

Object-Centric Framework for Video Moment Retrieval

Zongyao Li¹, Yongkang Wong², Satoshi Yamazaki¹, Jianquan Liu¹, Mohan Kankanhalli²

¹Visual Intelligence Research Laboratories, NEC Corporation

²National University of Singapore

zongyao-li@nec.com, yongkang.wong@nus.edu.sg, s-yamazaki31@nec.com, jqliu@nec.com, mohan@comp.nus.edu.sg

Abstract

Most existing video moment retrieval methods rely on temporal sequences of frame- or clip-level features that primarily encode global visual and semantic information. However, such representations often fail to capture fine-grained object semantics and appearance, which are crucial for localizing moments described by object-oriented queries involving specific entities and their interactions. In particular, temporal dynamics at the object level have been largely overlooked, limiting the effectiveness of existing approaches in scenarios requiring detailed object-level reasoning. To address this limitation, we propose a novel object-centric framework for moment retrieval. Our method first extracts query-relevant objects using a scene graph parser and then generates scene graphs from video frames to represent these objects and their relationships. Based on the scene graphs, we construct object-level feature sequences that encode rich visual and semantic information. These sequences are processed by a relational tracklet transformer, which models spatio-temporal correlations among objects over time. By explicitly capturing object-level state changes, our framework enables more accurate localization of moments aligned with object-oriented queries. We evaluated our method on three benchmarks: Charades-STA, QVHighlights, and TACoS. Experimental results demonstrate that our method outperforms existing state-of-the-art methods across all benchmarks.

1 Introduction

Video moment retrieval, also referred to as video temporal grounding, aims to localize temporal segments within a video that correspond to a given natural language query. Recent advances have primarily focused on enhancing cross-modal interactions between visual and textual modalities (Xu et al. 2024; Moon et al. 2023b,a), and leveraging prior knowledge from the large language models (LLMs) and event boundary cues (Jiang et al. 2024; Jang et al. 2023; Boris et al. 2024). Among existing architectures, the DETR-style structure (Carion et al. 2020), which encodes visual and textual information into a set of learnable queries, has gained popularity (Lei, Berg, and Bansal 2021). In parallel, other methods adopt FPN-style models (Lin et al. 2017a) to produce multi-scale temporal feature sequences (Yan et al.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

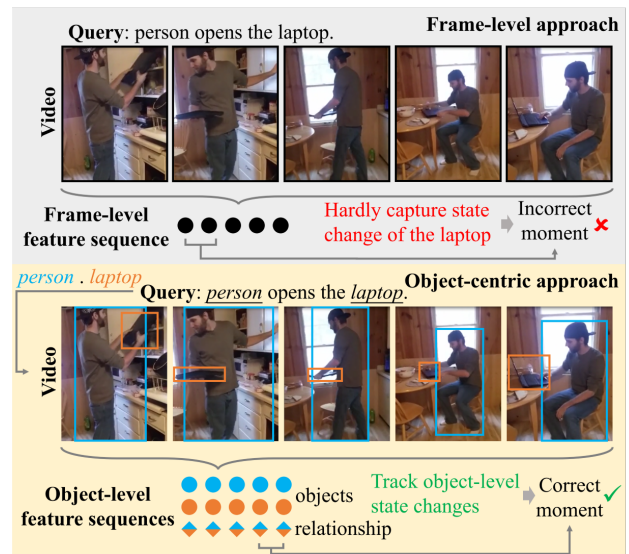


Figure 1: Comparison between the previous frame-level approach and the proposed object-centric approach. The frame-level approach fails to capture subtle yet critical state changes—such as the opening or closing of a laptop—resulting in inaccurate temporal localization. In contrast, the proposed object-centric approach explicitly tracks the states and interactions of relevant objects over time, enabling more precise moment localization.

2023; Mu, Mo, and Li 2024; Liu et al. 2024b). Despite architectural differences, most prior works rely on frame- or clip-level (hereafter referred to as frame-level) features extracted by pre-trained visual encoders such as CLIP (Radford et al. 2021) or I3D (Carreira and Zisserman 2017), which primarily capture global semantics and overlook fine-grained object-level dynamics.

Despite these advancements, a fundamental question remains underexplored: *Are frame-level features truly sufficient for video moment retrieval?* In many practical scenarios, queries are inherently object-oriented, often describing actions or events involving state changes at the object level. Nevertheless, such fine-grained queries have largely been addressed using object-agnostic approaches based on

coarse frame-level representations. While powerful visual encoders may implicitly capture some object-level temporal cues, frame-level features typically lack explicit local information—such as object semantics, appearance, and inter-object relationships—that is critical for accurately localizing moments described by object-oriented queries. This limitation poses a significant barrier to further performance improvement, particularly in cases requiring precise modeling of object-level dynamics.

To ameliorate the aforementioned limitations, we argue that object-oriented queries should be handled in an object-centric manner, with temporal modeling performed explicitly at the object level. Fig. 1 illustrates the contrast between our object-centric approach and the conventional frame-level approach. The example video depicts a person taking out a laptop, placing it on a table, and subsequently opening it. Given the query “*person opens the laptop*”, the frame-level approach struggles to capture the crucial state transition of the laptop—from closed to open—due to its reliance on coarse frame-level features. This leads to inaccurate moment localization, especially since the laptop occupies only a small portion of the visual scene. In contrast, the object-centric approach detects and tracks the relevant objects (i.e., “person” and “laptop”) and predicts their relationships throughout the video. Based on this, it constructs object-level feature sequences that encode semantic and appearance information for each object. Similarly, for each pair of related objects, a relational feature sequence is constructed to track the evolution of their interaction. Temporal modeling over these object and relationship sequences enables the model to explicitly capture state transitions aligned with the query, offering a finer and more accurate understanding than modeling over frame-level features alone.

As previously stated, this paper proposes a novel object-centric moment retrieval framework. Specifically, the framework first extracts query-relevant objects by parsing the input query into a scene graph using a scene graph parser (Li et al. 2023). An open-vocabulary scene graph generation (OVSGG) model (Chen et al. 2023) is then employed to produce scene graphs for each video frame, along with associated object and relationship features. These features are organized into temporally ordered sequences using a tracking algorithm (Zhang et al. 2022). To capture object-level state changes aligned with the query, a relational tracklet transformer is applied to the constructed object and relationship feature sequences. By modeling the spatio-temporal correlations among tracked objects and their interactions, the model enables more precise localization of moments corresponding to object-oriented queries.

The key contributions of this paper are as follows:

- We propose a novel object-centric framework for video moment retrieval, explicitly designed to capture object-level state changes that are often overlooked by prior frame-level methods.
- Within this framework, we introduce a scene-graph-based object-level feature sequence representation that encodes rich semantic and visual information. In conjunction, we design a relational tracklet transformer to

model spatio-temporal correlations among objects and their interactions.

- Extensive experiments on three benchmark datasets—Charades-STA, QVHighlights, and TACoS—demonstrate that our method consistently outperforms state-of-the-art methods, highlighting the importance of explicit object-level temporal modeling.

2 Related Work

Existing video moment retrieval methods differ in candidate generation strategies and can be broadly categorized into proposal-based two-stage methods (Gao et al. 2017; Gao and Xu 2021; Soldan et al. 2021; Xiao et al. 2021) and proposal-free one-stage methods (Lu et al. 2019; Mun, Cho, and Han 2020; Zeng et al. 2020; Zhang et al. 2020a). Two-stage methods first generate temporal proposals and then evaluate their relevance and boundaries, whereas one-stage methods directly regress the relevant temporal segment. Due to their superior efficiency and the strong temporal modeling capabilities of transformers, recent efforts have primarily focused on one-stage paradigms.

Most recent methods adopt either a point-based or query-based framework. The point-based framework (Yan et al. 2023; Mu, Mo, and Li 2024; Liu et al. 2024b) classifies temporal anchor points and regresses moment boundaries around them, drawing inspiration from point-based object detectors such as FCOS (Tian et al. 2019). In contrast, the query-based framework (Lei, Berg, and Bansal 2021; Moon et al. 2023b; Xu et al. 2024; Lee and Byun 2025), inspired by DETR (Carion et al. 2020), introduces learnable queries into a transformer to jointly encode visual and textual information for moment prediction. These DETR-style models are commonly used for both moment retrieval and highlight detection tasks, especially under the unified setting of the QVHighlights dataset (Lei, Berg, and Bansal 2021).

Several recent works extend moment retrieval beyond its traditional formulation, unifying it with other video understanding tasks such as temporal action localization (Zeng et al. 2025) and video summarization (Lin et al. 2023), aiming for improved generalization. Some works focus on enhancing cross-modal interaction: QD-DETR (Moon et al. 2023b) and MH-DETR (Xu et al. 2024) investigate advanced cross-attention mechanisms, while EaTR (Jang et al. 2023) and LLMEPET (Jiang et al. 2024) incorporate event boundary information to reduce distraction from irrelevant segments. Additionally, LLMEPET integrates an LLM encoder to refine cross-modal features. Unlike these methods, which typically adopt frozen backbones, R²-Tuning (Liu et al. 2024b) enables efficient transfer to video-language tasks by finetuning CLIP with a lightweight module.

MMRG (Zeng et al. 2021) is most relevant to our work. It captures object-level state changes by constructing a multi-modal relational graph based on interactions among visual and textual objects. However, MMRG employs a simplified graph structure that focuses on a single entity for each object category extracted from the query. In contrast, our object-centric framework enables fine-grained modeling of spatio-temporal correlations across multiple objects, allowing bet-

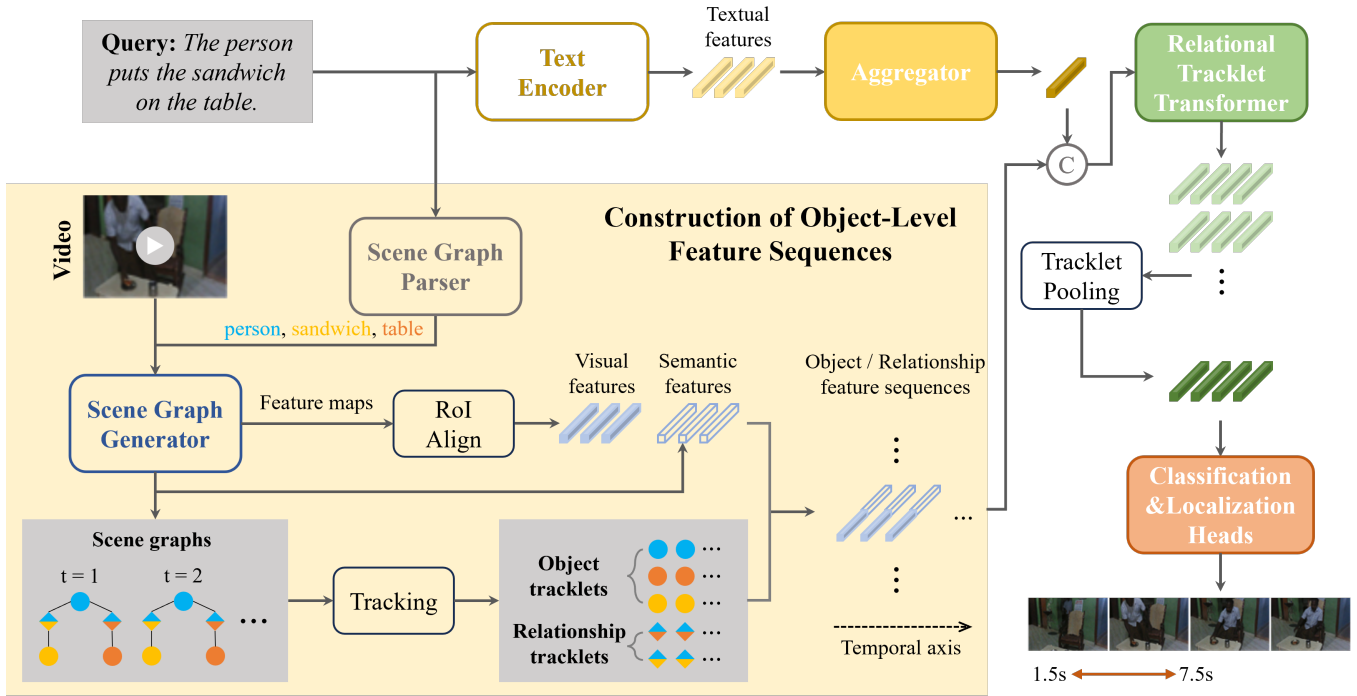


Figure 2: Overview of our object-centric moment retrieval framework. Given a video and a query, the framework first extracts query-relevant objects by parsing the query using a scene graph parser. It then constructs object-level feature sequences by embedding both visual and semantic information of these objects and their relationships into tracklets obtained from scene graph generation and object tracking. These feature sequences are concatenated with the query’s textual feature and passed to a relational tracklet transformer, which models the spatio-temporal correlations among objects and relationships. The resulting representation enables accurate moment classification and localization by capturing query-relevant object state changes.

ter generalization to complex scenarios. Graph-based representations have also been explored in highlight detection and video QA. VH-GNN (Zhang et al. 2020b) constructs spatial graphs over object nodes within frames and temporal graphs over frame nodes. However, as spatial graphs are aggregated independently into frames, it only models temporal dynamics at the frame level. Our method, by contrast, captures spatio-temporal correlations directly at the object level. VGT (Xiao et al. 2022) introduces a video graph for question answering, but is constrained to short clips and lacks the long-range modeling capacity required for moment retrieval. Moreover, unlike these prior graph-based methods, our framework leverages explicit relationship information from scene graphs, enabling a more precise understanding of object states and their temporal evolution.

3 Object-Centric Framework for Moment Retrieval

Given a video \mathbf{V} represented as a sequence of L frames $[\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_L]$, the goal of video moment retrieval is to identify the temporal segment that corresponds to a natural language query. Fig. 2 presents an overview of the proposed object-centric framework. Unlike conventional approaches that rely solely on frame-level features, our method constructs feature sequences for both objects relevant to the

query and their relationships. To model spatial and temporal correlations among these entities, we design a relational tracklet transformer that operates over the object-level feature sequences, enabling the precise localization of moments involving object-level state changes.

In Section 3.1, we introduce a frame-level baseline that serves as a comparative foundation. Section 3.2 highlights the limitations of frame-level representations and motivates our object-centric approach. We then describe the construction of object-level inputs in Section 3.3, followed by the architecture of our object-centric model in Section 3.4.

3.1 Baseline with Frame-Level Features

Following previous works (Yan et al. 2023; Mu, Mo, and Li 2024), we implement a baseline that utilizes frame-level feature sequences processed by a transformer model. Given a video $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_L]$, a pre-trained visual encoder (e.g., CLIP (Radford et al. 2021), I3D (Carreira and Zisserman 2017)) is employed to extract a sequence of frame-level features $\mathbf{F} \in \mathbb{R}^{L \times d_f}$, denoted as $[\mathbf{f}_1, \mathbf{f}_2, \dots, \mathbf{f}_L]$. When using a clip-level encoder such as I3D, each \mathbf{v}_i corresponds to a short video segment rather than a single frame.

In parallel, a pre-trained text encoder (e.g., CLIP (Radford et al. 2021), BERT (Kenton and Toutanova 2019)) encodes the input query into a sequence of token-level features $\mathbf{T} = [\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_K] \in \mathbb{R}^{K \times d_t}$. To enable multimodal fu-

sion, we aggregate the token embeddings into a single textual feature vector $\bar{\mathbf{t}} \in \mathbb{R}^{d_t}$ using a learnable linear attention mechanism. Specifically, a weight vector $\mathbf{w} \in \mathbb{R}^{d_t \times 1}$ computes attention scores $\mathbf{a} = \mathbf{T} \cdot \mathbf{w}$, which are then used to derive the aggregated representation $\bar{\mathbf{t}} = \mathbf{a}^\top \cdot \mathbf{T}$ via weighted averaging. This aggregated feature is concatenated with each frame feature, yielding the fused visual-textual sequence $\tilde{\mathbf{F}} = [\tilde{\mathbf{f}}_1, \tilde{\mathbf{f}}_2, \dots, \tilde{\mathbf{f}}_L]$

The sequence $\tilde{\mathbf{F}}$, augmented with temporal positional embeddings, is fed into a transformer consisting of stacked self-attention blocks. To handle moments of varying temporal durations, we generate multi-scale temporal features with resolutions $L, L/2, L/4$, etc., from intermediate transformer layers. Each element in the multi-scale sequences is treated as a moment candidate. A classification head predicts a relevance score p for each candidate, while a localization head regresses the corresponding start and end times (s, e) . The top- N scoring candidates $\{(p_i, s_i, e_i)\}_{i=1}^N$ are selected and post-processed using Soft-NMS (Bodla et al. 2017) to suppress redundancy.

During training, the transformer and prediction heads are jointly optimized using Focal Loss (Lin et al. 2017b) for classification and Distance-IoU Loss (Zheng et al. 2020) for localization. We adopt the Center Sampling strategy (Zhang, Wu, and Li 2022), in which only candidates near ground-truth moment centers are considered positive. The localization loss is computed solely for these positive samples.

3.2 Defects of Frame-Level Approaches

Most existing moment retrieval methods rely on frame-level features, assuming global representations suffice for aligning visual content with textual queries. However, this assumption often breaks down for object-oriented queries. We identify three core limitations:

Lack of Fine-Grained Visual Detail. Frame-level features lack localized visual information. Pre-trained visual encoders are optimized for image- or video-level objectives, which prioritize global semantics over local object appearance. As a result, the final-layer features predominantly encode high-level concepts while overlooking subtle visual cues such as an object’s state (e.g., whether a laptop is open or closed). Since state changes may not be accompanied by semantic shifts, capturing such transitions requires fine-level visual detail, which frame-level features often miss.

Insufficient Modeling of Object Relationships. Frame-level representations also fail to capture inter-object relationships, which are essential for reasoning about object states in context (e.g., “person opens laptop” implies physical interaction). This shortcoming arises because visual encoders are rarely trained to model structured relations among objects. Consequently, relational information critical for understanding events and actions is underrepresented.

Inability to Track Multiple Objects. Finally, frame-level approaches struggle to maintain consistent representations of multiple objects over time. Even if the dominant object in a scene is adequately captured, other relevant entities may be spatially small or neglected due to attention dilution. As a result, modeling the dynamics of multiple ob-

jects becomes infeasible with frame-level features alone.

To overcome these limitations, we propose an object-centric framework that explicitly models object-level appearance and relational dynamics. By leveraging scene graph generation (SGG) to detect objects and relationships at both the visual and semantic levels, and organizing them into temporally coherent tracklets, our method provides fine-grained, structured inputs for downstream modeling.

3.3 Object-Level Input Construction

As illustrated in Fig. 2, the construction of object-level inputs begins with parsing the query into a scene graph. We use a pre-trained scene graph parser (Li et al. 2023), implemented as a lightweight language model, to extract object nodes and their relationships from the natural language query. The resulting object class list is used to prompt an open-vocabulary scene graph generation (OVSGG) model (Chen et al. 2023), which detects instances of the relevant objects and their pairwise relationships in each video frame. We adopt the OVSGG model pre-trained on the Visual Genome (VG) dataset (Krishna et al. 2017). In addition to the object prompt, the OVSGG model requires a relationship class list as input. Since state-of-the-art models are still unable to reliably detect open world relationships, we use the fixed set of relationship classes from the VG dataset rather than deriving them from the query.

The OVSGG model outputs detection results along with intermediate feature maps from its Swin Transformer (Liu et al. 2021) backbone. We apply RoIAlign (He et al. 2017) to extract RoI features for each detected object or relationship. For relationships, RoI features are computed using the union of the subject and object bounding boxes. In parallel, we extract semantic features aligned with the text prompts from the final output of the OVSGG model. These visual and semantic features are complementary, together capturing both appearance-level and concept-level aspects of object states.

After detection, we apply a tracking algorithm (Zhang et al. 2022) to associate detected objects and relationships across time. Due to the low frame rate used in our setup (e.g., 1 or 0.5 fps), we perform feature-based matching using cosine similarity rather than IoU. The resulting object trajectories are organized into object tracklets, while relationships with confidence scores above a threshold are grouped into relationship tracklets, each corresponding to a specific object pair. Finally, we construct object-level feature sequences by embedding the concatenation of visual and semantic features into the corresponding tracklets. For time steps where a particular object or relationship is not present, we insert zero vectors to maintain temporal alignment.

3.4 Object-Centric Model Architecture

The object-level feature sequences introduced in Section 3.3 require a model architecture capable of effectively capturing spatio-temporal correlations among objects and their interactions. To this end, we propose a relational tracklet transformer specifically designed to process temporally ordered object and relationship tracklets.

Our model builds upon the TimeSformer architecture (Bertasius, Wang, and Torresani 2021), which factorizes

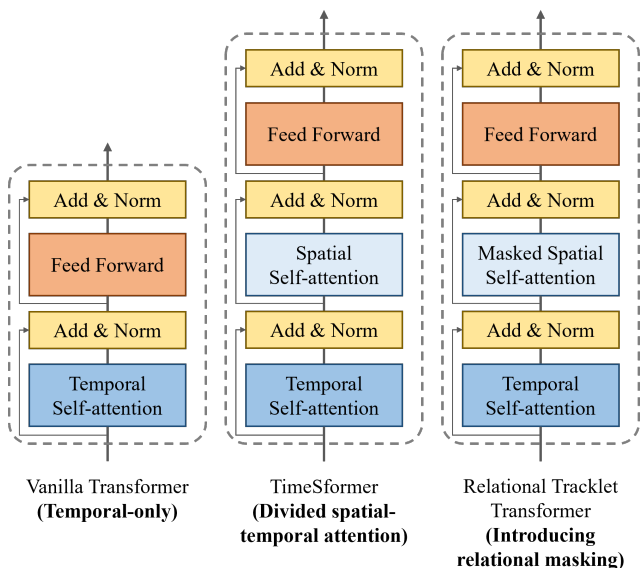


Figure 3: Comparison between three transformer block variants. The proposed variant (right) integrates a relational mask, derived from scene graph information, into the spatial self-attention mechanism, enabling the model to attend more effectively to semantically relevant object pairs.

joint spatio-temporal self-attention into separate temporal and spatial attention modules, thereby improving efficiency. To better understand the differences, Fig. 3 compares our relational tracklet transformer with a vanilla transformer and TimeSformer. A temporal-only vanilla transformer can be directly applied to object-level feature sequences but lacks the capacity to model spatial correlations. While TimeSformer addresses this by adding spatial attention, it treats all tracklets equally during spatial modeling, ignoring their semantic relationships.

To overcome this limitation, we introduce a *graph-aware spatial attention mechanism*. Specifically, we construct a binary attention mask that reflects the underlying scene graph structure and restricts spatial attention to semantically related entities. The input to the transformer is a concatenated sequence of object and relationship tracklets:

$$[\mathcal{O}; \mathcal{R}] = [\mathbf{O}_1, \mathbf{O}_2, \dots, \mathbf{O}_{N_{\mathcal{O}}}, \mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_{N_{\mathcal{R}}}],$$

where \mathcal{O} and \mathcal{R} denote the sets of object and relationship feature sequences, respectively. A binary mask $\mathbf{m} \in \mathbb{R}^{(N_{\mathcal{O}}+N_{\mathcal{R}}) \times (N_{\mathcal{O}}+N_{\mathcal{R}})}$ is constructed such that $\mathbf{m}_{i, N_{\mathcal{O}}+j} = \mathbf{m}_{N_{\mathcal{O}}+j, i} = 1$ if object tracklet \mathbf{O}_i is the subject or object of relationship tracklet \mathbf{R}_j , and 0 otherwise. Diagonal entries are set to 1 to preserve self-attention. This mask encodes the scene graph’s relational structure and is applied during spatial self-attention.

More concretely, for each temporal position j , the spatial self-attention is computed over the set of feature vectors

$$[\mathbf{o}_{1,j}, \mathbf{o}_{2,j}, \dots, \mathbf{o}_{N_{\mathcal{O}},j}, \mathbf{r}_{1,j}, \mathbf{r}_{2,j}, \dots, \mathbf{r}_{N_{\mathcal{R}},j}],$$

where $\mathbf{o}_{i,j}$ and $\mathbf{r}_{i,j}$ are the object and relationship features at time j from \mathbf{O}_i and \mathbf{R}_i , respectively. In contrast, temporal

self-attention is applied independently within each tracklet across time.

The textual feature aggregator and prediction heads in the object-centric model adopt the same design as those in the frame-level baseline (Section 3.1). Since the relational tracklet transformer produces one feature sequence per input tracklet, we apply max pooling along the tracklet dimension to obtain a unified temporal feature sequence. This sequence is then forwarded to the classification and localization heads. To mitigate potential detection failures and to leverage the complementary global context provided by frame-level features, we incorporate an optional parallel branch based on frame-level representations. The outputs from the object-centric and frame-level branches are fused via element-wise addition before being passed to the prediction heads. The training objectives remain consistent with those of the frame-level baseline.

4 Experiments

4.1 Experiment Setup

We evaluate our method on three widely used benchmarks: Charades-STA (Gao et al. 2017), QVHighlights (Lei, Berg, and Bansal 2021), and TACoS (Regneri et al. 2013). Charades-STA contains 6,672 videos (≈ 30 s each) and 16,128 text queries. QVHighlights comprises 10,148 videos and 10,310 queries, each possibly aligned with multiple moments. TACoS includes 127 long cooking videos (average length 4.8 minutes) and 18,818 query-moment pairs.

Video frames are extracted at 1 fps for Charades-STA and TACoS, and at 0.5 fps for QVHighlights. For Charades-STA and TACoS, we report Recall@1 (R1) at various temporal IoU (tIoU) thresholds and mean IoU (mIoU) between the top-1 prediction and ground truth. For QVHighlights, where queries can correspond to multiple segments, we follow prior works and use mean Average Precision (mAP) across tIoU thresholds from 0.5 to 0.95.

In all experiments, we use CLIP-B (Radford et al. 2021) as the text encoder. For the frame-level visual encoder, we evaluate two settings: (1) a combination of SlowFast-R50 (Feichtenhofer et al. 2019) and CLIP-B’s visual branch, and (2) VideoMAEv2-b (Wang et al. 2023). Additional implementation details are provided in Appendix A.

4.2 Comparison with State-of-the-Art Methods

Tab. 1 reports performance comparisons between our method and recent state-of-the-art methods across the three benchmarks. Most prior methods operate on frame-level features extracted using encoders such as CLIP (Radford et al. 2021), SlowFast (Feichtenhofer et al. 2019), I3D (Carreira and Zisserman 2017), or C3D (Tran et al. 2015). Our object-centric framework alone outperforms existing methods on all datasets. When combined with the frame-level baseline (as described in Section 3.4), further improvements are observed on Charades-STA and QVHighlights, though a performance drop occurs on TACoS.

On Charades-STA, the combined model (“Ours (w/ frame-level)”) surpasses prior methods by at least 4.2% in R1@0.7, 2.2% in R1@0.5, and 1.7% in R1@0.3. It also

Method	Charades-STA				TACoS				QVHighlights				
	R1			mIoU	R1			mIoU	R1		mAP		
	@0.3	@0.5	@0.7		@0.3	@0.5	@0.7		@0.5	@0.7	@0.5	@0.75	Avg.
MMRG (Zeng et al. 2021)	71.6	44.3	-	-	<u>57.8</u>	39.3	-	-	-	-	-	-	-
UnLoc (Yan et al. 2023)	-	60.8	38.4	-	-	-	-	-	66.1	46.7	-	-	-
UniVTG (Lin et al. 2023)	72.6	60.2	38.6	52.2	56.1	43.4	24.3	38.6	65.4	50.1	64.1	45.0	43.6
MH-DETR (Xu et al. 2024)	-	56.4	36.0	-	-	-	-	-	60.1	42.5	60.8	38.1	38.4
QD-DETR (Moon et al. 2023b)	-	57.3	32.6	-	-	-	-	-	62.4	45.0	62.5	39.9	39.9
CG-DETR (Moon et al. 2023a)	70.4	58.4	36.3	50.1	54.4	39.5	23.4	37.4	65.4	48.4	64.5	42.8	42.9
UVCOM (Xiao et al. 2024)	-	59.3	36.6	-	-	36.4	23.3	-	63.6	47.5	63.4	42.7	43.2
LLMEPET (Jiang et al. 2024)	70.9	-	36.5	50.3	52.7	-	22.8	36.6	<u>66.7</u>	49.9	65.8	43.9	44.1
SnAG* (Mu, Mo, and Li 2024)	73.8	62.9	41.9	54.1	50.9	41.2	30.9	38.6	-	-	-	-	-
BAM-DETR (Lee and Byun 2025)	72.9	60.0	39.4	52.3	56.7	41.5	26.8	39.3	62.7	48.6	64.4	46.3	45.4
R ² -Tuning (Liu et al. 2024b)	70.9	59.8	37.0	50.9	49.7	38.7	25.1	35.9	68.0	49.4	<u>69.0</u>	47.6	46.2
Ours (w/o frame-level)	<u>75.1</u>	<u>64.6</u>	<u>44.4</u>	<u>55.7</u>	58.5	47.7	34.1	43.1	65.1	<u>50.6</u>	68.4	<u>47.8</u>	<u>46.4</u>
Ours (w/ frame-level)	75.5	65.1	46.1	56.2	56.3	<u>45.2</u>	<u>33.6</u>	<u>41.9</u>	66.2	52.3	70.1	50.8	49.2

Table 1: Experimental results on Charades-STA (Gao et al. 2017), TACoS (Regneri et al. 2013), and QVHighlights (*test split*) (Lei, Berg, and Bansal 2021). The best and second-best values are highlighted in bold and underlined, respectively. “Ours (w/o frame-level)” denotes the object-centric framework without the frame-level branch, while “Ours (w/ frame-level)” includes the integration of a frame-level branch. For the latter, frame-level features are extracted using a combination of CLIP-B and SlowFast, following the common practice adopted by most previous methods. *For fair comparison, SnAG is re-evaluated on Charades-STA and TACoS using the same feature lengths as those used in our method and most of the previous methods.

achieves a 2.1% gain in mIoU, indicating better overall localization accuracy. On TACoS, our object-centric framework significantly outperforms prior methods, with margins of 3.2% (R1@0.7), 4.3% (R1@0.5), and 3.8% (mIoU). However, integrating frame-level features results in degraded performance. We hypothesize that this is due to the limited visual variance in TACoS, which consists mostly of static-camera recordings. In such cases, frame-level features struggle to capture subtle changes in small or rare objects, failing to complement the object-centric representation. On QVHighlights, our method achieves the highest average mAP, leading the second-best method by 3.0%. Although several prior methods outperform ours in R1@0.5, they lag behind in all other metrics.

4.3 Analysis of Feature Efficacy

To assess the impact of object-level and frame-level features, we conduct ablation studies summarized in Tab. 2. In addition to SlowFast+CLIP-B (SF+C), we also evaluate VideoMAEv2-b, a powerful video encoder capable of modeling motion through multi-frame inputs.

The first two rows in Tab. 2 show results for the frame-level baseline. VideoMAEv2 consistently outperforms SF+C, especially on Charades-STA, where motion understanding is critical. The smaller gap on QVHighlights reflects its greater emphasis on object semantics rather than motion. Results in the next three rows demonstrate the importance of both visual and semantic features in the object-centric framework: removing either component leads to noticeable drops in performance, confirming their complementary roles in modeling object state changes.

The final two rows evaluate the combined model (object-

centric + frame-level), which achieves consistent improvements over the object-centric model alone on both datasets. Notably, on Charades-STA, adding VideoMAEv2 features boosts mIoU by 5.6%, compared to only 0.5% with SF+C. In contrast, QVHighlights shows comparable mAP gains (2.8% vs. 3.7%).

Comparing the last two rows to the first two highlights the overall value of the object-centric framework. On Charades-STA, incorporating it yields mIoU improvements of 2.8% (SF+C) and 1.4% (VideoMAEv2). The smaller gain for VideoMAEv2 reflects its strong performance in person-centric videos, where motion modeling alone suffices. On QVHighlights, however, the object-centric framework consistently enhances both baselines, with mAP (Avg.) increases of 5.8% and 4.7%, further validating its contribution.

4.4 Impact of Scene Graph Quality

Since the object-centric framework relies on scene graphs, we investigate how scene graph quality affects performance. We train an OVSGG model on the Action Genome (AG) dataset (Ji et al. 2020), which provides scene graph annotations for keyframes in Charades-STA. This model achieves better accuracy on Charades-STA than the Visual Genome (VG)-trained version and produces more reliable scene graphs. We also experiment with ground-truth (GT) scene graphs from AG to approximate the upper bound, although they are temporally sparse and limited in vocabulary.

Tab. 3 compares results using scene graphs from three sources: VG, AG, and GT. Without frame-level features, using AG and GT graphs yields mIoU improvements of 0.3% and 11.7%, respectively, over VG. When frame-level features are included, consistent gains are observed: +0.9% and

Object-level features		Frame-level features		Charades-STA				QVHighlights				
Visual	Semantic	SF+C	VideoMAEv2	R1			mIoU	R1		mAP		
				@0.3	@0.5	@0.7		@0.5	@0.7	@0.5	@0.75	Avg.
		✓		72.5	61.8	43.1	53.4	59.4	46.1	63.8	45.4	44.1
			✓	79.1	69.5	52.1	59.9	60.7	47.6	65.2	47.9	46.1
✓				72.5	61.8	42.0	53.4	61.8	48.5	65.2	46.6	44.8
	✓			72.6	60.9	39.4	52.3	61.0	47.6	64.8	45.7	44.1
✓	✓			75.1	64.6	44.4	55.7	66.1	51.4	68.5	49.1	47.1
✓	✓	✓		75.5	65.1	46.1	56.2	67.4	54.9	69.7	52.1	49.9
✓	✓		✓	80.9	70.3	52.9	61.3	67.7	55.7	69.7	52.4	50.8

Table 2: Ablation study on object-level and frame-level features on Charades-STA and QVHighlights (*val* split). Frame-level features are extracted using two settings: (1) a combination of SlowFast and CLIP-B (SF+C), and (2) VideoMAEv2-b.

Visual encoder		Scene graph			R1			mIoU
SF+C	VMAE	VG	AG	GT	@0.3	@0.5	@0.7	
		✓			75.1	64.6	44.4	55.7
			✓		75.4	65.2	45.2	56.0
				✓	86.6	78.4	59.9	67.4
✓		✓			75.5	65.1	46.1	56.2
✓			✓		76.4	66.4	47.4	57.1
✓				✓	87.6	80.3	63.6	69.1
	✓	✓			80.9	70.3	52.9	61.3
	✓		✓		80.8	72.0	54.4	61.8
	✓			✓	88.0	81.1	64.9	69.8

Table 3: Impact of scene graph quality on performance on Charades-STA. We evaluate the effect of different scene graph sources: pre-trained models on Visual Genome (VG) and Action Genome (AG), as well as ground-truth (GT) scene graphs. Frame-level features are extracted using two settings: (1) a combination of SlowFast and CLIP-B (SF+C), and (2) VideoMAEv2-b (VMAE).

+12.9% for SF+C, and +0.5% and +8.5% for VideoMAEv2. These results highlight the strong potential of our framework to further improve as better scene graph generation models become available.

4.5 Computational Cost Analysis

The object-centric framework introduces additional components beyond frame-level methods, increasing computational cost. We analyze inference-time processing time and GPU memory usage of each major module on Charades-STA, as summarized in Tab. 4. The main source of overhead lies in the scene graph generator, which consumes more GPU memory than the visual encoder but keeps total usage around 10.7 GB—acceptable for most practical settings. In terms of processing time, the scene graph generator is about $2.76\times$ slower than the visual encoder, mainly due to its reliance on a heavy object detector. We expect this overhead can be mitigated by future OVSGG models based on real-

Module	Time cost (seconds/video)	GPU memory usage (Mb)
Scene graph parser	0.142	1549
Text encoder	0.005	885
Visual encoder	2.290	2886
Scene graph generator	6.321	3413
Tracklet transformer	0.115	1932

Table 4: Computational cost of major modules during inference on Charades-STA (measured on a single NVIDIA L40S GPU). The scene graph parser is FACTUAL (Flan-T5-base); the visual encoder is a combination of CLIP-B and SlowFast-R50; the text encoder is CLIP-B; and the scene graph generator is OvSGTR (Swin-B). The cost of the tracking process and the RoIAlign operation is negligible and accounted for in the runtime of the relational tracklet transformer and scene graph generator, respectively.

time detectors such as YOLO-World (Cheng et al. 2024), which are significantly faster than GroundingDINO (Liu et al. 2024a), used in our current setup. The costs of previous methods (Lee and Byun 2025; Jiang et al. 2024) mainly come from the visual encoder which is the same as ours, while their transformer modules are generally less costly. Therefore, Tab. 4 allows a rough cost comparison between previous methods and our method.

5 Conclusion

We proposed a novel object-centric framework for video moment retrieval that addresses the limitations of prior frame-level approaches in modeling object-level state changes. Our method constructs object-level feature sequences and captures their spatio-temporal dynamics via a relational tracklet transformer. Experiments on three benchmarks demonstrate that our method outperforms state-of-the-art methods. Additionally, we show that the framework’s performance can be further improved as open-vocabulary scene graph generation models continue to advance.

References

- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding? In *ICML*, volume 2, 4.
- Bodla, N.; Singh, B.; Chellappa, R.; and Davis, L. S. 2017. Soft-NMS—improving object detection with one line of code. In *ICCV*, 5561–5569.
- Boris, M.; Anil, B.; Anna, R.; and Marcus, R. 2024. The surprising effectiveness of multimodal large language models for video moment retrieval. *arXiv preprint arXiv:2406.18113*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *ECCV*, 213–229.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 6299–6308.
- Chen, Z.; Wu, J.; Lei, Z.; Zhang, Z.; and Chen, C. 2023. Expanding Scene Graph Boundaries: Fully Open-vocabulary Scene Graph Generation via Visual-Concept Alignment and Retention. *arXiv preprint arXiv:2311.10988*.
- Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024. YOLO-World: Real-time open-vocabulary object detection. In *CVPR*.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slow-fast networks for video recognition. In *ICCV*, 6202–6211.
- Gao, J.; Sun, C.; Yang, Z.; and Nevatia, R. 2017. Tall: Temporal activity localization via language query. In *ICCV*, 5267–5275.
- Gao, J.; and Xu, C. 2021. Fast video moment retrieval. In *ICCV*, 1523–1532.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *ICCV*, 2961–2969.
- Jang, J.; Park, J.; Kim, J.; Kwon, H.; and Sohn, K. 2023. Knowing where to focus: Event-aware transformer for video grounding. In *ICCV*, 13846–13856.
- Ji, J.; Krishna, R.; Fei-Fei, L.; and Niebles, J. C. 2020. Action genome: Actions as compositions of spatio-temporal scene graphs. In *CVPR*, 10236–10247.
- Jiang, Y.; Zhang, W.; Zhang, X.; Wei, X.-Y.; Chen, C. W.; and Li, Q. 2024. Prior Knowledge Integration via LLM Encoding and Pseudo Event Regulation for Video Moment Retrieval. In *ACMMM*, 7249–7258.
- Kenton, J. D. M.-W. C.; and Toutanova, L. K. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2.
- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123: 32–73.
- Lee, P.; and Byun, H. 2025. Bam-detr: Boundary-aligned moment detection transformer for temporal sentence grounding in videos. In *ECCV*, 220–238.
- Lei, J.; Berg, T. L.; and Bansal, M. 2021. Detecting moments and highlights in videos via natural language queries. In *NeurIPS*, 11846–11858.
- Li, Z.; Chai, Y.; Zhuo, T. Y.; Qu, L.; Haffari, G.; Li, F.; Ji, D.; and Tran, Q. H. 2023. FACTUAL: A Benchmark for Faithful and Consistent Textual Scene Graph Parsing. In *ACL*.
- Lin, K. Q.; Zhang, P.; Chen, J.; Pramanick, S.; Gao, D.; Wang, A. J.; Yan, R.; and Shou, M. Z. 2023. Univtg: Towards unified video-language temporal grounding. In *ICCV*, 2794–2804.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *CVPR*, 2117–2125.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollar, P. 2017b. Focal Loss for Dense Object Detection. In *ICCV*, 2980–2988.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024a. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *ECCV*, 38–55.
- Liu, Y.; He, J.; Li, W.; Kim, J.; Wei, D.; Pfister, H.; and Chen, C. W. 2024b. R^2 -Tuning: Efficient Image-to-Video Transfer Learning for Video Temporal Grounding. *arXiv preprint arXiv:2404.00801*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 10012–10022.
- Lu, C.; Chen, L.; Tan, C.; Li, X.; and Xiao, J. 2019. Debug: A dense bottom-up grounding approach for natural language video localization. In *EMNLP*, 5144–5153.
- Moon, W.; Hyun, S.; Lee, S.; and Heo, J.-P. 2023a. Correlation-guided query-dependency calibration in video representation learning for temporal grounding. *arXiv preprint arXiv:2311.08835*.
- Moon, W.; Hyun, S.; Park, S.; Park, D.; and Heo, J.-P. 2023b. Query-dependent video representation for moment retrieval and highlight detection. In *CVPR*, 23023–23033.
- Mu, F.; Mo, S.; and Li, Y. 2024. SnAG: Scalable and Accurate Video Grounding. In *CVPR*, 18930–18940.
- Mun, J.; Cho, M.; and Han, B. 2020. Local-global video-text interactions for temporal grounding. In *CVPR*, 10810–10819.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763.
- Regneri, M.; Rohrbach, M.; Wetzel, D.; Thater, S.; Schiele, B.; and Pinkal, M. 2013. Grounding action descriptions in videos. *TACL*, 1: 25–36.
- Soldan, M.; Xu, M.; Qu, S.; Tegner, J.; and Ghanem, B. 2021. Vlg-net: Video-language graph matching network for video grounding. In *ICCV*, 3224–3234.
- Tian, Z.; Shen, C.; Chen, H.; and He, T. 2019. FCOS: Fully Convolutional One-Stage Object Detection. In *ICCV*, 9626–9635.

Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 4489–4497.

Wang, L.; Huang, B.; Zhao, Z.; Tong, Z.; He, Y.; Wang, Y.; Wang, Y.; and Qiao, Y. 2023. Videomae v2: Scaling video masked autoencoders with dual masking. In *CVPR*, 14549–14560.

Xiao, J.; Zhou, P.; Chua, T.-S.; and Yan, S. 2022. Video graph transformer for video question answering. In *ECCV*, 39–58.

Xiao, S.; Chen, L.; Zhang, S.; Ji, W.; Shao, J.; Ye, L.; and Xiao, J. 2021. Boundary proposal network for two-stage natural language video localization. In *AAAI*, 2986–2994.

Xiao, Y.; Luo, Z.; Liu, Y.; Ma, Y.; Bian, H.; Ji, Y.; Yang, Y.; and Li, X. 2024. Bridging the gap: A unified video comprehension framework for moment retrieval and highlight detection. In *CVPR*, 18709–18719.

Xu, Y.; Sun, Y.; Zhai, B.; Jia, Y.; and Du, S. 2024. Mh-detr: Video moment and highlight detection with cross-modal transformer. In *IJCNN*, 1–8.

Yan, S.; Xiong, X.; Nagrani, A.; Arnab, A.; Wang, Z.; Ge, W.; Ross, D.; and Schmid, C. 2023. Unloc: A unified framework for video localization tasks. In *ICCV*, 13623–13633.

Zeng, R.; Xu, H.; Huang, W.; Chen, P.; Tan, M.; and Gan, C. 2020. Dense regression network for video grounding. In *CVPR*, 10287–10296.

Zeng, Y.; Cao, D.; Wei, X.; Liu, M.; Zhao, Z.; and Qin, Z. 2021. Multi-modal relational graph for cross-modal video moment retrieval. In *CVPR*, 2215–2224.

Zeng, Y.; Zhong, Y.; Feng, C.; and Ma, L. 2025. Unimd: Towards unifying moment retrieval and temporal action detection. In *ECCV*, 286–304.

Zhang, C.-L.; Wu, J.; and Li, Y. 2022. Actionformer: Localizing moments of actions with transformers. In *ECCV*, 492–510. Springer.

Zhang, H.; Sun, A.; Jing, W.; and Zhou, J. T. 2020a. Span-based Localizing Network for Natural Language Video Localization. In *ACL*, 6543–6554.

Zhang, Y.; Gao, J.; Yang, X.; Liu, C.; Li, Y.; and Xu, C. 2020b. Find objects and focus on highlights: Mining object semantics for video highlight detection via graph neural networks. In *AAAI*, 12902–12909.

Zhang, Y.; Sun, P.; Jiang, Y.; Yu, D.; Weng, F.; Yuan, Z.; Luo, P.; Liu, W.; and Wang, X. 2022. Bytetrack: Multi-object tracking by associating every detection box. In *ECCV*, 1–21.

Zheng, Z.; Wang, P.; Liu, W.; Li, J.; Ye, R.; and Ren, D. 2020. Distance-IoU loss: Faster and better learning for bounding box regression. In *AAAI*, 12993–13000.