

Query-centric Audio-Visual Cognition Network for Moment Retrieval, Segmentation and Step-Captioning

Yunbin Tu¹, Liang Li^{2,1*}, Li Su^{1,3*}, Qingming Huang¹

¹ School of Computer Science and Technology, University of Chinese Academy of Sciences, Beijing, China

² Key Laboratory of AI Safety of CAS, Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

³ Peng Cheng Laboratory, Shenzhen, China

tuyunbin22@mails.ucas.ac.cn, liang.li@ict.ac.cn, {suli,qmhuang}@ucas.ac.cn

Abstract

Video has emerged as a favored multimedia format on the internet. To better gain video contents, a new topic HIREST is presented, including video retrieval, moment retrieval, moment segmentation, and step-captioning. The pioneering work chooses the pre-trained CLIP-based model for video retrieval, and leverages it as a feature extractor for other three challenging tasks solved in a multi-task learning paradigm. Nevertheless, this work struggles to learn the comprehensive cognition of user-preferred content, due to disregarding the hierarchies and association relations across modalities. In this paper, guided by the shallow-to-deep principle, we propose a query-centric audio-visual cognition (QUAG) network to construct a reliable multi-modal representation for moment retrieval, segmentation and step-captioning. Specifically, we first design the modality-synergistic perception to obtain rich audio-visual content, by modeling global contrastive alignment and local fine-grained interaction between visual and audio modalities. Then, we devise the query-centric cognition that uses the deep-level query to perform the temporal-channel filtration on the shallow-level audio-visual representation. This can cognize user-preferred content and thus attain a query-centric audio-visual representation for three tasks. Extensive experiments show QUAG achieves the SOTA results on HIREST. Further, we test QUAG on the query-based video summarization task and verify its good generalization.

Code — <https://github.com/tuyunbin/QUAG>

Extended version — <https://arxiv.org/pdf/2412.13543>

Introduction

Recently, we have been witnessing an exponential growth of videos, along with the advance of generative AI (e.g., Sora) and video platforms. Accordingly, a plethora of studies (Wu et al. 2023; Li et al. 2024; Wang et al. 2024c) have been proposed to enhance video retrieval capabilities. However, given a text query, returning the whole video is not always satisfied. Sometimes, users want to directly localize the moment most related to the query. For instance, when learning how to make a strawberry pie, users tend to only focus on the instructional moment, as shown in Figure 1. Moreover, if this moment is intricate, we anticipate machines to segment

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

*Corresponding authors.

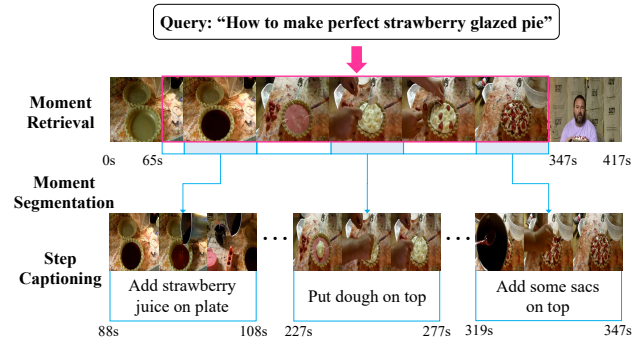


Figure 1: The illustrative example consisting of moment retrieval, moment segmentation, and step-captioning. First, given a text query “How to make perfect strawberry glazed pie”, the model is required to localize the most related moment in the video (moment retrieval). Then, the model proceeds to break down the moment into finer-level steps (moment segmentation). Finally, the model should describe each step with a concise sentence (step-captioning).

it into finer-level steps, and caption each with a succinct sentence for a better understanding, such as “Add strawberry juice on plate”, “Put dough on top”, etc. This shift in user preference sparks an emerging research topic: hierarchical retrieval and step-captioning (HIREST) (Zala et al. 2023).

This topic benchmarks four tasks of video retrieval, moment retrieval, moment segmentation, and step captioning. In the pioneering work (Zala et al. 2023), Zala *et al.* implement video retrieval by the pre-trained CLIP-based model and use it as a feature extractor for the three downstream tasks. Their focus is to jointly address the three tasks in a single architecture. Toward this end, they propose a *Joint* model that first leverages pre-trained models (Fang et al. 2023; Radford et al. 2023; Reimers and Gurevych 2019) to produce the multi-modal representation for visual frames, audio and query. Next, this representation is used to predict the boundaries of moments and steps, as well as yielding each step caption. In this process, the model is trained via a multi-task setup in a round-robin way (Cho et al. 2021).

Despite the promising results, the above work has the limitations in learning a comprehensive cognition of user-

preferred video content, owing to the naive multi-modal fusion strategy. First, the representations of three modalities are directly fused without distinction, which disregards the hierarchies across modalities. Studies in psychology (Tacca 2011; Yang, Zhuang, and Pan 2021) have revealed that human perception and cognition is in a shallow-to-deep fashion. When watching a video, we usually first obtain shallow-level sensory information, *e.g.*, the appearance and sound of objects, to obtain an intuitive perception for audio-visual content in the video. Then, we incorporate deeper-level knowledge (*e.g.*, intentions) into the sensory information to learn the comprehensive cognition of the particular content of interest. Hence, it is beneficial to model the modality hierarchies based on this shallow-to-deep principle.

Second, the fusion strategies (element-wise multiplication and summation) among three modalities fail to make full use of their association relations. The shallow-level visual and audio modalities can represent video content from different aspects, but direct element-wise summation may lose their synergistic relation. Afterwards, the deep-level query can filter trivial details and highlight the important ones within the shallow-level audio-visual content. Nevertheless, element-wise multiplication cannot model such a filtration relation that helps cognize user-preferred video content. As such, it is warranted to progressively capture these association relations during modeling hierarchies across modalities.

In this paper, we propose a **QUery-centric Audio-visual coGnition (QUAG)** network to cognize user-preferred video content and learn an effective multi-modal representation for addressing moment retrieval, segmentation and step-captioning. QUAG consists of modality-synergistic perception (MSP) and query-centric cognition (QC²) modules. Concretely, given visual frame and audio representations, MSP first models their global contrastive alignment to make them reside in the same embedding space; then learns their local fine-grained interaction to mine their joint representations, which are fused as the audio-visual representation. Afterwards, guided by the deep-level query, QC² performs a temporal-channel filtration on the shallow-level audio-visual representation, thus highlighting user-requested details. Next, the query representation is injected into the filtered audio-visual representation to construct a query-centric audio-visual representation. This is finally fed into the multi-modal encoder, multi-task prediction heads and text decoder for addressing the three challenging tasks.

Our key contributions are summarized as follows:

- Based on the investigation for human perception and cognition, we follow the shallow-to-deep principle to propose QUAG, which jointly solves three challenging tasks via learning a query-centric audio-visual representation.
- In QUAG, we first design MSP to attain the audio-visual representation by modeling global contrastive alignment and local fine-grained interaction between visual and audio representations. Then, we devise QC² to attend to the query-centric audio-visual representation, by implementing temporal-channel filtration from the deep-level query to the shallow-level audio-visual representation.
- Extensive experiments show QUAG achieves state-of-

the-art results for the moment retrieval, segmentation and step-captioning on the HIREST dataset. Moreover, we test QUAG on the TVSum dataset for the query-based video summarization, which validates its generalization.

Related Work

The task of hierarchical retrieval and step-captioning belongs to the multi-modal learning community (Cong et al. 2022, 2023; Tu et al. 2023a, 2024a,b,c). In the following, we review the related works from four dimensions.

Video-moment retrieval. Cross-modal retrieval is a fundamental problem in artificial intelligence (Zha et al. 2019; Zhang et al. 2020; Liu et al. 2022a; Zhang et al. 2024; Yue et al. 2023; Wang et al. 2024a,b; Tang et al. 2024; Yue et al. 2024). Text-based video retrieval (Dong et al. 2022; Xie et al. 2024; Li et al. 2023) is to find related videos from a corpora by a text query. Recently, benefiting from the contrastive learning, most text-to-image/video retrieval models (*e.g.*, CLIP) have been designed (Radford et al. 2021; Fang et al. 2023). With these models, the cosine similarity between the text query and video can be easily computed, thus finding the most related video to the query. However, since there are some query-irrelevant parts in the whole video, moment retrieval has been studied by most methods (Sun et al. 2022; Lei, Berg, and Bansal 2021; Moon et al. 2023), which is to localize the query-related span in the video.

Video summarization. Traditional video summarization methods (Song et al. 2015; Xiong and Grauman 2014) are to condense lengthy videos by extracting important information. However, these methods ignore the users' various preferences over the summaries. To address this limitation, the query-focused approaches (Sharghi, Gong, and Shah 2016; Sharghi, Laurel, and Gong 2017; Narasimhan et al. 2022) incorporate users' preferences through text queries, thus identifying the most relevant frames within the video. In other words, these query-focused methods can find a collection of moments that are most related to the users' preferences.

Video Captioning. Traditional methods (Li et al. 2022; Tu et al. 2023b) generate a concise sentence for a short video. To describe longer videos, dense video captioning is studied by some methods (Yang et al. 2023; Kim et al. 2024), which is to generate a paragraph for the long video. On the other hand, there are some methods (Lei et al. 2020; Li et al. 2020; Tu et al. 2022) that try to study TV show captioning. Instead of only using visual modalities, they consider introducing text features (*e.g.*, subtitles from the actors' dialogues) to augment video features, thus modeling a visual-linguistic representation for caption generation.

Hierarchical Retrieval and Step-Captioning. The above tasks are individually studied by previous works. In fact, these tasks share a common goal to extract information from a video corpus. Thus, Zala *et al.* (Zala et al. 2023) combine these tasks as a new task called HIREST to cater to users' various preferences. In this pioneering work, the pre-trained CLIP-based model is leveraged for video retrieval, and used as the feature extractor for the other three tasks. Further, a unified framework is proposed to jointly address the three tasks. In this work, we follow this paradigm to address the three tasks in a unified architecture. Compared

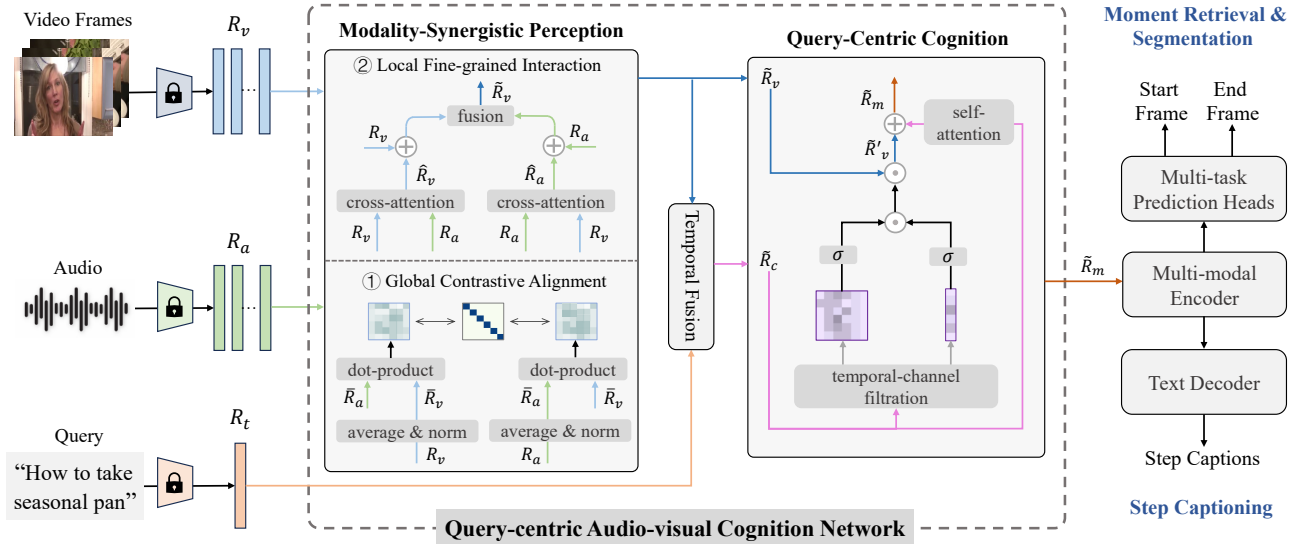


Figure 2: The overview of our method. Based on the principle of shallow-to-deep, we propose a query-centric audio-visual cognition (QUAG) network, where the core modules are the **modality-synergistic perception** and **query-centric cognition**. QUAG aims to learn a comprehensive cognition of user-preferred video content, and thus attain a query-centric audio-visual representation for jointly addressing the moment retrieval, moment segmentation, and step-captioning.

to the pioneering work that directly fuses multiple modalities by element-wise summation/multiplication, we learn a query-centric audio-visual representation by modeling the hierarchies and association relations across modalities, so as to better solve the three downstream tasks.

Methodology

As shown in Figure 2, the overall framework of our method includes the following parts. First, given the visual frames, audio, and query, we first extract the visual, audio, and text representations by pre-trained models. Then, we feed them into QUAG to produce a query-centric audio-visual representation. Finally, this representation is fed into a multi-modal encoder for intra-relation modeling, and used to 1) predict the boundaries of moments and steps, and 2) prompt the text decoder to generate the caption for each step.

Multi-modal Input Embedding

Given an untrimmed video containing N_v frames and a text query consisting of N_t tokens, we first employ pre-trained models to extract the visual representation R_v and text representation R_t , separately. Then, considering that the audio information helps perceive the main objects in the video, we also extract the speech transcription from the audio and embed it as the audio representation R_a , whose length is equal to the video representation, *i.e.*, N_v . Next, we project three modality representations into the same embedding space by three linear transformation functions, *i.e.*, $R_v \in \mathbb{R}^{N_v \times D}$, $R_a \in \mathbb{R}^{N_v \times D}$, $R_t \in \mathbb{R}^D$.

Query-centric Audio-visual Cognition

After obtaining the visual, audio, and text representations, we propose a query-centric audio-visual cognition (QUAG)

network that utilizes the modality-synergistic perception (MSP) and query-centric cognition (QC²) modules to learn a query-centric audio-visual representation based on the shallow-to-deep principle, so as to learn the comprehensive cognition of user-preferred video content.

Modality-Synergistic Perception Given the pair-wise visual representation R_v and audio representation R_a , we design the MSP to first make them reside in the same embedding space by maximizing their global contrastive alignment. Specifically, we compute their global features by the mean-pooling operation over the length dimension:

$$\bar{R}_v = \frac{1}{N_v} \sum_{N_v} R_v, \bar{R}_a = \frac{1}{N_v} \sum_{N_v} R_a, \quad (1)$$

Where $\bar{R}_v \in \mathbb{R}^D$ and $\bar{R}_a \in \mathbb{R}^D$. Then, in a training batch, we sample B pair-wise globally visual and audio features. For k -th globally visual feature \bar{R}_v^k , k -th globally audio feature \bar{R}_a^k is its positive, while other r ($r \neq k$) globally audio features will be the negatives. That is, there are a total of B positive sample pairs and $B \times (B - 1)$ negative sample pairs in this batch. Next, we employ the InfoNCE loss (Oord, Li, and Vinyals 2018) to maximize the bi-directional similarities between the positive pairs \bar{R}_v^k and \bar{R}_a^k , while minimizing the similarity between the negative pairs:

$$\begin{aligned} \mathcal{L}_{v2a} &= -\frac{1}{B} \sum_k \log \frac{e^{(\text{sim}(\bar{R}_v^k, \bar{R}_a^k)/\tau)}}{\sum_r e^{(\text{sim}(\bar{R}_v^k, \bar{R}_r^k)/\tau)}}, \\ \mathcal{L}_{a2v} &= -\frac{1}{B} \sum_k \log \frac{e^{(\text{sim}(\bar{R}_a^k, \bar{R}_v^k)/\tau)}}{\sum_r e^{(\text{sim}(\bar{R}_a^k, \bar{R}_r^k)/\tau)}}, \\ \mathcal{L}_{msp} &= \frac{1}{2} (\mathcal{L}_{v2a} + \mathcal{L}_{a2v}), \end{aligned} \quad (2)$$

where τ is the temperature hyper-parameter. ‘‘sim’’ means the dot-product operation. This loss formulates a self-supervisory signal to enhance the global alignment between the visual and audio representations, which facilitates the subsequent local fine-grained interaction.

Subsequently, we interact the local features on the visual representation R_v and audio representation R_a to learn their fine-grained synergy and mine the joint representations. This is performed by the multi-head cross-attention mechanism (MHCA) (Vaswani et al. 2017):

$$\begin{aligned}\hat{R}_v &= \text{MHCA}(R_v, R_a, R_a), \\ \hat{R}_a &= \text{MHCA}(R_a, R_v, R_v).\end{aligned}\quad (3)$$

In the bracket of each equation, the first term represents the query, while the last two terms denote the key and value representations. After that, we concatenate the joint representations from two modalities as the audio-visual representation, which is performed by a fully-connected layer:

$$\tilde{R}_v = \left([\hat{R}_v; \hat{R}_a] W_c + b_c \right), \quad (4)$$

where $\tilde{R}_v \in \mathbb{R}^{N_v \times D}$. $W_c \in \mathbb{R}^{2D \times D}$ and $b_c \in \mathbb{R}^D$. $[\cdot]$ is a concatenation operation.

Query-Centric Cognition We devise QC² to use the deep-level query representation $R_t \in \mathbb{R}^D$ to attend to the related information over the shallow-level audio-visual representation $\tilde{R}_v \in \mathbb{R}^{N_v \times D}$. QC² first measures the relevance between R_t and \tilde{R}_v from both temporal and channel dimensions. The temporal and channel relation matrices \mathbf{A}_{te} and \mathbf{A}_{ch} are computed as follows:

$$\begin{aligned}\tilde{R}_c &= \left([\tilde{R}_v; R_t] W_d + b_d \right), \\ \mathbf{A}_{te} &= \text{sigmoid} \left(\mathcal{F}_{te} \left(\frac{1}{D} \sum_j \tilde{R}_c(:, j) \right) \right), \\ \mathbf{A}_{ch} &= \text{sigmoid} \left(\mathcal{F}_{ch} \left(\frac{1}{N_v} \sum_j \tilde{R}_c(j, :) \right) \right),\end{aligned}\quad (5)$$

where $W_d \in \mathbb{R}^{2D \times D}$ and $b_d \in \mathbb{R}^D$. $[\cdot]$ is a concatenation operation and we broadcast the text representation $R_t \in \mathbb{R}^D$ as $R_t \in \mathbb{R}^{N_v \times D}$, to temporally match its shape with the audio-visual representation $\tilde{R}_v \in \mathbb{R}^{N_v \times D}$. $\mathbf{A}_{te} \in \mathbb{R}^{N_v \times 1}$, $\mathbf{A}_{ch} \in \mathbb{R}^{1 \times D}$. $\mathcal{F}_{te} \in \mathbb{R}^{1 \times 1}$ and $\mathcal{F}_{ch} \in \mathbb{R}^{D \times D}$ are linear transformation functions.

Next, QC² computes the temporal-channel relation matrix by fusing \mathbf{A}_{te} and \mathbf{A}_{ch} together, which is implemented with the element-wise multiplication function:

$$\mathbf{A}_{tc} = \mathbf{A}_{te} \odot \mathbf{A}_{ch}, \quad (6)$$

where $\mathbf{A}_{tc} \in \mathbb{R}^{N_v \times D}$. Guided by \mathbf{A}_{tc} , the model can filter the trivial details in the audio-visual representation \tilde{R}_v :

$$\tilde{R}'_v = \mathbf{A}_{tc} \odot \tilde{R}_v, \quad (7)$$

where $\tilde{R}'_v \in \mathbb{R}^{N_v \times D}$. Subsequently, we integrate \tilde{R}_c into the filtered audio-visual representation \tilde{R}'_v to model a query-centric audio-visual representation:

$$\tilde{R}_m = \tilde{R}'_v + \phi(\tilde{R}_c), \quad (8)$$

where ϕ refers to the multi-head self-attention. Finally, \tilde{R}_m is fed into a transformer-based multi-modal encoder for enhancement, and then is utilized for the moment retrieval, moment segmentation and step-captioning.

Moment Retrieval, Moment Segmentation, and Step-captioning

Moment Retrieval is to find the moment most relevant to the query and output its span. Specifically, given \tilde{R}_m , we use a learnable prediction head consisting of two linear layers to predict the start and end boundaries at the same time, where the frame inputs are not masked. The probability distributions of the start and end indexes over the entire video are computed as follows:

$$\begin{aligned}\mathbf{P}_{\text{start}} &= \text{Softmax} \left(\tilde{R}_m W_s + b_s \right), \\ \mathbf{P}_{\text{end}} &= \text{Softmax} \left(\tilde{R}_m W_e + b_e \right),\end{aligned}\quad (9)$$

where $W_s \in \mathbb{R}^{D \times 1}$, $W_e \in \mathbb{R}^{D \times 1}$, $b_s \in \mathbb{R}^1$, and $b_e \in \mathbb{R}^1$. **Moment Segmentation** refers to identifying ‘‘key steps’’ within the retrieved moment that does not have repetitive frames, where each step’s end timestamp is the next step’s start timestamp. Concretely, given \tilde{R}_m , we use a learnable prediction head with a linear layer to predict the boundary of each step in an auto-regressive manner, where the previously predicted boundaries are integrated into \tilde{R}_m as the part of the input. The frames outside of the moment and in the previous steps are masked out. The probability distribution of boundary index for each step is computed as follows:

$$\mathbf{P}_{\text{step}} = \text{Softmax} \left(\tilde{R}_m W_t + b_t \right), \quad (10)$$

where $W_t \in \mathbb{R}^{D \times 1}$ and $b_t \in \mathbb{R}^1$.

Step-captioning aims to describe each segmented step with a sentence. To this end, we first use the pre-trained captioning model CLIP4Caption (Tang et al. 2021) to initialize the parameters of text decoder. Then, we feed \tilde{R}_m into the text decoder that yields the text descriptions through an auto-regressive way. This process can be formulated as follows:

$$\hat{\mathbf{w}}_i = \mathcal{F}_{De} \left(\tilde{R}_m, \mathbf{w}_1, \dots, \mathbf{w}_{i-1} \right), \quad (11)$$

where $\hat{\mathbf{w}}_i \in \mathbb{R}^{N_t}$ is the i -th decoded distribution over a dictionary of N_t vocabularies; $\{\mathbf{w}_1, \dots, \mathbf{w}_{i-1}\}$ are previously generated words (ground-truth words during training and predicted words during inference).

Training. For moment retrieval, given the ground-truth start and end indices, we define the training loss (to be minimized) as the sum of the negative log probabilities by the predicted distributions, which are then averaged over all examples in a batch:

$$\mathcal{L}_{ret} = -\frac{1}{B} \sum_i \log \left(\mathbf{p}_{start}^{y_i^{start}} \right) - \frac{1}{B} \sum_i \log \left(\mathbf{p}_{end}^{y_i^{end}} \right), \quad (12)$$

where B is the batch size; y_i^{start} and y_i^{end} are the ground-truth start and end indices of the i -th example, respectively.

For moment segmentation, given ground-truth step indices, we define the training loss (to be minimized) as the sum of the negative log probabilities by the predicted distributions, which are averaged over all examples in a batch:

$$\mathcal{L}_{seg} = -\frac{1}{B} \sum_i^B \log \left(\mathbf{p}_{step}^{y_i^{step}} \right), \quad (13)$$

where B is the batch size and y_i^{step} is the ground-truth step index of the i -th example, respectively.

For step-captioning, given the ground-truth step caption words (w_1^*, \dots, w_m^*) , we define the training loss (to be minimized) as the sum of the negative log probabilities by the predicted distributions, which are then averaged over all examples in a batch:

$$\mathcal{L}_{cap} = -\sum_{i=1}^m \log p_{\theta}(w_i^* | w_{<i}^*) \quad (14)$$

where $p_{\theta}(w_i^* | w_{<i}^*)$ is computed by Eq. (11). m is the length of a step caption.

Our method is trained in a multi-task configuration, where a round-robin manner is leveraged and a batch is sampled from one of the three data loaders at each iteration (Zala et al. 2023; Cho et al. 2021). That is, the training loss for each iteration is defined as:

$$\mathcal{L}(\theta) = \mathcal{L}_i + \lambda \mathcal{L}_{msp}, \quad (15)$$

where i refers to *ret*, *seg*, or *cap*. θ are a set of trainable weights. λ is a trade-off parameter to balance the contribution between one of the task losses and the contrastive loss, which is discussed in the supplementary material*.

Experiments

Datasets

HIREST consists of the tasks of video retrieval, moment retrieval, moment segmentation, and step-captioning. It is comprised of 3.4K text-video pairs, 1.8K moments, and 8.6K step captions. We use the official split with 1,507 video-query pairs for training, 477 video-query pairs for validation and 1,391 video-query pairs for testing.

TVSum contains the task of query-based video summarization, which is relevant to moment segmentation. This dataset includes 10 varying categories of videos, and each category consists of 5 videos. For a fair-comparison, we follow QD-DETR (Moon et al. 2023) to utilize 80% videos for training and the remaining for testing.

Evaluation Metrics

HIREST: (1) We validate moment retrieval by evaluating the outputs of model against the ground-truth moment spans via Recall of Intersection over Union (IoU) thresholds (0.5 and 0.7). (2) For moment segmentation, models are assessed based on the similarity between the generated step spans and the ground-truth step spans, via Recall and Precision metrics over IoU thresholds (0.5 and 0.7). (3) For

*The supplementary material is in the extended version.

step-captioning, the traditional metrics are used to evaluate the yielded sentence: METEOR (Banerjee and Lavie 2005), ROUGE-L (Lin 2004), CIDEr (Vedantam, Lawrence Zitnick, and Parikh 2015), and SPICE (Anderson et al. 2016). Besides, Following (Zala et al. 2023), we introduce the ELMo (Peters et al. 2018)-based Decomposable Attention model (Parikh et al. 2016) to compute the entailment of generated sentences against the ground-truth captions.

TVSum: Following previous works (Moon et al. 2023; Liu et al. 2022b), we leverage the top-5 mAP as the main metric.

Implementation Details

First, we use pre-trained EVA-CLIP (Fang et al. 2023) to extract the visual representation for visual frames, and the text encoder of EVA-CLIP to map a text query as the text representation. Then, we use pre-trained Whisper (Radford et al. 2023) to extract speech transcription from audio, and use pre-trained MiniLM (Reimers and Gurevych 2019) text encoder to map the transcription into a representation, called audio representation. The hidden size is set to 768. The text decoder is first initialized from pre-trained CLIP4Caption (Tang et al. 2021), and then is fine-tuned on the HIREST dataset. During training, the batch size is set to 5 and learning rate is set to 1×10^{-5} , AdamW optimizer (Loshchilov and Hutter 2018) is used to minimize the training loss defined in Eq. (15). More details are shown in supplementary.

Performance Comparison on HIREST

Results on the Moment Retrieval In this task, we compare QUAG with the pioneering work *Joint* model (Zala et al. 2023). We also compare QUAG with three task-specific moment retrieval models: QD-DETR (Moon et al. 2023), TR-DETR (Sun et al. 2024), and UVCOM (Xiao et al. 2024). The mentioned models are implemented based on their released codes. Besides, following Zala et al. (Zala et al. 2023), we compare QUAG with a text-to-image retrieval model EVA-CLIP (Fang et al. 2023).

The comparison results are shown in Table 1. Our QUAG obtains the best performances in moment retrieval at both metrics. These indicate our method has the high recall in identifying relevant moments with the 50% overlap and 70% overlap with the ground-truth spans. Especially, our method has the performance improvements of 2.2% and 4.2% against the *Joint* model at R@0.5 and R@0.7, respectively, which further validates the effectiveness of QUAG.

Model	Recall @ IoU	
	0.5	0.7
EVA-CLIP-G/14 (Fang et al. 2023)	38.27	19.33
QD-DETR (Moon et al. 2023)	71.52	38.34
UVCOM (Xiao et al. 2024)	<u>72.12</u>	<u>38.53</u>
TR-DETR (Sun et al. 2024)	72.07	38.40
<i>Joint</i> (Zala et al. 2023)	70.98	37.31
QUAG	72.54	38.86

Table 1: Moment retrieval results on the HIREST test set.

Model	Recall @ IoU		Precision @ IoU	
	0.5	0.7	0.5	0.7
QD-DETR	34.24	16.27	28.32	12.05
UVCOM	37.03	16.79	<u>30.38</u>	<u>13.80</u>
TR-DETR	37.07	16.35	29.71	12.93
<i>Joint</i>	<u>36.23</u>	<u>16.05</u>	28.52	12.84
QUAG	39.27	17.35	31.68	14.81

Table 2: Moment segmentation results on the test set of the HIREST dataset.

Model	M	R	C	S	Entail. (%)
BMT	3.84	-	6.72	1.05	30.68
SwinBERT	5.94	-	24.66	6.67	35.09
<i>Joint</i>	4.14	11.85	21.19	3.02	35.97
QUAG	<u>4.76</u>	13.20	25.44	<u>4.49</u>	40.10

Table 3: Step-captioning results on the HIREST test set. M, R, C, S indicate METEOR, ROUGE-L, CIDEr, and SPICE.

Results on the Moment Segmentation In this task, we compare the proposed QUAG with the *Joint* model (Zala et al. 2023). Also, three task-specific models QD-DETR (Moon et al. 2023), TR-DETR (Sun et al. 2024), and UVCOM (Xiao et al. 2024) are implemented based on their codes and compared with our QUAG.

The results are shown in Table 2. The QUAG excels in Recall at IoU thresholds of 0.5 and 0.7. QUAG’s high recall rates at both IoU thresholds shows its strong ability for moment segmentation. For the precision metric, our QUAG obtains the best result at IoU thresholds of both 0.5 and 0.7, which indicates our method can segment a moment with a high degree of accuracy. Overall, the superior results on both recall and precision demonstrate QUAG’s effectiveness in achieving accurate and reliable moment segmentation, which is essential for the subsequent step-captioning task.

Results on the Step-captioning In this task, we compare the proposed QUAG with the *Joint* model (Zala et al. 2023), BMT (Iashin and Rahtu 2020) that is a dense video captioning model pre-trained on the ActivityNet dataset (Krishna et al. 2017)), and SwinBERT (Lin et al. 2022) that is a pre-trained video captioning model on the YouCook2 dataset (Zhou, Xu, and Corso 2018). The comparison results are shown in Table 3.

It is noted that the overall performance of our QUAG is better in this task. Our QUAG outperforms the *Joint* model on every metric by a large margin. For two captioning-based models BMT and SwinBERT, our QUAG outperforms BMT on all metrics and achieves the matching performance against SwinBERT. Especially, on the caption-centric metric CIDEr, our QUAG obtains the relative improvements of 3.2% and 280% over SwinBERT and BMT, respectively. On the entailment score, QUAG also obtains the relative improvements of 14.3% and 30.7% over SwinBERT and BMT,

Model	MSP	QC ²	M	R	C	S
<i>Joint</i>	×	×	4.14	11.85	21.19	3.02
QUAG	✓	×	4.70	13.05	23.18	3.49
QUAG	×	✓	4.50	12.65	20.43	3.87
QUAG	✓	✓	4.76	13.20	25.44	4.49

Table 4: Ablation study for each module on step-captioning.

separately. The comparison results show the good generalization of our method for caption generation.

Overall Performance Analysis. Compared with the pioneering work *Joint* on HIREST, our method surpasses it in every metric for each sub-task, proving its effectiveness. To comprehensively evaluate our approach, it is also compared with some task-specific methods. We notice our method shows limited improvements in certain metrics of step-captioning, which is actually a foreseen situation. The ultimate goal of HIREST is to train a generic model that can generalize to several related video understanding tasks, so it could perform worse than certain task-specific methods.

Ablation Study

We conduct the ablation study to investigate the contribution of each module and the full model. Here, we choose the step-captioning task as the major ablation task, as it requires the model have a comprehensive capability for the cognition, alignment and generation for multi-modal content. The baseline is *Joint*, which directly fuses three modality representations by element-wise summation and multiplication.

The ablative results are shown in Table 4. The performances of each module and the full QUAG are better than that of *Joint* baseline, which validate their effectiveness. However, the performance improvement is not significant when each module is individually used. When only QC² is used, there is even the performance decreasing on the CIDEr metric. Our conjecture is that each module only learns either synergistic relation or filtration relations. Besides, the hierarchies across modalities are overlooked. When the full QUAG model is used, its performance significantly surpasses the others, in particular with the increases of 20.1% and 48.7% on the CIDEr and SPICE over the *Joint* baseline. The ablation results indicate each designed module not only plays its unique role, but also supplements each other.

Performance Comparison on TVsum

We further validate the generalization ability of QUAG for the query-based video summarization that is related to moment segmentation. The experiment is conducted on the TVsum dataset and the compared methods are TGG (Ye et al. 2021), UMT (Liu et al. 2022b), QD-DETR (Moon et al. 2023), UVCOM (Xiao et al. 2024), and TR-DETR (Sun et al. 2024). The inputs for all compared methods are visual, audio, and text query representations.

The results are shown in Table 5. QUAG obtains the best results on the 5 categories, which is better than the others. Besides, the average overall performance of our method is in par with TR-DETR (87.0 vs. 87.1), which shows its good

Model	VT	VU	GA	MS	PK	PR	FM	BK	BT	DS	Avg.
TGG (Ye et al. 2021)	85.0	71.4	81.9	78.6	80.2	75.5	71.6	77.3	78.6	68.1	76.8
UMT (Liu et al. 2022b)	87.5	81.5	88.2	78.8	81.4	<u>87.0</u>	76.0	86.9	84.4	<u>79.6</u>	83.1
QD-DETR (Moon et al. 2023)	87.6	91.7	90.2	88.3	84.1	88.3	<u>78.7</u>	91.2	87.8	<u>77.7</u>	86.6
UVCOM (Xiao et al. 2024)	87.6	91.6	91.4	<u>86.7</u>	<u>86.9</u>	86.9	<u>76.9</u>	92.3	87.4	75.6	86.3
TR-DETR (Sun et al. 2024)	90.6	92.4	91.7	81.3	<u>86.9</u>	85.5	79.8	93.4	<u>88.3</u>	81.0	87.1
QUAG	89.8	93.1	92.9	85.6	87.9	86.1	76.7	91.4	89.0	77.7	<u>87.0</u>

Table 5: Performance comparison on the TVsum dataset.

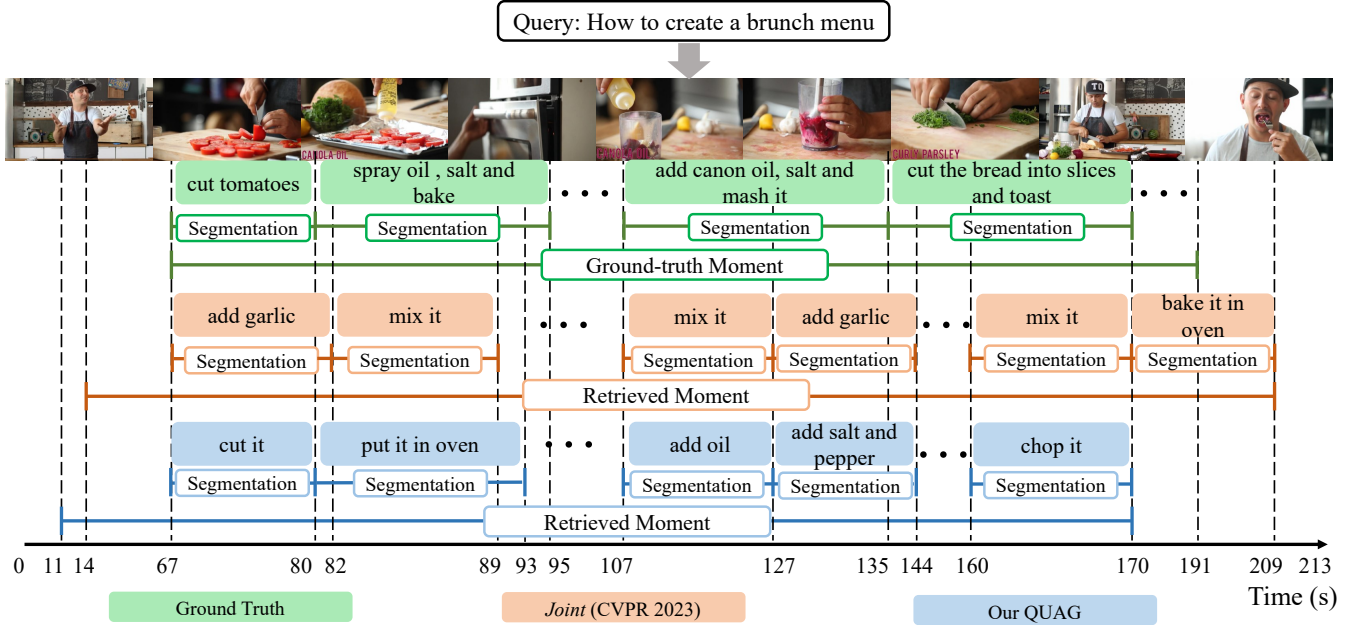


Figure 3: Given a text query “How to create a brunch menu”, and the ground-truth annotations, we compare the predicted and generated outputs of our QUAG and *Joint* (Zala et al. 2023) for moment retrieval, segmentation, and step-captioning.

generalization. The superior results mainly benefits from that our method construct a query-centric audio-visual representation. As such, the model can better learn which part of a video is important, and thus extracts the most relevant frames from a video based on the query.

Qualitative Analysis

To obtain an intuitive observation for our method, in Figure 3, we provide a visualization case for moment retrieval, segmentation, and step-captioning. We observe that the retrieved moment by our QUAG contains the less trivial details than that of the *Joint*. Besides, QUAG better predicts the step boundaries and generates the step captions. For instance, the yielded step captions by *Joint* model are redundant, which lack the details in the retrieved moment. By contrast, our QUAG generates “cut it” that matches the ground-truth caption “cut tomatoes” (67-80s); “add oil” and “add salt and pepper” that show the semantic consistency with the ground-truth caption “add canon oil, salt and mash it” (107-135s). Besides, our QUAG can describe fine-grained content of “chop it” during the actor cuts the bread into slices.

In short, the visualization results show the effectiveness of QUAG. The superiority benefits from the fact of progressively modeling a query-centric audio-visual representation, rather than directly fusing them without distinction. More examples are shown in the supplementary material.

Conclusion

This paper follows the shallow-to-deep principle to propose QUAG, which learns a query-centric audio-visual representation for moment retrieval, segmentation, and step-captioning. In QUAG, we first devise MSP to gain the audio-visual representation via modeling the global contrastive alignment and local fine-grained interaction between visual and audio modalities. Then, QC² is designed to learn a comprehensive cognition for the user-preferred content, by modeling temporal-channel filtration from the deep-level query to shallow-level audio-visual content. Extensive experiments show QUAG achieves SOTA performances for the three challenging tasks on the HIREST dataset. Moreover, we test QUAG for query-based video summarization task. The results verify its good generalization capability.

Acknowledgements

This work was supported by National Natural Science Foundation of China: 62322211, 62336008, U21B2038, “Pioneer” and “Leading Goose” R&D Program of Zhejiang Province (2024C01023), Fundamental Research Funds for the Central Universities (E2ET1104), Key Laboratory of Intelligent Processing Technology for Digital Music (Zhejiang Conservatory of Music), Ministry of Culture and Tourism (2023DMKLB004).

References

- Anderson, P.; Fernando, B.; Johnson, M.; and Gould, S. 2016. Spice: Semantic propositional image caption evaluation. In *ECCV*, 382–398.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *ACL*, 65–72.
- Cho, J.; Lei, J.; Tan, H.; and Bansal, M. 2021. Unifying vision-and-language tasks via text generation. In *ICML*, 1931–1942.
- Cong, G.; Li, L.; Liu, Z.; Tu, Y.; Qin, W.; Zhang, S.; Yan, C.; Wang, W.; and Jiang, B. 2022. Ls-gan: iterative language-based image manipulation via long and short term consistency reasoning. In *ACM MM*, 4496–4504.
- Cong, G.; Li, L.; Qi, Y.; Zha, Z.-J.; Wu, Q.; Wang, W.; Jiang, B.; Yang, M.-H.; and Huang, Q. 2023. Learning to dub movies via hierarchical prosody models. In *CVPR*, 14687–14697.
- Dong, J.; Chen, X.; Zhang, M.; Yang, X.; Chen, S.; Li, X.; and Wang, X. 2022. Partially relevant video retrieval. In *ACM MM*, 246–257.
- Fang, Y.; Wang, W.; Xie, B.; Sun, Q.; Wu, L.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2023. Eva: Exploring the limits of masked visual representation learning at scale. In *CVPR*, 19358–19369.
- Iashin, V.; and Rahtu, E. 2020. A Better Use of Audio-Visual Cues: Dense Video Captioning with Bi-modal Transformer. In *BMVC*.
- Kim, M.; Kim, H. B.; Moon, J.; Choi, J.; and Kim, S. T. 2024. Do You Remember? Dense Video Captioning with Cross-Modal Memory Retrieval. In *CVPR*, 13894–13904.
- Krishna, R.; Hata, K.; Ren, F.; Fei-Fei, L.; and Carlos Niebles, J. 2017. Dense-captioning events in videos. In *ICCV*, 706–715.
- Lei, J.; Berg, T. L.; and Bansal, M. 2021. Detecting moments and highlights in videos via natural language queries. *NeurIPS*, 34: 11846–11858.
- Lei, J.; Yu, L.; Berg, T. L.; and Bansal, M. 2020. Tvr: A large-scale dataset for video-subtitle moment retrieval. In *ECCV*, 447–463.
- Li, H.; Yang, S.; Zhang, Y.; Tao, D.; and Yu, Z. 2023. Progressive Feature Mining and External Knowledge-Assisted Text-Pedestrian Image Retrieval. *arXiv preprint arXiv:2308.11994*.
- Li, L.; Chen, Y.-C.; Cheng, Y.; Gan, Z.; Yu, L.; and Liu, J. 2020. HERO: Hierarchical Encoder for Video+ Language Omni-representation Pre-training. In *EMNLP*, 2046–2065.
- Li, L.; Gao, X.; Deng, J.; Tu, Y.; Zha, Z.-J.; and Huang, Q. 2022. Long short-term relation transformer with global gating for video captioning. *IEEE Transactions on Image Processing*, 31: 2726–2738.
- Li, Q.; Su, L.; Zhao, J.; Xia, L.; Cai, H.; Cheng, S.; Tang, H.; Wang, J.; and Yin, D. 2024. Text-Video Retrieval via Multi-Modal Hypergraph Networks. In *WSDM*, 369–377.
- Lin, C.-Y. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, 74–81.
- Lin, K.; Li, L.; Lin, C.-C.; Ahmed, F.; Gan, Z.; Liu, Z.; Lu, Y.; and Wang, L. 2022. Swinbert: End-to-end transformers with sparse attention for video captioning. In *CVPR*, 17949–17958.
- Liu, X.; Li, L.; Wang, S.; Zha, Z.-J.; Li, Z.; Tian, Q.; and Huang, Q. 2022a. Entity-enhanced adaptive reconstruction network for weakly supervised referring expression grounding. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3): 3003–3018.
- Liu, Y.; Li, S.; Wu, Y.; Chen, C.-W.; Shan, Y.; and Qie, X. 2022b. Umt: Unified multi-modal transformers for joint video moment retrieval and highlight detection. In *CVPR*, 3042–3051.
- Loshchilov, I.; and Hutter, F. 2018. Decoupled Weight Decay Regularization. In *ICLR*.
- Moon, W.; Hyun, S.; Park, S.; Park, D.; and Heo, J.-P. 2023. Query-dependent video representation for moment retrieval and highlight detection. In *CVPR*, 23023–23033.
- Narasimhan, M.; Nagrani, A.; Sun, C.; Rubinstein, M.; Darrell, T.; Rohrbach, A.; and Schmid, C. 2022. Tl; dw? summarizing instructional videos with task relevance and cross-modal saliency. In *ECCV*, 540–557.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Parikh, A.; Täckström, O.; Das, D.; and Uszkoreit, J. 2016. A Decomposable Attention Model for Natural Language Inference. In *EMNLP*, 2249–2255.
- Peters, M. E.; Neumann, M.; Iyyer, M.; Gardner, M.; Clark, C.; Lee, K.; and Zettlemoyer, L. 2018. Deep Contextualized Word Representations. In *NAACL-HLT*, 2227–2237.
- 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.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2023. Robust speech recognition via large-scale weak supervision. In *ICML*, 28492–28518.
- Reimers, N.; and Gurevych, I. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *EMNLP-IJCNLP*, 3982–3992.
- Sharghi, A.; Gong, B.; and Shah, M. 2016. Query-focused extractive video summarization. In *ECCV*, 3–19.

- Sharghi, A.; Laurel, J. S.; and Gong, B. 2017. Query-focused video summarization: Dataset, evaluation, and a memory network based approach. In *CVPR*, 4788–4797.
- Song, Y.; Vallmitjana, J.; Stent, A.; and Jaimes, A. 2015. Tv-sum: Summarizing web videos using titles. In *CVPR*, 5179–5187.
- Sun, H.; Zhou, M.; Chen, W.; and Xie, W. 2024. Tr-detr: Task-reciprocal transformer for joint moment retrieval and highlight detection. In *AAAI*, volume 38, 4998–5007.
- Sun, X.; Wang, X.; Gao, J.; Liu, Q.; and Zhou, X. 2022. You need to read again: Multi-granularity perception network for moment retrieval in videos. In *SIGIR*, 1022–1032.
- Tacca, M. C. 2011. Commonalities between perception and cognition. *Frontiers in psychology*, 2: 358.
- Tang, M.; Wang, Z.; Liu, Z.; Rao, F.; Li, D.; and Li, X. 2021. Clip4caption: Clip for video caption. In *ACM MM*, 4858–4862.
- Tang, W.; Li, L.; Liu, X.; Jin, L.; Tang, J.; and Li, Z. 2024. Context Disentangling and Prototype Inheriting for Robust Visual Grounding. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 46(05): 3213–3229.
- Tu, Y.; Li, L.; Su, L.; Gao, S.; Yan, C.; Zha, Z.-J.; Yu, Z.; and Huang, Q. 2022. I2 Transformer: Intra-and inter-relation embedding transformer for TV show captioning. *IEEE Transactions on Image Processing*, 31: 3565–3577.
- Tu, Y.; Li, L.; Su, L.; Yan, C.; and Huang, Q. 2024a. Distractors-Immune Representation Learning with Cross-modal Contrastive Regularization for Change Captioning. In *ECCV*, 311–328.
- Tu, Y.; Li, L.; Su, L.; Zha, Z.-J.; and Huang, Q. 2024b. SMART: Syntax-Calibrated Multi-Aspect Relation Transformer for Change Captioning. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 46(07): 4926–4943.
- Tu, Y.; Li, L.; Su, L.; Zha, Z.-J.; Yan, C.; and Huang, Q. 2023a. Self-supervised cross-view representation reconstruction for change captioning. In *ICCV*, 2805–2815.
- Tu, Y.; Li, L.; Su, L.; Zha, Z.-J.; Yan, C.; and Huang, Q. 2024c. Context-aware Difference Distilling for Multi-change Captioning. In *ACL*, 7941–7956.
- Tu, Y.; Zhou, C.; Guo, J.; Li, H.; Gao, S.; and Yu, Z. 2023b. Relation-aware attention for video captioning via graph learning. *Pattern Recognition*, 136: 109204.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *NeurIPS*, 5998–6008.
- Vedantam, R.; Lawrence Zitnick, C.; and Parikh, D. 2015. Cider: Consensus-based image description evaluation. In *CVPR*, 4566–4575.
- Wang, L.; Huang, X.; Yu, Z.; Peng, H.; Gao, S.; Mao, C.; Huang, Y.; Dong, L.; and Yu, P. S. 2024a. Zero-Shot Text Normalization via Cross-Lingual Knowledge Distillation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32: 4631–4646.
- Wang, L.; Yu, Z.; Gao, S.; Mao, C.; and Huang, Y. 2024b. DETS: End-to-End Single-Stage Text-to-Speech Via Hierarchical Diffusion Gan Models. In *ICASSP*, 10916–10920.
- Wang, M.; Li, H.; Zhang, Y.; Li, J.; Xie, M.; and Tao, D. 2024c. Dual-task Mutual Reinforcing Embedded Joint Video Paragraph Retrieval and Grounding. *arXiv preprint arXiv:2411.17481*.
- Wu, W.; Luo, H.; Fang, B.; Wang, J.; and Ouyang, W. 2023. Cap4video: What can auxiliary captions do for text-video retrieval? In *CVPR*, 10704–10713.
- 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.
- Xie, M.; Wang, M.; Li, H.; Zhang, Y.; Tao, D.; and Yu, Z. 2024. Phrase Decoupling Cross-Modal Hierarchical Matching and Progressive Position Correction for Visual Grounding. *arXiv preprint arXiv:2410.23570*.
- Xiong, B.; and Grauman, K. 2014. Detecting snap points in egocentric video with a web photo prior. In *ECCV*, 282–298.
- Yang, A.; Nagrani, A.; Seo, P. H.; Miech, A.; Pont-Tuset, J.; Laptev, I.; Sivic, J.; and Schmid, C. 2023. Vid2seq: Large-scale pretraining of a visual language model for dense video captioning. In *CVPR*, 10714–10726.
- Yang, Y.; Zhuang, Y.; and Pan, Y. 2021. Multiple knowledge representation for big data artificial intelligence: framework, applications, and case studies. *Frontiers of Information Technology & Electronic Engineering*, 22(12): 1551–1558.
- Ye, Q.; Shen, X.; Gao, Y.; Wang, Z.; Bi, Q.; Li, P.; and Yang, G. 2021. Temporal cue guided video highlight detection with low-rank audio-visual fusion. In *ICCV*, 7950–7959.
- Yue, S.; Tu, Y.; Li, L.; Gao, S.; and Yu, Z. 2024. Multi-grained Representation Aggregating Transformer with Gating Cycle for Change Captioning. *ACM Transactions on Multimedia Computing, Communications and Applications*, 20(10): 321:1–321:23.
- Yue, S.; Tu, Y.; Li, L.; Yang, Y.; Gao, S.; and Yu, Z. 2023. I3n: Intra-and inter-representation interaction network for change captioning. *IEEE Transactions on Multimedia*, 25: 8828–8841.
- Zala, A.; Cho, J.; Kottur, S.; Chen, X.; Oguz, B.; Mehdad, Y.; and Bansal, M. 2023. Hierarchical video-moment retrieval and step-captioning. In *CVPR*, 23056–23065.
- Zha, Z.-J.; Liu, D.; Zhang, H.; Zhang, Y.; and Wu, F. 2019. Context-aware visual policy network for fine-grained image captioning. *IEEE transactions on pattern analysis and machine intelligence*, 44(2): 710–722.
- Zhang, B.; Li, L.; Wang, S.; Cai, S.; Zha, Z.-J.; Tian, Q.; and Huang, Q. 2024. Inductive State-Relabeling Adversarial Active Learning With Heuristic Clique Rescaling. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 46(12): 9780–9796.
- Zhang, D.; Zhang, H.; Tang, J.; Hua, X.-S.; and Sun, Q. 2020. Causal intervention for weakly-supervised semantic segmentation. *NeurIPS*, 33: 655–666.
- Zhou, L.; Xu, C.; and Corso, J. 2018. Towards automatic learning of procedures from web instructional videos. In *AAAI*, volume 32.