} \left(\lambda_{L1} \|\mathbf{e} - \mathbf{m}_k\|_1 + \lambda_{IoU} \mathcal{L}_{IoU}(\mathbf{e}, \mathbf{m}_k) \right), \quad (9)$$

where \mathcal{L}_{IoU} is the widely adopted Generalized Intersection Over Union (GIoU) loss [59], λ_{L1} and λ_{IoU} are the weights for the L1 norm and GIoU loss, respectively.

The aforementioned regulation is designed to occur after the prediction process, as a post-validation step. We have also devised a pre-reinforcement-based regulation technique to further enhance the alignment between predicted moments and events. The concept behind this technique involves adjusting the position embeddings of

frames by reinforcing the notion that frames within the same event should possess similar position embeddings. By doing so, the prediction process is encouraged to select frames from the same event for moment composition, rather than across different events. We set the target position embeddings as those of the fused embeddings after $F_{\theta_f}(\cdot)$. For an event \mathbf{e} , we denote the position embeddings of its member frames as a matrix \mathbf{P}_e , wherein the position embedding of the center frame is \mathbf{p}_e . The position regulation is then written as

$$\mathcal{L}_{pos} = \sum_{\mathbf{e} \in \{\mathbf{e}\}} \exp \left[\frac{1}{N} \mathbf{1}^\top \text{abs}(\mathbf{P}_e - \mathbf{p}_e \cdot \mathbf{1}^\top) \mathbf{1} \right], \quad (10)$$

where $\mathbf{1}$ is a vector of ones. The \mathcal{L}_{pos} encourages the position embeddings of the member frames to that of the center embedding within an event.

By incorporating both semantic refinement and pseudo-event regulation, the loss function is expanded as

$$\mathcal{L} = \mathcal{L}_{mnt} + \lambda_e \mathcal{L}_{eot} + \lambda_p \mathcal{L}_{pos}, \quad (11)$$

where λ_e and λ_p are hydrometers to balance the learning. The extended set of parameters to be tuned is denoted as $\theta = \{\theta_F, \theta_D, \theta_{L1}, \theta_{L2}\}$.

5 EXPERIMENTS

5.1 Dataset and Evaluation Metrics

To evaluate the performance of our proposed method, we conduct extensive experiments across multiple tasks, organized into three panels: 1) *Joint moment retrieval and highlight detection* on QVHighlights [29], which includes more than 10,000 videos with high-quality text queries. For moment retrieval, we report Recall@1 with IoU thresholds of 0.5 and 0.7, mean Advance Precision (mAP) with IoU thresholds of 0.5 and 0.75, and mean mAP across a range of IoU thresholds [0.5:0.05:0.95]. For highlight detection, we report mAP and HIT@1, which consider a segment as truly “positive” when its predicted saliency score is rated as “very good”. 2) *Individual moment retrieval* on Charades-STA [17] and TACoS [58]. The metrics involve Recall@1 and Recall@5 with IoU thresholds of 0.5 and 0.7. 3) *Highlight detection* on TVSum [67] and Youtube-HL [68], where Top-5 mAP and mAP are employed as performance measures, respectively. All metrics are aligned with previous studies [29, 43, 50].

5.2 Implementation Details

We utilize the SlowFast [16] and CLIP [56] models to extract video features on QVHighlights dataset every 2 seconds. We adopt SlowFast+Clip and VGG [64] on Charades-STA and TACoS datasets, where video features are extracted every 1 second and 2 seconds, respectively. For YouTube-HL and TVSum datasets, we extract clip-level features by a pre-trained I3D [7]. Following previous methods [43], each feature vector captures 32 consecutive frames and is considered as a clip when the overlap exceeds 50% with each other.

We utilize the 14th to 17th encoding layers of LLaMA (7B) across all configurations. The feature fusion encoder F and decoder D are implemented using 6-layer Transformer. The loss balancing hydrometers λ_e and λ_p are set to 0.1 and 0.001, respectively. We use AdamW [44] with a weight decay of $1e-5$ and apply an additional pre-discard rate of 0.5 to the visual inputs. We train our model across five datasets with specific settings for each: 1) QVHighlights:

Table 1: Results of joint moment retrieval and highlight detection (HD) on QVHighlights test split [28].

Method	Moment Retrieval					HD	
	R1		mAP			\geq Very Good	
	@ 0.5	@ 0.7	@ 0.5	@ 0.75	Avg.	mAP	HIT@1
BeautyThumb [66]	-	-	-	-	-	14.36	20.88
DVSE [42]	-	-	-	-	-	18.75	21.79
MCN [20]	11.41	2.72	24.94	8.22	10.67	-	-
CAL [15]	25.49	11.54	23.40	7.65	9.89	-	-
CLIP [56]	16.88	5.19	18.11	7.0	7.67	31.30	61.04
XML [30]	41.83	30.35	44.63	31.73	32.14	34.49	55.25
XML+ [30]	46.69	33.46	47.89	34.67	34.90	35.38	55.06
Moment-DETR [29]	52.89	33.02	54.82	29.40	30.73	35.69	55.60
UMT [43]	56.23	41.18	53.83	37.01	36.12	38.18	59.99
UniVTG [36]	58.86	40.86	57.60	35.59	35.47	38.20	60.96
QD-DETR [49]	62.40	44.98	62.52	39.88	39.86	38.94	62.40
CG-DETR [50]	65.43	48.38	64.51	42.77	42.86	40.33	66.21
Ours	66.73	49.94	65.76	43.91	44.05	40.33	65.69

Learning rate of $1e-4$, batch size of 32, for 200 epochs. 2) Charades-STA: Learning rate of $2e-4$, batch size of 32, for 200 epochs. 3) TACoS: Learning rate of $2e-4$, batch size of 16, for 200 epochs. 4) YouTube-HL: Learning rate of $1e-4$, batch size of 4, for 2,000 epochs. 5) TVSum: Learning rate of $1e-4$, batch size of 1, for 2,000 epochs.

5.3 Comparison with State-of-the-arts

Joint Moment Retrieval and Highlight Detection. Our evaluation is detailed in Tab. 1, specifically focusing on the QVHighlights test split. It is obvious that the proposed approach reaches the highest scores across almost all evaluation metrics. Specifically, we achieve an average mAP of 44.05 on moment retrieval and a HIT@1 of 65.69 on highlight detection, significantly surpassing traditional methods. When compared to contemporary models such as CG-DETR, our method still exhibits clear advantages, achieving superior mAP performance in terms of 2.78% on moment retrieval. This demonstrates the effectiveness of our approach in handling complex video analysis tasks.

Moment Retrieval. As shown in Tab. 2, we report the results of moment retrieval on two additional benchmarks, TACoS and Charades-STA. Notably, our approach outperforms all SOTA methods even without pre-training. Specifically, we exceed other methods by margins of 0.19% to 34.28% on TACoS and 0.6% to 21.82% on Charades-STA in terms of mIoU. The consistent outperformance across various benchmarks and comparison with a wide range of SOTA methods suggest the effectiveness of incorporating prior knowledge from large language models.

Highlight Detection. We conduct more experiments on highlight detection. Tab. 3 displays the Top-5 mAP on the TV-Sum dataset, while Tab. 4 details the mAP performance on YouTube-HL. From an overall perspective, our method demonstrates significant improvements, achieving an mAP of 88.1% on TV-Sum and 75.3% on YouTube-HL.

Table 2: Moment retrieval results tested on TACoS and Charades-STA datasets. Video features are extracted using Slowfast and CLIP.

Method	TACoS			Charades-STA		
	R@0.3	R@0.7	mIoU	R@0.3	R@0.7	mIoU
2D-TAN	40.01	12.92	27.22	58.76	27.50	41.25
VSLNet	35.54	13.15	24.99	60.30	24.14	41.58
Moment-DETR	37.97	11.97	25.49	65.83	30.59	45.54
QD-DETR	-	-	-	-	32.55	-
LLaViLo	-	-	-	-	33.43	-
UniVTG	51.44	17.35	33.60	70.81	35.65	50.10
CG-DETR	52.23	22.23	36.48	70.43	36.34	50.13
Ours	52.73	22.78	36.55	70.91	36.49	50.25

Table 3: Highlight detection results on TV-Sum, † denotes the methods that utilize the audio modality.

Method	VT	VU	GA	MS	PK	PR	FM	BK	BT	DS	Avg.
sLSTM [88]	41.1	46.2	46.3	47.7	44.8	46.1	45.2	40.6	47.1	45.5	45.1
SG [48]	42.3	47.2	47.5	48.9	45.6	47.3	46.4	41.7	48.3	46.6	46.2
LIM-S [78]	55.9	42.9	61.2	54.0	60.3	47.5	43.2	66.3	69.1	62.6	56.3
Trailer [73]	61.3	54.6	65.7	60.8	59.1	70.1	58.2	64.7	65.6	68.1	62.8
SL-Module [81]	86.5	68.7	74.9	86.2	79.0	63.2	58.9	72.6	78.9	64.0	73.3
QD-DETR [50]	88.2	87.4	85.6	85.0	85.8	86.9	76.4	91.3	89.2	73.7	85.0
UniVTG [36]	83.9	85.1	89.0	80.1	84.6	81.4	70.9	91.7	73.5	69.3	81.0
MINI-Net† [22]	80.6	68.3	78.2	81.8	78.1	65.8	57.8	75.0	80.2	65.5	73.2
TCG† [84]	85.0	71.4	81.9	78.6	80.2	75.5	71.6	77.3	78.6	68.1	76.8
Joint-VA† [2]	83.7	57.3	78.5	86.1	80.1	69.2	70.0	73.0	97.4	67.5	76.3
UMT† [43]	87.5	81.5	88.2	78.8	81.4	87.0	76.0	86.9	84.4	79.6	83.1
Ours	90.8	91.9	94.2	88.7	85.8	90.4	78.6	93.4	88.3	78.7	88.1

Table 4: Performance of mAP for highlight detection on YouTube-HL. † denotes using audio modality.

Method	Dog	Gym	Par.	Ska.	Ski.	Sur.	Avg.
RRAE [83]	49.0	35.0	50.0	25.0	22.0	49.0	38.3
GIFs [19]	30.8	33.5	54.0	55.4	32.8	54.1	46.4
LSVM [68]	60.0	41.0	61.0	62.0	36.0	61.0	53.6
LIM-S [78]	57.9	41.7	67.0	57.8	48.6	65.1	56.4
SL-Module [81]	70.8	53.2	77.2	72.5	66.1	76.2	69.3
QD-DETR [50]	72.2	77.4	71.0	72.7	72.8	80.6	74.4
UniVTG [36]	71.8	76.5	73.9	73.3	73.2	82.2	75.2
MINI-Net† [22]	58.2	61.7	70.2	72.2	58.7	65.1	64.4
TCG † [84]	55.4	62.7	70.9	69.1	60.1	59.8	63.0
Joint-VA† [2]	64.5	71.9	80.8	62.0	73.2	78.3	71.8
UMT † [43]	65.9	75.2	81.6	71.8	72.3	82.7	74.9
Ours	73.6	74.2	72.5	75.3	73.4	82.5	75.3

5.4 Ablation study

Due to the space limitation, we limit the ablation study on QVHighlights validation split. QD-DETR [50] is temporally used as the baseline, and later we will conduct a compatibility study to show that our proposed plug-in components can be used by other 5 different frameworks to improve their performances. The results are shown in Tab. 5, where a consistent performance gain is observed

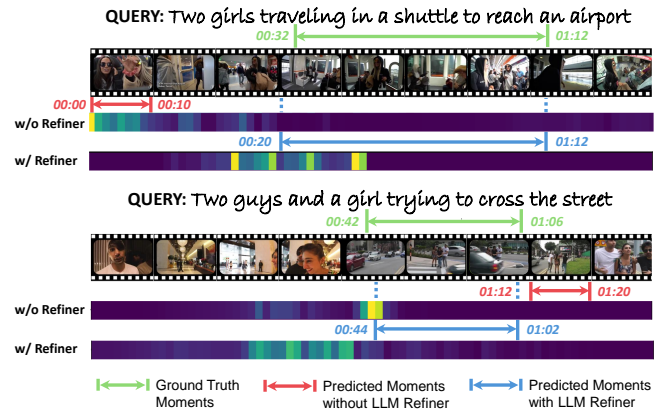


Figure 5: Illustration of the effectiveness of using the LLM as a relation refiner. The predictions with the LLM encoder are better aligned with the ground truth. The model without the refiner focuses more on the visually dominate concepts (e.g., girls, guys), while with the refiner, contextual concepts (e.g., traveling, crossing-street) can be further incorporated.

Table 5: Ablation study on QVHighlights validation split.

Models	R1@0.5	R1@0.7	mAP
1 Baseline	63.87	48.71	41.46
2 Baseline + LLM	66.71 _{2.8} ↑	49.42 _{0.7} ↑	45.69 _{4.2} ↑
3 Baseline + \mathcal{L}_{evt}	65.16 _{1.3} ↑	50.26 _{1.6} ↑	45.47 _{4.0} ↑
4 Baseline + \mathcal{L}_{pos}	64.58 _{0.7} ↑	49.48 _{0.8} ↑	43.88 _{2.4} ↑
5 Baseline + LLM + \mathcal{L}_{evt}	66.00 _{2.1} ↑	51.55 _{2.8} ↑	45.78 _{4.3} ↑
6 Baseline + LLM + \mathcal{L}_{evt} + \mathcal{L}_{pos}	66.58 _{2.7} ↑	51.10 _{2.4} ↑	46.24 _{4.8} ↑

when each component is combined into the framework, resulting in an improvement ranging from 0.7% to 4.8% in mAP.

Effectiveness of using LLM encoders as relation refiners: It is evident in Tab. 5 that LLM can bring further improvement gain due to its ability to refine the inter-frame relation. To grasp more insights, the predicted moments with and without the LLM refiner are shown in Fig. 5 on two exemplar video segments. In the example of the “two girls traveling in a shuttle”, the model without the LLM refiner focuses more on a moment that the two girls are visually dominant in the video, while the model with the LLM refiner focuses on a moment that the concepts of girls, traveling and, shuttle-to-airport all appeared with comparably less but equivalent salience in the video. This is an indication that the LLM refiner can help model the inter-concept relation in a better way. Similarly, in the second example, the model without the refiner focuses more on the concepts of two-guy and girl, and with the LLM refiner, the presence of the action of cross-street has been further considered.

Effectiveness of the pseudo-event regulation: In Fig. 6, we give two examples to illustrate the effectiveness. It is evident that the predicted moments align with the pseudo-event boundaries in a better way, eliminating the distractions from adjacent moments.

Effectiveness of position embedding regulation: After implementing pseudo-events that identify event boundaries, we can use these boundaries to refine the distribution of position embeddings.

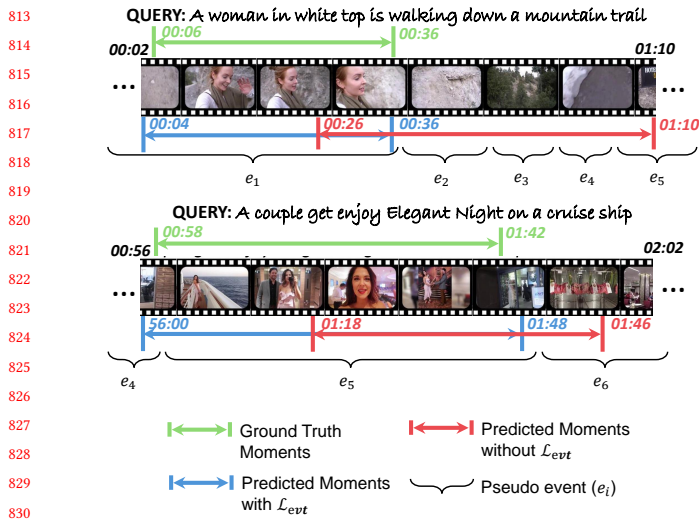


Figure 6: Illustration of the effectiveness of the event regulation. The regulation guides predicted moments to reside inside the event boundaries rather than crossing them, which eliminates the distractions from adjacent moments.

This adjustment aims for the position embeddings to align more closely with the event distribution in the video, as illustrated in Fig. 7. The impact of adopting \mathcal{L}_{pos} is distinctly visible: prior to its use, the distribution of position embeddings is notably dispersed, while with the regulated loss \mathcal{L}_{pos} , the position embeddings show better alignment to the event boundaries.

Study of the compatibility to different frameworks: By inserting the proposed integration components into various VMR frameworks, in Tab. 6, our method demonstrates high adaptability and compatibility with existing frameworks, indicated by a consistent performance gain over the original settings. The improvements observed are not merely instances of isolated success but reflect consistent effectiveness for the proposed integration components. **Qualitative Results:** Due the space limitation, we can study only a few examples in this section. We have included more illustrations and qualitative analysis in the Appendix.

6 CONCLUSION

In conclusion, our work presents an approach for enhancing video moment retrieval (VMR) by integrating large language model (LLM) encoders and pseudo-event regulation. The LLM encoders contribute to refining multimodal embeddings and their inter-concept relations, successfully applied to various embeddings such as CLIP and BLIP. We also use pseudo-events as temporal content distribution priors that aid in aligning moment predictions with actual event boundaries, addressing a previously underexplored aspect of VMR. The proposed methods serve as plug-in components, compatible with existing VMR frameworks, and have been empirically validated to achieve state-of-the-art performance. This study not only addresses the challenges posed by the vast and growing video content but also opens new avenues for the application of LLMs in fine-grained video analysis tasks.

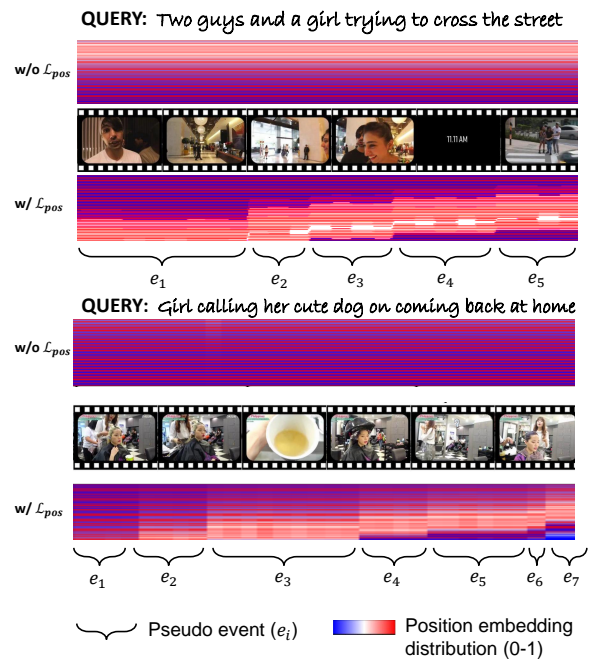


Figure 7: Illustration of the effectiveness of the position embedding regulation. The resulting embeddings align with the event distribution in a better way.

Table 6: Performance comparison of our integrated method with existing VMR frameworks (base models). Please note that the reported results of the base models might slightly differ from those in the original papers due to variations in the environment configuration. We have utilized the released source code, but the exact replication of the original environment is not feasible.

Method	Metric	Original	Integrated	Gain (%)
Moment-DETR [29]	MR-full-R1@0.5	53.90	59.23	+9.89
	MR-full-R1@0.7	34.80	38.52	+10.69
	MR-full-mAP	32.20	34.36	+6.71
QD-DETR [50]	MR-full-R1@0.5	62.70	62.90	+0.32
	MR-full-R1@0.7	46.70	46.45	-0.54
	MR-full-mAP	41.20	41.34	+0.34
EaTR [23]	MR-full-R1@0.5	57.42	63.03	+9.77
	MR-full-R1@0.7	42.58	46.06	+8.17
	MR-full-mAP	38.98	41.05	+5.31
UniVTG [36]	MR-full-R1@0.5	59.74	59.70	-0.07
	MR-full-R1@0.7	35.59	38.84	+9.13
	MR-full-mAP	36.13	36.36	+0.64
CG-DETR [49]	MR-full-R1@0.5	67.40	70.26	+4.24
	MR-full-R1@0.7	52.10	53.61	+2.90
	MR-full-mAP	44.90	45.78	+1.96

REFERENCES

- [1] 2023. GPT-4V(ision) System Card. <https://api.semanticscholar.org/CorpusID:263218031>
- [2] Taivanbat Badamdorj, Mrigank Roachan, Yang Wang, and Li Cheng. 2021. Joint Visual and Audio Learning for Video Highlight Detection. In *2021 IEEE/CVF*

- 929 *International Conference on Computer Vision (ICCV)*. 8107–8117. <https://doi.org/10.1109/ICCV48922.2021.00802>
- 930 [3] Taivanbat Badamdorj, Mrigank Rochan, Yang Wang, and Li Cheng. 2022. Con-
- 931 trastive learning for unsupervised video highlight detection. In *Proceedings of the*
- 932 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14042–14052.
- 933 [4] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang
- 934 Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-VL: A Versatile Vision-Language
- 935 Model for Understanding, Localization, Text Reading, and Beyond. *arXiv preprint*
- 936 *arXiv:2308.12966* (2023).
- 937 [5] Sijia Cai, Wangmeng Zuo, Larry S. Davis, and Lei Zhang. 2018. Weakly-
- 938 supervised Video Summarization using Variational Encoder-Decoder and Web
- 939 Prior. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- 940 [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexan-
- 941 der Kirillov, and Sergey Zagoruyko. 2020. End-to-End Object Detection with
- 942 Transformers. *arXiv:2005.12872* [cs.CV]
- 943 [7] João Carreira and Andrew Zisserman. 2017. Quo Vadis, Action Recognition? A
- 944 New Model and the Kinetics Dataset. In *2017 IEEE Conference on Computer Vision*
- 945 *and Pattern Recognition (CVPR)*. 4724–4733. <https://doi.org/10.1109/CVPR.2017.502>
- 946 [8] Jingyuan Chen, Xinpeng Chen, Lin Ma, Zequn Jie, and Tat-Seng Chua. 2018.
- 947 Temporally grounding natural sentence in video. In *Proceedings of the 2018*
- 948 *Conference on Empirical Methods in Natural Language Processing*. 162–171.
- 949 [9] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang,
- 950 Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed
- 951 Elhoseiny. 2023. MiniGPT-v2: large language model as a unified interface for
- 952 vision-language multi-task learning. *arXiv:2310.09478* [cs.CV]
- 953 [10] Shaoxiang Chen, Wenhao Jiang, Wei Liu, and Yu-Gang Jiang. 2020. Learning
- 954 modality interaction for temporal sentence localization and event captioning in
- 955 videos. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK,*
- 956 *August 23–28, 2020, Proceedings, Part IV*, Vol. 16. Springer, 333–351.
- 957 [11] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang,
- 958 Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica,
- 959 and Eric P. Xing. 2023. Vicuna: An Open-Source Chatbot Impressing GPT-4 with
- 960 90%* ChatGPT Quality. <https://lmsys.org/blog/2023-03-30-vicuna/>
- 961 [12] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav
- 962 Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebas-
- 963 tian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez,
- 964 Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran,
- 965 Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin,
- 966 Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay
- 967 Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin
- 968 Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek
- 969 Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani
- 970 Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan
- 971 Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Huang, Andrew Dai,
- 972 Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts,
- 973 Denny Zhou, Quoc V. Le, and Jason Wei. 2022. Scaling Instruction-Finetuned
- 974 Language Models. *arXiv:2210.11416* [cs.LG]
- 975 [14] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao,
- 976 Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. InstructBLIP:
- 977 Towards General-purpose Vision-Language Models with Instruction Tuning.
- 978 *arXiv:2305.06500* [cs.CV]
- 979 [15] Victor Escorcia, Mattia Soldan, Josef Sivic, Bernard Ghanem, and Bryan Russell.
- 980 2019. Temporal Localization of Moments in Video Collections with Natural
- 981 Language.
- 982 [16] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. 2019. Slow-
- 983 Fast Networks for Video Recognition. In *Proceedings of the IEEE/CVF International*
- 984 *Conference on Computer Vision (ICCV)*.
- 985 [17] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. 2017. Tall: Temporal
- 986 activity localization via language query. In *Proceedings of the IEEE international*
- 987 *conference on computer vision*. 5267–5275.
- 988 [18] Runzhou Ge, Jiyang Gao, Kan Chen, and Ram Nevatia. 2019. Mac: Mining activity
- 989 concepts for language-based temporal localization. In *2019 IEEE winter conference*
- 990 *on applications of computer vision (WACV)*. IEEE, 245–253.
- 991 [19] Michael Gygli, Yale Song, and Liangliang Cao. 2016. Video2gif: Automatic
- 992 generation of animated gifs from video. In *Proceedings of the IEEE conference on*
- 993 *computer vision and pattern recognition*. 1001–1009.
- 994 [20] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell,
- 995 and Bryan Russell. 2017. Localizing moments in video with natural language. In
- 996 *Proceedings of the IEEE international conference on computer vision*. 5803–5812.
- 997 [21] Fa-Ting Hong, Xuanteng Huang, Wei-Hong Li, and Wei-Shi Zheng. 2020. Mini-
- 998 net: Multiple instance ranking network for video highlight detection. In *European*
- 999 *Conference on Computer Vision*. Springer, 345–360.
- 990 [22] Fa-Ting Hong, Xuanteng Huang, Wei-Hong Li, and Wei-Shi Zheng. 2020.
- 991 MINI-Net: Multiple Instance Ranking Network for Video Highlight Detection.
- 992 *arXiv:2007.09833* [cs.CV]
- 993 [23] Jinhyun Jang, Jungin Park, Jin Kim, Hyeonjun Kwon, and Kwanghoon Sohn.
- 994 2023. Knowing Where to Focus: Event-aware Transformer for Video Grounding.
- 995 In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- 996 [24] Hyolim Kang, Jinwoo Kim, Taehyun Kim, and Seon Joo Kim. 2022. UBoCo:
- 997 Unsupervised Boundary Contrastive Learning for Generic Event Boundary Detec-
- 998 tion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
- 999 *Recognition (CVPR)*. 20073–20082.
- 990 [25] Aditya Khosla, Raffay Hamid, Chih-Jen Lin, and Neel Sundaresan. 2013. Large-
- 991 scale video summarization using web-image priors. In *Proceedings of the IEEE*
- 992 *conference on computer vision and pattern recognition*. 2698–2705.
- 993 [26] Jing Yu Koh, Daniel Fried, and Ruslan Salakhutdinov. 2023. Generating Images
- 994 with Multimodal Language Models. *arXiv:2305.17216* [cs.CL]
- 995 [27] Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. 2023. Grounding Language
- 996 Models to Images for Multimodal Inputs and Outputs. *arXiv:2301.13823* [cs.CL]
- 997 [28] Jie Lei, Tamara L. Berg, and Mohit Bansal. [n. d.]. QVhighlights test split. <https://codalab.lisn.upsaclay.fr/competitions/6937#results>
- 998 [29] Jie Lei, Tamara L. Berg, and Mohit Bansal. 2021. QVHighlights: Detecting
- 999 Moments and Highlights in Videos via Natural Language Queries. *arXiv:2107.09609* [cs.CV]
- 1000 [30] Jie Lei, Licheng Yu, Tamara L. Berg, and Mohit Bansal. 2020. TVR: A Large-Scale
- 1001 Dataset for Video-Subtitle Moment Retrieval. *arXiv:2001.09099* [cs.CV]
- 1002 [31] Chenliang Li, Haiyang Xu, Junfeng Tian, Wei Wang, Ming Yan, Bin Bi, Jiabo
- 1003 Ye, Hehong Chen, Guohai Xu, Zheng Cao, et al. 2022. mPLUG: Effective and
- 1004 Efficient Vision-Language Learning by Cross-modal Skip-connections. *arXiv*
- 1005 *preprint arXiv:2205.12005* (2022).
- 1006 [32] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. BLIP-2: Boot-
- 1007 strapping Language-Image Pre-training with Frozen Image Encoders and Large
- 1008 Language Models. *arXiv:2301.12597* [cs.CV]
- 1009 [33] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. BLIP: Bootstrapping
- 1010 Language-Image Pre-training for Unified Vision-Language Understanding and
- 1011 Generation. *arXiv:2201.12086* [cs.CV]
- 1012 [34] KunChang Li, Yanan He, Yi Wang, Yizhuo Li, Wenhao Wang, Ping Luo, Yali Wang,
- 1013 Limin Wang, and Yu Qiao. 2024. VideoChat: Chat-Centric Video Understanding.
- 1014 *arXiv:2305.06355* [cs.CV]
- 1015 [35] Yanwei Li, Chengyao Wang, and Jiaya Jia. 2023. LLaMA-VID: An Image is Worth
- 1016 2 Tokens in Large Language Models. *arXiv:2311.17043* [cs.CV]
- 1017 [36] Kevin Qinghong Lin, Pengchuan Zhang, Joya Chen, Shraman Pramanick, Difei
- 1018 Gao, Alex Jingpeng Wang, Rui Yan, and Mike Zheng Shou. 2023. UniVTG: Towards
- 1019 Unified Video-Language Temporal Grounding. *arXiv:2307.16715* [cs.CV]
- 1020 [37] Bingbin Liu, Serena Yeung, Edward Chou, De-An Huang, Li Fei-Fei, and Juan
- 1021 Carlos Niebles. 2018. Temporal modular networks for retrieving complex composi-
- 1022 tional activities in videos. In *Proceedings of the European Conference on Computer*
- 1023 *Vision (ECCV)*. 552–568.
- 1024 [38] Daizong Liu, Xiaoye Qu, Jianfeng Dong, Pan Zhou, Yu Cheng, Wei Wei, Zichuan
- 1025 Xu, and Yulai Xie. 2021. Context-aware biaffine localizing network for temporal
- 1026 sentence grounding. In *Proceedings of the IEEE/CVF Conference on Computer*
- 1027 *Vision and Pattern Recognition*. 11235–11244.
- 1028 [39] Daizong Liu, Xiaoye Qu, Xiao-Yang Liu, Jianfeng Dong, Pan Zhou, and Zichuan
- 1029 Xu. 2020. Jointly cross-and self-modal graph attention network for query-based
- 1030 moment localization. In *Proceedings of the 28th ACM International Conference on*
- 1031 *Multimedia*. 4070–4078.
- 1032 [40] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual
- 1033 Instruction Tuning. *arXiv:2304.08485* [cs.CV]
- 1034 [41] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng
- 1035 Chua. 2018. Attentive moment retrieval in videos. In *The 41st International ACM*
- 1036 *SIGIR Conference on Research Development in Information Retrieval*. 15–24.
- 1037 [42] Wu Liu, Tao Mei, Yongdong Zhang, Cherry Che, and Jiebo Luo. 2015. Multi-
- 1038 task deep visual-semantic embedding for video thumbnail selection. In *2015*
- 1039 *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3707–3715.
- 1040 <https://doi.org/10.1109/CVPR.2015.7298994>
- 1041 [43] Ye Liu, Siyuan Li, Yang Wu, Chang Wen Chen, Ying Shan, and Xiaohu Qie. 2022.
- 1042 UMT: Unified Multi-modal Transformers for Joint Video Moment Retrieval and
- 1043 Highlight Detection. In *Proceedings of the IEEE/CVF Conference on Computer*
- 1044 *Vision and Pattern Recognition (CVPR)*. 3042–3051.
- 1045 [44] Ilya Loshchilov and Frank Hutter. 2017. Fixing Weight Decay Regularization in
- 1046 Adam. *CoRR* abs/1711.05101 (2017). [arXiv:1711.05101](http://arxiv.org/abs/1711.05101) <http://arxiv.org/abs/1711.05101>
- 1047 [45] Chujie Lu, Long Chen, Chilie Tan, Xiaolin Li, and Jun Xiao. 2019. Debug: A
- 1048 dense bottom-up grounding approach for natural language video localization. In
- 1049 *Proceedings of the 2019 Conference on Empirical Methods in Natural Language*
- 1050 *Processing and the 9th International Joint Conference on Natural Language Processing*

- 1045 (EMNLP-IJCNLP). 5144–5153.
- 1046 [46] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. 2023. Video-ChatGPT: Towards Detailed Video Understanding via Large Vision and Language Models. arXiv:2306.05424 [cs.CV]
- 1047 [47] Behrooz Mahasseni, Michael Lam, and Sinisa Todorovic. 2017. Unsupervised video summarization with adversarial lstm networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*. 202–211.
- 1048 [48] Behrooz Mahasseni, Michael Lam, and Sinisa Todorovic. 2017. Unsupervised Video Summarization with Adversarial LSTM Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2982–2991. <https://doi.org/10.1109/CVPR.2017.318>
- 1049 [49] WonJun Moon, Sangeek Hyun, SuBeen Lee, and Jae-Pil Heo. 2023. Correlation-guided Query-Dependency Calibration in Video Representation Learning for Temporal Grounding. *arXiv preprint arXiv:2311.08835* (2023).
- 1050 [50] WonJun Moon, Sangeek Hyun, SangUk Park, Dongchan Park, and Jae-Pil Heo. 2023. Query-dependent video representation for moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 23023–23033.
- 1051 [51] OpenAI. 2023. Chatgpt. <https://chat.openai.com/>
- 1052 [52] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. arXiv:2203.02155 [cs.CL]
- 1053 [53] Rameswar Panda, Abir Das, Ziyan Wu, Jan Ernst, and Amit K. Roy-Chowdhury. 2017. Weakly Supervised Summarization of Web Videos. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 3677–3686. <https://doi.org/10.1109/ICCV.2017.395>
- 1054 [54] Ziqi Pang, Ziyang Xie, Yunze Man, and Yu-Xiong Wang. 2023. Frozen Transformers in Language Models Are Effective Visual Encoder Layers. arXiv:2310.12973 [cs.CV]
- 1055 [55] Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. 2023. Kosmos-2: Grounding Multimodal Large Language Models to the World. arXiv:2306.14824 [cs.CL]
- 1056 [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning Transferable Visual Models From Natural Language Supervision. arXiv:2103.00020 [cs.CV]
- 1057 [57] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* 21, 140 (2020), 1–67. <http://jmlr.org/papers/v21/20-074.html>
- 1058 [58] Michaela Regneri, Marcus Rohrbach, Dominikus Wetzels, Stefan Thater, Bernt Schiele, and Manfred Pinkal. 2013. Grounding Action Descriptions in Videos. *Transactions of the Association for Computational Linguistics* 1 (2013), 25–36. https://doi.org/10.1162/tacl_a_00207
- 1059 [59] Hamid Rezaatoughi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. 2019. Generalized Intersection Over Union: A Metric and a Loss for Bounding Box Regression. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 658–666. <https://doi.org/10.1109/CVPR.2019.00075>
- 1060 [60] Mrigank Rochan, Mahesh Kumar Krishna Reddy, Linwei Ye, and Yang Wang. 2020. Adaptive video highlight detection by learning from user history. In *European conference on computer vision*. Springer, 261–278.
- 1061 [61] Anna Rohrbach, Marcus Rohrbach, Wei Qiu, Annemarie Friedrich, Manfred Pinkal, and Bernt Schiele. 2014. *Coherent Multi-sentence Video Description with Variable Level of Detail*. Springer International Publishing, 184–195. https://doi.org/10.1007/978-3-319-11752-2_15
- 1062 [62] Dian Shao, Yu Xiong, Yue Zhao, Qingqiu Huang, Yu Qiao, and Dahua Lin. 2018. Find and focus: Retrieve and localize video events with natural language queries. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 200–216.
- 1063 [63] Mike Zheng Shou, Stan Weixian Lei, Weiyao Wang, Deepti Ghadyaram, and Matt Feiszli. 2021. Generic Event Boundary Detection: A Benchmark for Event Segmentation. arXiv:2101.10511 [cs.CV]
- 1064 [64] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:1409.1556 [cs.CV]
- 1065 [65] Mattia Soldan, Mengmeng Xu, Sisi Qu, Jesper Tegner, and Bernard Ghanem. 2021. Vlg-net: Video-language graph matching network for video grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 3224–3234.
- 1066 [66] Yale Song, Miriam Redi, Jordi Vallmitjana, and Alejandro Jaimes. 2016. To Click or Not To Click: Automatic Selection of Beautiful Thumbnails from Videos. arXiv:1609.01388 [cs.MM]
- 1067 [67] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. 2015. TVSum: Summarizing Web Videos Using Titles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- 1068 [68] Min Sun, Ali Farhadi, and Steve Seitz. 2014. Ranking Domain-specific Highlights by Analyzing Edited Videos. In *ECCV*.
- [69] Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. 2023. Generative Pretraining in Multimodality. arXiv:2307.05222 [cs.CV]
- [70] Gemini Team. 2024. Gemini: A Family of Highly Capable Multimodal Models. arXiv:2312.11805 [cs.CL]
- [71] InternLM Team. 2023. InternLM: A Multilingual Language Model with Progressively Enhanced Capabilities. <https://github.com/InternLM/InternLM-techreport>.
- [72] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288 [cs.CL]
- [73] Lezi Wang, Dong Liu, Rohit Puri, and Dimitris N. Metaxas. 2020. Learning Trailer Moments in Full-Length Movies. arXiv:2008.08502 [cs.CV]
- [74] Wenhui Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, and Jifeng Dai. 2023. VisionLLM: Large Language Model is also an Open-Ended Decoder for Vision-Centric Tasks. arXiv:2305.11175 [cs.CV]
- [75] Weihang Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, Jiazheng Xu, Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie Tang. 2023. CogVLM: Visual Expert for Pretrained Language Models. arXiv:2311.03079 [cs.CV]
- [76] Shaoning Xiao, Long Chen, Songyang Zhang, Wei Ji, Jian Shao, Lu Ye, and Jun Xiao. 2021. Boundary proposal network for two-stage natural language video localization. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 2986–2994.
- [77] Bo Xiong, Yannis Kalantidis, Deepti Ghadyaram, and Kristen Grauman. 2019. Less is more: Learning highlight detection from video duration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 1258–1267.
- [78] Bo Xiong, Yannis Kalantidis, Deepti Ghadyaram, and Kristen Grauman. 2019. Less is More: Learning Highlight Detection from Video Duration. arXiv:1903.00859 [cs.CV]
- [79] Huijuan Xu, Kun He, Bryan A Plummer, Leonid Sigal, Stan Sclaroff, and Kate Saenko. 2019. Multilevel language and vision integration for text-to-clip retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 9062–9069.
- [80] Minghao Xu, Hang Wang, Bingbing Ni, Riheng Zhu, Zhenbang Sun, and Changhu Wang. 2021. Cross-category video highlight detection via set-based learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 7970–7979.
- [81] Minghao Xu, Hang Wang, Bingbing Ni, Riheng Zhu, Zhenbang Sun, and Changhu Wang. 2021. Cross-category Video Highlight Detection via Set-based Learning. arXiv:2108.11770 [cs.CV]
- [82] Shen Yan, Xuehan Xiong, Arsha Nagrani, Anurag Arnab, Zhonghao Wang, Weina Ge, David Ross, and Cordelia Schmid. 2023. Unloc: A unified framework for video localization tasks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- [83] Huan Yang, Baoyuan Wang, Stephen Lin, David Wipf, Minyi Guo, and Baining Guo. 2015. Unsupervised Extraction of Video Highlights Via Robust Recurrent Auto-encoders. arXiv:1510.01442 [cs.CV]
- [84] Qinghao Ye, Xiyue Shen, Yuan Gao, Zirui Wang, Qi Bi, Ping Li, and Guang Yang. 2021. Temporal Cue Guided Video Highlight Detection with Low-Rank Audio-Visual Fusion. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. 7930–7939. <https://doi.org/10.1109/ICCV48922.2021.00785>
- [85] Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. 2019. Semantic conditioned dynamic modulation for temporal sentence grounding in videos. *Advances in Neural Information Processing Systems* 32 (2019).
- [86] Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S Davis. 2019. Man: Moment alignment network for natural language moment retrieval via iterative graph adjustment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1247–1257.
- [87] Hao Zhang, Aixun Sun, Wei Jing, and Joey Tianyi Zhou. 2020. Span-based localizing network for natural language video localization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Online, 6543–6554.
- [88] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. 2016. Video Summarization with Long Short-term Memory. arXiv:1605.08110 [cs.CV]

1161	[89]	Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. 2020. Learning 2D temporal adjacent networks for moment localization with natural language. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , Vol. 34. 12870–12877.		
1162				
1163	[90]	Songyang Zhang, Jinsong Su, and Jiebo Luo. 2019. Exploiting temporal relationships in video moment localization with natural language. In <i>Proceedings of the 27th ACM International Conference on Multimedia</i> . 1230–1238.		
1164				
1165	[91]	Zhu Zhang, Zhijie Lin, Zhou Zhao, and Zhenxin Xiao. 2019. Cross-modal interaction networks for query-based moment retrieval in videos. In <i>Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval</i> . 655–664.		
1166				
1167				
1168				
1169				
1170				
1171				
1172				
1173				
1174				
1175				
1176				
1177				
1178				
1179				
1180				
1181				
1182				
1183				
1184				
1185				
1186				
1187				
1188				
1189				
1190				
1191				
1192				
1193				
1194				
1195				
1196				
1197				
1198				
1199				
1200				
1201				
1202				
1203				
1204				
1205				
1206				
1207				
1208				
1209				
1210				
1211				
1212				
1213				
1214				
1215				
1216				
1217				
1218				
	[92]	Kaizhi Zheng, Xuehai He, and Xin Eric Wang. 2024. MiniGPT-5: Interleaved Vision-and-Language Generation via Generative Vokens. arXiv:2310.02239 [cs.CV]		1219
				1220
	[93]	Hou Zhijian, Ngo Chong-Wah, and Chan Wing-Kwong. 2021. Conquer: Contextual query-aware ranking for video corpus moment retrieval. In <i>Proceedings of the 29th ACM International Conference on Multimedia</i> .		1221
				1222
				1223
	[94]	Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models. arXiv:2304.10592 [cs.CV]		1224
				1225
				1226
				1227
				1228
				1229
				1230
				1231
				1232
				1233
				1234
				1235
				1236
				1237
				1238
				1239
				1240
				1241
				1242
				1243
				1244
				1245
				1246
				1247
				1248
				1249
				1250
				1251
				1252
				1253
				1254
				1255
				1256
				1257
				1258
				1259
				1260
				1261
				1262
				1263
				1264
				1265
				1266
				1267
				1268
				1269
				1270
				1271
				1272
				1273
				1274
				1275
				1276