

# Query-Dependent Video Representation for Moment Retrieval and Highlight Detection

WonJun Moon<sup>1,\*</sup>, Sangeek Hyun<sup>1,\*</sup>, SangUk Park<sup>2</sup>, Dongchan Park<sup>2</sup>, Jae-Pil Heo<sup>1,\*</sup>

<sup>1</sup>Sungkyunkwan University, <sup>2</sup>Pyler

{wjun0830, hsi1032, jaepilheo}@g.skku.edu, {psycoder, cto}@pyler.tech

## Abstract

Recently, video moment retrieval and highlight detection (MR/HD) are being spotlighted as the demand for video understanding is drastically increased. The key objective of MR/HD is to localize the moment and estimate clip-wise accordance level, i.e., saliency score, to the given text query. Although the recent transformer-based models brought some advances, we found that these methods do not fully exploit the information of a given query. For example, the relevance between text query and video contents is sometimes neglected when predicting the moment and its saliency. To tackle this issue, we introduce Query-Dependent DETR (QD-DETR), a detection transformer tailored for MR/HD. As we observe the insignificant role of a given query in transformer architectures, our encoding module starts with cross-attention layers to explicitly inject the context of text query into video representation. Then, to enhance the model's capability of exploiting the query information, we manipulate the video-query pairs to produce irrelevant pairs. Such negative (irrelevant) video-query pairs are trained to yield low saliency scores, which in turn, encourages the model to estimate precise accordance between query-video pairs. Lastly, we present an input-adaptive saliency predictor which adaptively defines the criterion of saliency scores for the given video-query pairs. Our extensive studies verify the importance of building the query-dependent representation for MR/HD. Specifically, QD-DETR outperforms state-of-the-art methods on QVHighlights, TVSum, and Charades-STA datasets. Codes are available at [github.com/wjun0830/QD-DETR](https://github.com/wjun0830/QD-DETR).

## 1. Introduction

Along with the advance of digital devices and platforms, video is now one of the most desired data types for consumers [2, 48]. Although the large information capacity of videos might be beneficial in many aspects, e.g., informative

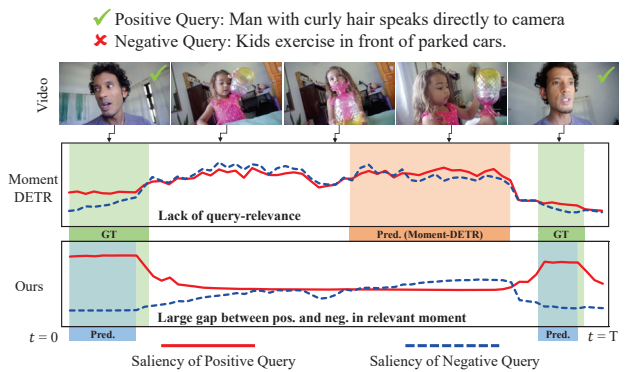


Figure 1. Comparison of highlight-ness (saliency score) when relevant and non-relevant queries are given. We found that the existing work only uses queries to play an insignificant role, thereby may not be capable of detecting negative queries and video-query relevance; saliency scores for clips in ground-truth (GT) moments are low and equivalent for positive and negative queries. On the other hand, query-dependent representations of QD-DETR result in corresponding saliency scores to the video-query relevance and precisely localized moments.

and entertaining, inspecting the videos is time-consuming, so that it is hard to capture the desired moments [1, 2].

Indeed, the need to retrieve user-requested or highlight moments within videos is greatly raised. Numerous research efforts were put into the search for the requested moments in the video [1, 12, 13, 30] and summarizing the video highlights [4, 33, 47, 57]. Recently, Moment-DETR [23] further spotlighted the topic by proposing a QVHighlights dataset that enables the model to perform both tasks, retrieving the moments with their highlight-ness, simultaneously.

When describing the moment, one of the most favored types of query is the natural language sentence (text) [1]. While early methods utilized convolution networks [14, 44, 59], recent approaches have shown that deploying the attention mechanism of transformer architecture is more effective to fuse the text query into the video representation. For example, Moment-DETR [23] introduced the transformer

\* Equal contribution

\* Corresponding author

architecture which processes both text and video tokens as input by modifying the detection transformer (DETR), and UMT [31] proposed transformer architectures to take multi-modal sources, e.g., video and audio. Also, they utilized the text queries in the transformer decoder. Although they brought breakthroughs in the field of MR/HD with seminal architectures, they overlooked the role of the text query. To validate our claim, we investigate the Moment-DETR [23] in terms of the impact of text query in MR/HD (Fig. 1). Given the video clips with a relevant positive query and an irrelevant negative query, we observe that the baseline often neglects the given text query when estimating the query-relevance scores, i.e., saliency scores, for each video clip.

To this end, we propose Query-Dependent DETR (QD-DETR) that produces query-dependent video representation. Our key focus is to ensure that the model's prediction for each clip is highly dependent on the query. First, to fully utilize the contextual information in the query, we revise the transformer encoder to be equipped with cross-attention layers at the very first layers. By inserting a video as the query and a text as the key and value of the cross-attention layers, our encoder enforces the engagement of the text query in extracting video representation. Then, in order to not only inject a lot of textual information into the video feature but also make it fully exploited, we leverage the negative video-query pairs generated by mixing the original pairs. Specifically, the model is learned to suppress the saliency scores of such negative (irrelevant) pairs. Our expectation is the increased contribution of the text query in prediction since the videos will be sometimes required to yield high saliency scores and sometimes low ones depending on whether the text query is relevant or not. Lastly, to apply the dynamic criterion to mark highlights for each instance, we deploy a saliency token to represent the entire video and utilize it as an input-adaptive saliency criterion. With all components combined, our QD-DETR produces query-dependent video representation by integrating source and query modalities. This further allows the use of positional queries [29] in the transformer decoder. Overall, our superior performances over the existing approaches validate the significance of the role of text query for MR/HD.

## 2. Related Work

### 2.1. Moment Retrieval and Highlight Detection

MR is the task of localizing the moment relevant to the given text description. Popular approaches are modeling the cross-modal interaction between text query-video pair [32, 54, 56] or understanding the context of the temporal relation among video clips [1, 59]. On the other hand, TVT [24] exploited the additional data, i.e., subtitle, to capture the moment, and FVMR [14] enhanced the model in terms of inference speed for efficient MR.

Different from the MR, HD aims to measure the clip-wise importance level of the given video [41, 52]. Due to its popularity and applicability, HD can be divided into several branches. From the perspective of annotation, we can categorize HD into supervised, weakly supervised, and unsupervised HD. Supervised HD [15, 41, 50] utilizes fine-grained highlight scores, which are very expensive to collect and annotate [49]. On the other hand, weakly supervised HD [5, 36, 49] learns to detect segments as highlights with video event labels, and finally, unsupervised HD [4, 22, 33, 38] does not require any annotations. Also, while the task is often implemented only with the video, there are works to employ the extra data modalities. Generally, multi-modality was taken into account by using the natural language query to find the desired thumbnail [30] and using additional sources, i.e., audio, to predict highlights [3, 53].

Although the MR/HD share a common objective to localize or discover the desired part of the given video, they have been studied separately. To handle these tasks at once, Moment-DETR [23] proposed the QVHighlights dataset, which contains a human-written text query and its corresponding moment with clip-level saliency labels. They also introduced the modified version of detection transformer (DETR [6]) to localize the query-relevant moments and their saliency scores. Following them, UMT [31] focused on processing multi-modal data by utilizing both video and audio features. Different from recent works deploying transformer architectures, here we concentrate on producing a query-dependent representation with the transformer.

### 2.2. Detection Transformers

DETR [6], an end-to-end object detector based on vision transformers, is one of the very recent works that utilize the transformer architectures for computer vision [11, 42]. Although DETR suffered from slow convergence, it simplifies the prediction process by eliminating the need for anchor generation and non-maximum suppression. Since then, along with the advance in DETR [10, 25, 60], DETR-like architectures have been popular in downstream tasks in both the image [8, 9, 20, 58] and video domains [23, 51]. Some of these works focused on analyzing the role of the decoder query and discovered that using the positional information speeds up the training and also enhances the detection performance [29, 34]. On the other hand, there are trials to extend the application of DETR on multi-modal data [21, 23], especially dealing with the query from different modalities, i.e., text, for detection (or retrieval). They generally handle the multi-modal data by simply forwarding them together to the transformer. In this paper, we also focus on handling the multi-modal data based on DETR-like architecture. However, different from the aforementioned techniques, we concentrate on the query-dependency of the prediction results.

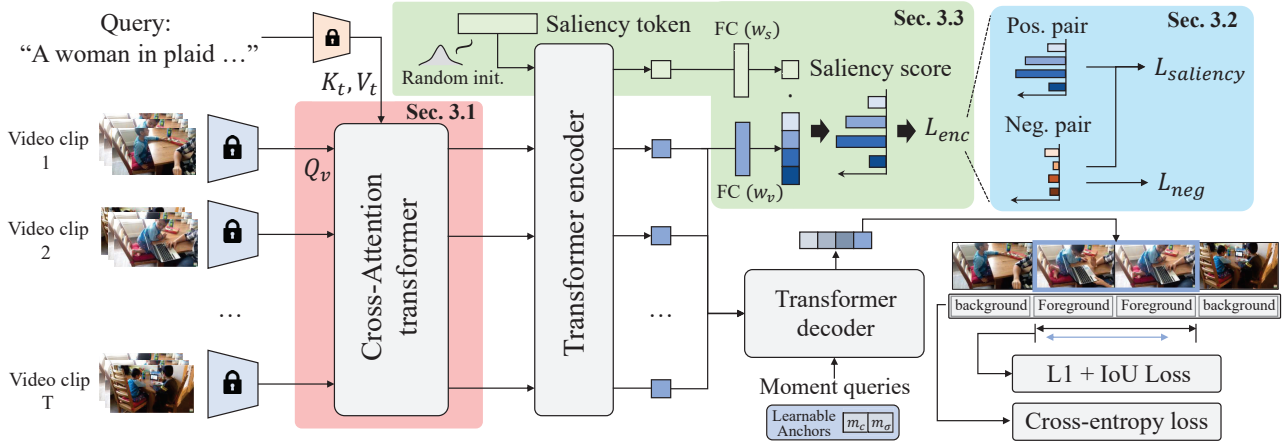


Figure 2. Overview of the proposed QD-DETR architecture. Given a video and text query, we first extract video and text features from the frozen backbones. These video and text features are forwarded into the cross-attention transformer (Sec. 3.1). This process ensures the consistent contribution of text queries to the video tokens and together with negative pair learning (Sec. 3.2), builds query-dependent video representation. Then, accompanied by the saliency token (Sec. 3.3), video tokens are given to the transformer encoder. In this procedure, the saliency token is transformed into the adaptive saliency prediction criteria. The outputs of the encoder are then, processed to compute losses for both HD and MR. Specifically, the encoder’s output tokens are directly projected to saliency scores and optimized for HD, and also provided to the transformer decoder with the learnable moment queries to estimate the query-described moments. Finally, losses for MR are computed by the discrepancy between predicted and their corresponding GT moments.

### 3. Query-Dependent DETR

Moment retrieval and highlight detection have the common objective to find preferred moments with the text query. Given a video of  $L$  clips and a text query with  $N$  words, we denote their representations as  $\{v_1, v_2, \dots, v_L\}$  and  $\{t_1, t_2, \dots, t_N\}$  extracted by frozen video and text encoders, respectively. With these representations, the main objective is to localize the center coordinate  $m_c$  and width  $m_\sigma$  within the video and rank the highlight score (saliency score)  $\{s_1, s_2, \dots, s_L\}$  for each clip. A straightforward approach to utilize transformer [11, 43] for the MR is to make a moment-wise prediction as a set of clips [23], or generating the moment according to the clip-wise predictions [31]. To exploit the multi-modal information, e.g., video and text query, they either simply concatenated the features across the modalities or inserted the texts to form the moment query to the transformer decoder. However, we claim that the relationship between the video and text query should be carefully considered rather than a simple concatenation since MR/HD requires every video clip to be conditionally assessed with the text queries.

Our overall architecture is described in Fig. 2, following the design of concrete baseline, Moment-DETR [23]. Given a video and query representation extracted from fixed backbones, QD-DETR first transforms the video representation to be query-dependent using cross-attention layers. To further enhance the query-awareness of video representations, we incorporate irrelevant video-query pairs with a low saliency

for the learning objective. Then, along with the transformer encoder-decoder architectures, the saliency token is defined that turns into an adaptive saliency predictor when attended by the specific video instance.

#### 3.1. Cross-Attentive Transformer Encoder

In this subsection, we use italic letters to represent *query*, *key*, and *value* of the cross-attention layers. The key objective of the encoder for MR/HD is to produce clip-wise representations equipped with information regarding the degree of query-relevance since these features are directly used for retrieving the query-matched moments and predicting clip-wise saliency scores. However, the encoding process of existing works may not ensure the query conditioning on every clip. For example, Moment-DETR [23] naively concatenated the video with the query for input to the self-attention layers, which may result in an insignificant role of the query if the high similarities among the video clips overwhelm the contribution of the text query. On the other hand, UMT [31] utilizes the text query only for the synthesis of a moment query in the transformer decoder so thus resulting video representations are not associated with the text query.

To take the textual contexts into every video clip representation, we deploy cross-attention layers between the source and the query modalities at the very first layers of the encoder. This ensures the consistent contribution of the query, thereby extracting query-dependent video representation. In detail, whereas the *query* for cross-attention layers is prepared by projecting the video clips as  $Q_v = [p_q(v_1), \dots, p_q(v_L)]$ , the

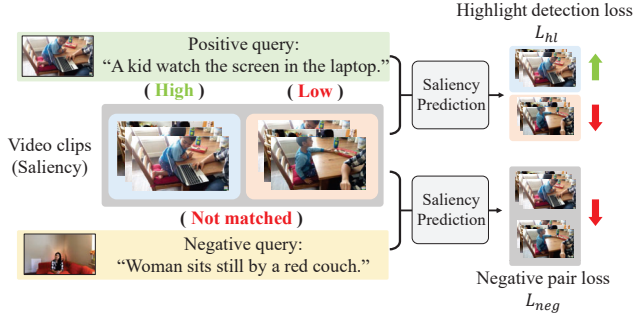


Figure 3. Illustration of negative pair learning. Typical HL loss is defined only with a positive video-query pair, which is insufficient to learn various degrees of query-relevance. On the other hand, our negative pair learning enforces the model to yield different scores for a video depending on the query and to be learned to suppress saliency scores for the negative query.

key and value are computed with the query text features as  $K_t = [p_k(t_1), \dots, p_k(t_N)]$  and  $V_t = [p_v(t_1), \dots, p_v(t_N)]$ .  $p_q(\cdot)$ ,  $p_k(\cdot)$ , and  $p_v(\cdot)$  are projection layers for query, key, and value. Then, the cross-attention layer operates as follows:

$$\text{Attention}(Q_v, K_t, V_t) = \text{softmax}\left(\frac{Q_v K_t^T}{\sqrt{d}}\right) V_t, \quad (1)$$

where  $d$  is the dimension of the projected key, value, and query. Since the softmax scores are distributed only over the query elements, video clips are expressed with the weighted sum of the text queries in proportion to the similarity to texts. Attention scores are then projected through MLP and integrated into the original video representations as the typical transformer layers. For the rest of the paper, we define the query-dependent video tokens, i.e., the output of cross-attention layers, as  $X = \{x_v^1, x_v^2, \dots, x_v^L\}$ .

### 3.2. Learning from Negative Relationship

While the cross-attention layers explicitly fuse the video and query features for intermediate video clip representations to engage the query information in an architectural way, we argue that given video-text pairs lack diversity to learn the general relationship. For instance, many consecutive clips in a single video often share similar appearances, and the similarity to a specific query will not be highly distinguishable, thereby, the text query may not much affect the prediction.

Thus, we consider the relationships between irrelevant pairs of videos inspired by many recognition practices [17, 26, 27, 35] that learn discriminative features across different categories. To implement such relationships, we define given training video-query pairs as positive pairs and mix the video and query from different pairs to construct negative pairs. Fig. 3 illustrates the ways to augment such negative pairs

and utilize them with positive pairs in training. While the video clips in positive pairs are trained to yield segmented saliency scores according to the query-relevance, irrelevant negative video-query pairs are enforced to have the lowest saliency scores. Formally, the loss function for suppressing the saliency of negative pairs  $x_v^{\text{neg}}$  are expressed as follows:

$$L_{\text{neg}} = -\log(1 - S(x_v^{\text{neg}})), \quad (2)$$

where  $S(\cdot)$  is the saliency score predictor. This training scheme can also prevent the model from predicting the moments and highlights solely based on the inter-relationship among video clips without consideration of the query-relevance since the same video instance should be predicted differently depending on whether the positive or negative query is given.

### 3.3. Input-Adaptive Saliency Predictor

Naive implementation for saliency predictor  $S(\cdot)$  would be stacking one or more fully-connected layers. However, such a general head provides identical criteria for the saliency prediction of every video-query pair, neglecting the diverse nature of video and natural language query pairs. This violates our key idea to extract query-dependent video representation.

Thus, we define the saliency token  $x_s$  to be utilized as an input-adaptive saliency predictor. Briefly, the saliency token is a randomly initialized learnable vector that becomes an input-adaptive predictor when added to the sequence of encoded video tokens and projected through the transformer encoder. To illustrate, as shown in Fig. 2, we first concatenate the saliency token with the query-dependent video tokens  $X$ . We process these tokens to the transformer encoder which makes the saliency token to be re-organized with the input-dependent contexts. Consequently, saliency and video tokens are projected by a corresponding single fully-connected layer with weights,  $w_s$  and  $w_v$ , respectively, where their scaled-dot product becomes the saliency scores. Formally, saliency score  $S(x_v^i)$  is computed as follows:

$$S(x_v^i) = \frac{w_s^T x_s \cdot w_v^T x_v^i}{\sqrt{d}}, \quad (3)$$

where  $d$  is the channel dimension of projected tokens.

### 3.4. Decoder and Objectives

**Transformer Decoder.** Recently, understanding the role of the query in the detection transformer is being spotlighted [29, 34]. It is verified that designing the query with the positional information helps not only for acceleration of training but also for enhancing accuracy. Yet, it is hard to directly employ these studies in tasks handling multi-modal data, e.g., MR/HD, since multi-modal data often have different definitions of position; the position can be understood as time in the video and word order in the text.



On the contrary, our architectural design eliminates the need to feed the text query to the decoder since the query information is already taken into the video representations. To this end, we modify the 2D dynamic anchor boxes [29] to represent 1D moments in the video. Specifically, we utilize the center coordinate  $m_c$  and the duration  $m_\sigma$  of the moments to design the queries. Similarly to the previous way in the image domain, we pool the features around the center coordinate and modulate the cross-attention map with the moment duration. Then, the coordinates and durations are layer-wisely revised.

**Loss Functions.** Training objectives for QD-DETR include loss functions for MR/HD, respectively. First, objective functions for MR, in which the key focus is to locate the desired moments, are adopted from the baseline [23]. Moment retrieval loss  $L_{mr}$  measures the discrepancy between the GT moment and the predicted counterpart. It consists of a  $L1$  loss and a generalized IoU loss  $L_{gIoU}(\cdot)$  from previous work [37] with minor modification to localize temporal moments. Additionally, the cross-entropy loss is used to classify the predicted moments as  $\hat{y}$  either to foreground and background by  $L_{CE} = -\sum_{y \in Y} y \log(\hat{y})$  where  $\{fg, bg\} \subset Y$ . Thus,  $L_{mr}$  is defined as follows:

$$L_{mr} = \lambda_{L1} \|m - \hat{m}\| + \lambda_{gIoU} L_{gIoU}(m, \hat{m}) + \lambda_{CE} L_{CE}, \quad (4)$$

where  $m$  and  $\hat{m}$  are ground-truth moment and its correspond prediction containing center coordinate  $m_c$  and duration  $m_\sigma$ . Also,  $\lambda_*$  are hyperparameters for balancing the losses.

Loss functions for HD are to estimate the saliency score. It comprises two components; margin ranking loss  $L_{margin}$  and rank-aware contrastive loss  $L_{cont}$ . Following [23], the margin rank loss operates with two pairs of high-rank and low-rank clips. To be specific, the high-rank clips are ensured to retain higher saliency scores than both the low-rank clips within the GT moment and the negative clips outside the GT moment. In short,  $L_{margin}$  is defined as:

$$L_{margin} = \max(0, \Delta + S(x^{\text{low}}) - S(x^{\text{high}})) \quad (5)$$

where  $\Delta$  is the margin,  $S(\cdot)$  is the saliency score estimator, and  $x^{\text{high}}$  and  $x^{\text{low}}$  are video tokens from two pairs of high and low-rank clips, respectively. In addition to margin loss which only indirectly guides the saliency predictor, we employ rank-aware contrastive loss [18] to learn the precisely segmented saliency levels with the contrastive loss. Given the maximum rank value  $R$ , each clip in the mini-batch has a saliency score lower than  $R$ . Then, we iterate the batch for  $R$  times, each time utilizing the samples with higher saliency scores than the iteration index ( $r \in \{0, 1, \dots, R-1\}$ ) to build the positive set  $X_r^{\text{pos}}$ . Samples with a lower rank than the iteration index are included in the negative set  $X_r^{\text{neg}}$ . Then,

the rank-aware contrastive loss  $L_{cont}$  is defined as:

$$L_{cont} = -\sum_{r=1}^R \log \frac{\sum_{x \in X_r^{\text{pos}}} \exp(S(x)/\tau)}{\sum_{x \in (X_r^{\text{pos}} \cup X_r^{\text{neg}})} \exp(S(x)/\tau)} \quad (6)$$

where  $\tau$  is a temperature scaling parameter. Note that,  $X_r^{\text{neg}}$  also include all clips in negative pairs  $x_v^{\text{neg}}$  defined in Sec. 3.2. Finally with margin loss and rank-aware contrastive loss,  $L_{hl}$  and total loss function  $L_{total}$  are defined as follows:

$$L_{hl} = \lambda_{margin} L_{margin} + \lambda_{cont} L_{cont}, \quad (7)$$

$$L_{total} = L_{hl} + L_{mr} + \lambda_{neg} L_{neg}. \quad (8)$$

## 4. Evaluation

### 4.1. Experimental Settings

**Dataset and Evaluation Metrics.** For the evaluation, we validate the effectiveness of query-dependent source representation on QVHighlights [23], TVSum [40], CharadesSTA [13]. **QVHighlights** is the most recently publicized dataset for both moment retrieval and highlight detection. It is also the only dataset that has annotations for both tasks. In detail, QVHighlights consists of over 10,000 videos annotated with human-written text queries. It provides a fair benchmark as the evaluation for the test split can only be measured through submitting the prediction to the QVHighlights server<sup>1</sup>. **Charades-STA** and **TVSum** are the dataset for moment retrieval and video summarization, respectively. Each of them contains 9,848 videos regarding indoor activities and 50 videos of various genres, e.g., news, documentary, and vlog. For all datasets, we follow the data splits from the existing works [23, 31].

To measure the performances, we use the same evaluation metrics used in the baselines. Specifically, recall@1 with IoU thresholds 0.5 and 0.7, and mean average precision (mAP) at different thresholds. Similarly, we use mAP and HIT@1 for evaluating the highlight detection. HIT@1 is computed through the hit ratio of the highest-scored clip.

### 4.2. Experimental Results

We compare QD-DETR against baselines in MR and HD throughout Tab. 1, Tab. 3, and Tab. 2. Our experiments with multi-modal sources, i.e., video with audio, are implemented by simply concatenating the video and audio along the channel axis. Throughout the tables, we use bolds to denote the best scores.

In Tab. 1, the task is to jointly learn and predict MR/HD. As observed, our QD-DETR outperforms state-of-the-art (SOTA) approaches with all evaluation metrics. Among methods utilizing the video source, QD-DETR shows a dramatic increase with stricter metrics with high IOU; it outperforms previous SOTA by large margins up

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

Table 1. Performance comparison on QVHighlights *test* split. V and A in the Src column denote video and audio, respectively, representing the modalities of the source data. Our experiments are averaged over five runs and ‘±’ denotes the standard deviation.

Method	Src	MR					HD	
		R1		mAP			>= Very Good	
		@0.5	@0.7	@0.5	@0.75	Avg.	mAP	HIT@1
BeautyThumb [39]	V	-	-	-	-	-	14.36	20.88
DVSE [30]	V	-	-	-	-	-	18.75	21.79
MCN [1]	V	11.41	2.72	24.94	8.22	10.67	-	-
CAL [12]	V	25.49	11.54	23.40	7.65	9.89	-	-
XML [24]	V	41.83	30.35	44.63	31.73	32.14	34.49	55.25
XML+ [24]	V	46.69	33.46	47.89	34.67	34.90	35.38	55.06
Moment-DETR [23]	V	52.89 $\pm$ 2.3	33.02 $\pm$ 1.7	54.82 $\pm$ 1.7	29.40 $\pm$ 1.7	30.73 $\pm$ 1.4	35.69 $\pm$ 0.5	55.60 $\pm$ 1.6
QD-DETR (Ours)	V	<b>62.40</b> $\pm$ 1.1	<b>44.98</b> $\pm$ 0.8	<b>62.52</b> $\pm$ 0.6	<b>39.88</b> $\pm$ 0.7	<b>39.86</b> $\pm$ 0.6	<b>38.94</b> $\pm$ 0.4	<b>62.40</b> $\pm$ 1.4
UMT [31]	V+A	56.23	41.18	53.38	37.01	36.12	38.18	59.99
QD-DETR (Ours)	V+A	<b>63.06</b> $\pm$ 1.0	<b>45.10</b> $\pm$ 0.7	<b>63.04</b> $\pm$ 0.9	<b>40.10</b> $\pm$ 1.0	<b>40.19</b> $\pm$ 0.6	<b>39.04</b> $\pm$ 0.3	<b>62.87</b> $\pm$ 0.6

Table 2. Highlight detection performance comparison on TVsum dataset.

Method	Src	VT	VU	GA	MS	PK	PR	FM	BK	BT	DS	Avg.
sLSTM [57]	V	41.1	46.2	46.3	47.7	44.8	46.1	45.2	40.6	47.1	45.5	45.1
SG [33]	V	42.3	47.2	47.5	48.9	45.6	47.3	46.4	41.7	48.3	46.6	46.2
LIM-S [49]	V	55.9	42.9	61.2	54.0	60.3	47.5	43.2	66.3	69.1	62.6	56.3
Trailer [45]	V	61.3	54.6	65.7	60.8	59.1	70.1	58.2	64.7	65.6	68.1	62.8
SL-Module [50]	V	86.5	68.7	74.9	<b>86.2</b>	79.0	63.2	58.9	72.6	78.9	64.0	73.3
QD-DETR (Ours)	V	<b>88.2</b>	<b>87.4</b>	<b>85.6</b>	85.0	<b>85.8</b>	<b>86.9</b>	<b>76.4</b>	<b>91.3</b>	<b>89.2</b>	<b>73.7</b>	<b>85.0</b>
MINI-Net [19]	V+A	80.6	68.3	78.2	81.8	78.1	65.8	57.8	75.0	80.2	65.5	73.2
TCG [53]	V+A	85.0	71.4	81.9	78.6	80.2	75.5	71.6	77.3	78.6	68.1	76.8
Joint-VA [3]	V+A	83.7	57.3	78.5	86.1	80.1	69.2	70.0	73.0	<b>97.4</b>	67.5	76.3
UMT [31]	V+A	87.5	81.5	88.2	78.8	81.4	87.0	76.0	86.9	84.4	<b>79.6</b>	83.1
QD-DETR (Ours)	V+A	<b>87.6</b>	<b>91.7</b>	<b>90.2</b>	<b>88.3</b>	<b>84.1</b>	<b>88.3</b>	<b>78.7</b>	<b>91.2</b>	87.8	77.7	<b>86.6</b>

Table 3. Charades dataset. † denotes the method using the video and audio as the source. SF+C stands for Slowfast and CLIP features.

Method	feat	R1@0.5	R1@0.7	Method	feat	R1@0.5	R1@0.7
SAP	VGG	27.42	13.36	CTRL	C3D	23.63	8.89
TripNet	VGG	36.61	14.50	ACL	C3D	30.48	12.20
SM-RL	VGG	24.36	11.17	RWM-RL	C3D	36.70	-
MAN	VGG	41.24	20.54	MAN	C3D	46.53	22.72
2D-TAN	VGG	40.94	22.85	DEBUG	C3D	37.39	17.69
FVMR	VGG	42.36	24.14	VSLNet	C3D	47.31	30.19
UMT†	VGG	48.31	29.25	<b>Ours</b>	C3D	<b>50.67</b>	<b>31.02</b>
<b>Ours</b>	VGG	52.77	31.13	M-DETR	SF+C	53.63	31.37
<b>Ours†</b>	VGG	<b>55.51</b>	<b>34.17</b>	<b>Ours</b>	SF+C	<b>57.31</b>	<b>32.55</b>

to 36% in R1@0.7 and mAP@0.75. On the other hand, QD-DETR with video and audio sources boosts 11.84% on average of the metrics for MR compared to the SOTA method employing the multi-modal source data. These results verify the importance of emphasizing the source (video-only or video+audio) descriptive contexts in the text queries.

Results in Tab. 3 also compare MR performances against the models using VGG [7, 14, 31, 46, 55, 59], C3D [13, 16, 28, 32, 55, 56], and Slowfast (SF) and CLIP features [23] on Charades dataset. For a fair comparison, we enumerate

each method with its backbone and compare within it. For each feature from VGG, C3D, and SF+C, we follow the data preparation settings from UMT [31], VSLNet [56], and Moment-DETR [23]. As reported, we validate that our model surpasses the existing SOTA methods in every type of feature.

For video highlight detection in Tab. 2, we follow the protocols from the previous work [31]. Specifically, we train the model for each category and average the mAP scores. Out of 10 categories, QD-DETR outperforms baselines on 9 categories when the only video source is available, and 8 categories when both video and audio are available. Overall, compared to methods with video-only and multi-modal sources, QD-DETR establishes new SOTA performances by improving by 4.2% in average compared to the previous SOTA model.

### 4.3. Ablation study

To investigate the effectiveness of each component in our work, we conduct an extensive ablation study in Tab. 4. Note that, CATE and DAM denote cross-attentive transformer encoder and dynamic anchor moments, respectively. Rows (b) to (e) show the effectiveness of each component compared to

Table 4. Ablation study on QVHighlights *val* split. CATE and DAM stands for cross-attentive transformer encoder and using dynamic anchor moments as the decoder query, respectively. All the quantities are averaged over 5 runs.

	CATE	Neg. Pair	Saliency Token	DAM	MR					HD	
					R1		mAP		Avg.	>= Very Good	
					@0.5	@0.7	@0.5	@0.75		mAP	HIT@1
(a)					52.89	33.02	54.82	29.40	30.73	35.69	55.60
(b)	✓				56.16	38.71	56.48	33.42	34.07	37.14	58.34
(c)		✓			58.69	39.83	58.39	34.84	35.40	39.02	62.81
(d)			✓		55.48	37.00	55.81	26.75	32.84	37.48	58.59
(e)				✓	53.19	35.91	55.58	32.55	33.33	35.68	55.56
(f)	✓			✓	57.72	42.35	59.10	38.16	38.03	36.56	57.44
(g)	✓	✓			59.57	42.12	59.19	36.63	36.76	38.64	61.62
(h)		✓	✓		60.00	40.97	59.21	35.41	35.89	39.06	62.88
(i)	✓	✓	✓		60.32	42.39	59.47	36.79	36.93	<b>39.21</b>	62.76
(j)	✓	✓	✓	✓	<b>62.68</b>	<b>46.66</b>	<b>62.23</b>	<b>41.82</b>	<b>41.22</b>	39.13	<b>63.03</b>

Table 5. Ablation study on cross-attention transformer encoder. We compare ours against the deepened transformer encoder with self-attention layers to validate that the performance gain does not come from additional parameters. SATE and CATE each indicate the transformer encoder only with self-attention layers and our transformer encoder. The numbers in the parenthesis denote the number of layers. For the experiment with †, we only use the query features as the condition in the encoder and only the video representations are processed by the decoder.

T2V	MR					HD	
	R1		mAP		Avg.	>= Very Good	
	@0.5	@0.7	@0.5	@0.75		mAP	HIT@1
Moment-DETR (SATE 2)	52.89 $\pm$ 2.3	33.02 $\pm$ 1.7	54.82 $\pm$ 1.7	29.40 $\pm$ 1.7	30.73 $\pm$ 1.4	35.69 $\pm$ 0.5	55.60 $\pm$ 1.6
Moment-DETR (SATE 4)	53.60 $\pm$ 1.2	35.81 $\pm$ 0.9	54.55 $\pm$ 0.8	30.64 $\pm$ 0.7	31.74 $\pm$ 0.4	35.96 $\pm$ 0.2	56.56 $\pm$ 0.9
Moment-DETR (CATE 4)	55.10 $\pm$ 0.7	37.02 $\pm$ 0.9	56.21 $\pm$ 0.3	32.00 $\pm$ 0.9	33.19 $\pm$ 0.6	36.43 $\pm$ 0.3	56.98 $\pm$ 0.6
Moment-DETR (CATE 4)†	56.16 $\pm$ 1.2	38.71 $\pm$ 1.1	56.48 $\pm$ 0.8	33.42 $\pm$ 0.7	34.07 $\pm$ 0.6	37.14 $\pm$ 0.4	58.34 $\pm$ 0.4
QD-DETR (SATE 4)†	60.48 $\pm$ 0.7	45.21 $\pm$ 1.0	60.84 $\pm$ 0.5	40.45 $\pm$ 0.7	40.12 $\pm$ 0.6	38.66 $\pm$ 0.2	61.29 $\pm$ 1.0
QD-DETR (CATE 4)†	<b>62.68<math>\pm</math>1.1</b>	<b>46.66<math>\pm</math>0.6</b>	<b>62.23<math>\pm</math>1.0</b>	<b>41.82<math>\pm</math>0.9</b>	<b>41.22<math>\pm</math>0.4</b>	<b>39.13<math>\pm</math>0.3</b>	<b>63.03<math>\pm</math>0.5</b>

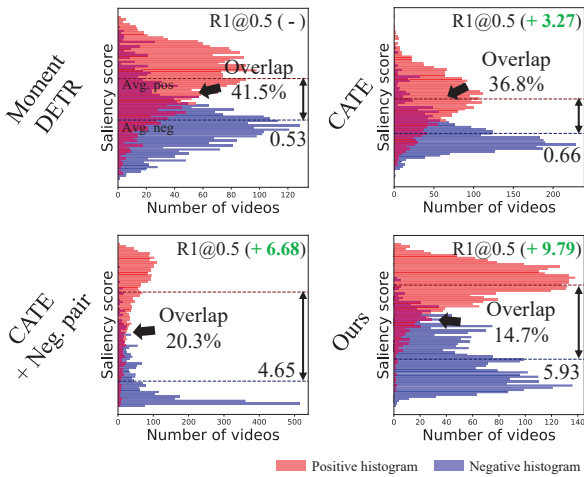


Figure 4. Ablation study in terms of saliency scores. We plot the histograms of the average value of saliency scores in each video when the positive and negative text queries are given. For positive scores, we only account the scores within GT moments. The average value of each histogram are visualized by the dotted line. The decrease in the overlap between histograms and the increase in the gap between average values confirms the gradual improvements of significance of the query in extracting video representation.

the baseline (a). To explain, whereas (e) only boosts the MR performances since it only affects the transformer decoder, (b), (c), and (d) are especially beneficial for both MR/HD tasks since they are focused on query-dependent video representations ; (b) ensures the contributions of text query in the video representation, (c) fully exploits the contexts of the text query, and (d) provides input-adaptive saliency predictor instead of MLP. Moreover, while our components are verified that they are all complementary to others, DAM’s effectiveness is especially dependent on the usage of CATE (compare between {(a, e)} and {(b, f), (i, j)}). We claim that this is because DAM exploits the position information of the input tokens to capture the corresponding moments. However, without CATE, input tokens are a mixture of multi-modal tokens, thereby providing confusing position information.

To provide in-depth examinations of each component, we inspect the difference between the positive and the negative saliency scores in Fig. 4. Since the role of text query is trivial in our baseline, each distribution significantly overlies on top of the other. Then, as we add CATE and negative pair learning, we observe a consistent decrease in overlapped

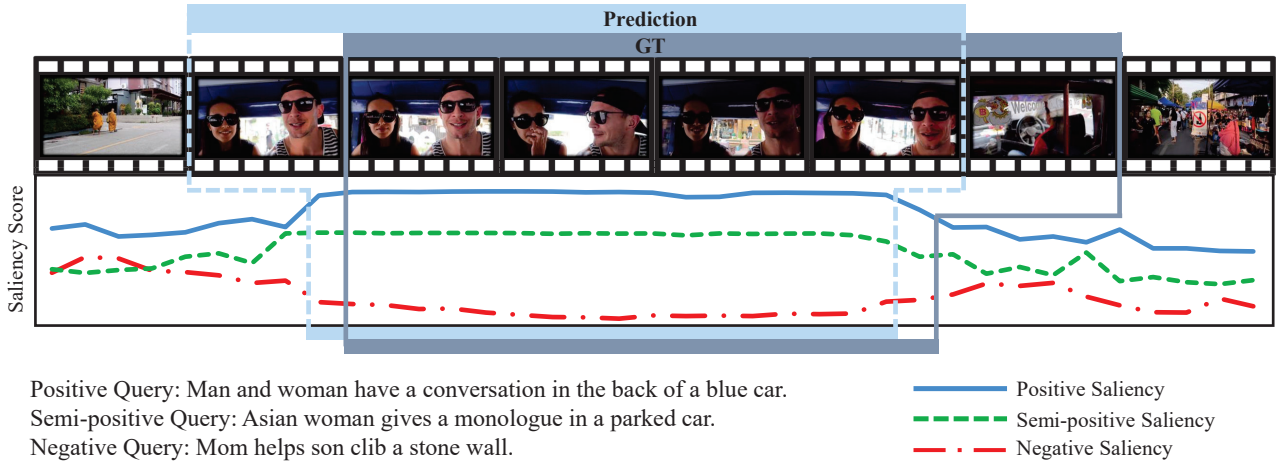


Figure 5. Visualization of results predicted by QD-DETR. Predicted and ground-truth moments are bounded by the lines. Blue, green, and red lines indicate the saliency scores for positive, semi-positive, and negative queries. The positive saliency scores are consistently higher than the others, while the scores for semi-positive are higher than the ones for the negative.

areas and a larger gap between the average saliency scores of positive and negative histograms. Also, we believe that the widely-distributed histogram of saliency scores for the 'CATE+Neg.pair' is due to using an identical criterion for saliency prediction for diverse video-query representations. By employing an input-adaptive saliency predictor, we notice that scores for positive queries are in almost optimal shape.

In addition, some might ask whether CATE benefits the training because of additional encoder layers. To answer this, we conduct another ablation study in Tab. 5. Briefly, since our transformer encoder utilizes 2 cross-attention layers and 2 self-attention layers, we conduct comparisons against the transformer encoder composed of 4 self-attention layers (SATE). First, we compare CATE and SATE on Moment-DETR; by comparing the results in 2<sup>nd</sup> and 3<sup>rd</sup> rows, we find that CATE is much more beneficial than SATE even with the same number of layers. Furthermore, the last two rows show comparisons within QD-DETR architecture that has the same tendency. These results clearly demonstrate that the improvements from CATE are mainly from emphasizing the role of the text query rather than additional layers.

#### 4.4. Qualitative Results

In this subsection, we study how the query-dependent video representation sensitively reacts to the change in the contexts of the text query. In Fig. 5, the measured saliency scores according to the video-query relevance are visualized. We found that the more the query is relevant to the video clips, the higher the saliency scores retained for the query. For instance, whereas the negative query that is totally irrelevant to the video instance has the lowest scores, the scores for semi-positive reside between the positive and the negative ones. Also, we find that QD-DETR sometimes provides a more precise moment prediction than a given ground-truth

moment, as can be seen with the temporal box bounded by the dotted lines. We believe that the tendency of a bit higher saliency scores at non-relevant clips for a positive query is due to the information mixing in the self-attention layers.

## 5. Limitation and Conclusion

**Limitation** As elaborated in the paper, we aim to highlight the role of the text query in retrieving the relevant moments and estimating their accordance level with the given text query. Likewise, the proposed components expect a given query to maintain a meaningful context. If not, and noisy text queries are provided, i.e., mismatched or irrelevant ground truth texts, the training may not be effective as reported.

**Conclusion** Although the advent of transformer architecture has been powerful for MR/HD, investigation of the role of text query has been lacking in such architectures. Therefore, we focused on studying the role of the text query. As we found that the textual information is not fully exploited in expressing the video representations, we designed the cross-attentive transformer encoder and proposed a negative-pair training scheme. Cross-attentive encoder assures the query's contributions while extracting video representation, and negative-pair training enforces the model to learn the relationship between query and video by preventing solving the problems without consideration of the query. Finally, to preserve the diversity of query-dependent video representation, we defined the saliency token to be an input-adaptive saliency predictor. Extensive experiments validated the strength of QD-DETR with superior performances.

**Acknowledgements.** This work was supported in part by MSIT/IITP (No. 2022-0-00680, 2019-0-00421, 2020-0-01821, 2021-0-02068), and MSIT&KNPA/KIPoT (Police Lab 2.0, No. 210121M06).



## References

- [1] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, pages 5803–5812, 2017. 1, 2, 6
- [2] Evlampios Apostolidis, Eleni Adamantidou, Alexandros I Metsai, Vasileios Mezaris, and Ioannis Patras. Video summarization using deep neural networks: A survey. *Proceedings of the IEEE*, 109(11):1838–1863, 2021. 1
- [3] Taivanbat Badamdorj, Mrigank Rochan, Yang Wang, and Li Cheng. Joint visual and audio learning for video highlight detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8127–8137, 2021. 2, 6
- [4] Taivanbat Badamdorj, Mrigank Rochan, Yang Wang, and Li Cheng. Contrastive learning for unsupervised video highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14042–14052, 2022. 1, 2
- [5] Sijia Cai, Wangmeng Zuo, Larry S Davis, and Lei Zhang. Weakly-supervised video summarization using variational encoder-decoder and web prior. In *Proceedings of the European conference on computer vision (ECCV)*, pages 184–200, 2018. 2
- [6] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 2
- [7] Shaoxiang Chen and Yu-Gang Jiang. Semantic proposal for activity localization in videos via sentence query. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019. 6
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, 2022. 2
- [9] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. 2
- [10] Xiyang Dai, Yinpeng Chen, Jianwei Yang, Pengchuan Zhang, Lu Yuan, and Lei Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2988–2997, 2021. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 3
- [12] Victor Escorcia, Mattia Soldan, Josef Sivic, Bernard Ghanem, and Bryan Russell. Temporal localization of moments in video collections with natural language. *arXiv preprint arXiv:1907.12763*, 2019. 1, 6
- [13] Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE international conference on computer vision*, pages 5267–5275, 2017. 1, 5, 6
- [14] Junyu Gao and Changsheng Xu. Fast video moment retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1523–1532, 2021. 1, 2, 6
- [15] Michael Gygli, Yale Song, and Liangliang Cao. Video2gif: Automatic generation of animated gifs from video. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1001–1009, 2016. 2
- [16] Dongliang He, Xiang Zhao, Jizhou Huang, Fu Li, Xiao Liu, and Shilei Wen. Read, watch, and move: Reinforcement learning for temporally grounding natural language descriptions in videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8393–8400, 2019. 6
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 4
- [18] David T Hoffmann, Nadine Behrmann, Juergen Gall, Thomas Brox, and Mehdi Noroozi. Ranking info noise contrastive estimation: Boosting contrastive learning via ranked positives. In *AAAI Conference on Artificial Intelligence*, 2022. 5
- [19] Fa-Ting Hong, Xuanteng Huang, Wei-Hong Li, and Wei-Shi Zheng. Mini-net: Multiple instance ranking network for video highlight detection. In *European Conference on Computer Vision*, pages 345–360. Springer, 2020. 6
- [20] Kuan-Chih Huang, Tsung-Han Wu, Hung-Ting Su, and Winston H Hsu. Monodtr: Monocular 3d object detection with depth-aware transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4012–4021, 2022. 2
- [21] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdettr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1780–1790, 2021. 2
- [22] Aditya Khosla, Raffay Hamid, Chih-Jen Lin, and Neel Sundaresan. Large-scale video summarization using web-image priors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2698–2705, 2013. 2
- [23] Jie Lei, Tamara L Berg, and Mohit Bansal. Detecting moments and highlights in videos via natural language queries. *Advances in Neural Information Processing Systems*, 34:11846–11858, 2021. 1, 2, 3, 5, 6
- [24] Jie Lei, Licheng Yu, Tamara L Berg, and Mohit Bansal. Tvr: A large-scale dataset for video-subtitle moment retrieval. In *European Conference on Computer Vision*, pages 447–463. Springer, 2020. 2, 6
- [25] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13619–13627, 2022. 2

- [26] Junnan Li, Richard Socher, and Steven CH Hoi. Dividemix: Learning with noisy labels as semi-supervised learning. *arXiv preprint arXiv:2002.07394*, 2020. 4
- [27] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 4
- [28] Meng Liu, Xiang Wang, Liqiang Nie, Xiangnan He, Baoquan Chen, and Tat-Seng Chua. Attentive moment retrieval in videos. In *The 41st international ACM SIGIR conference on research & development in information retrieval*, pages 15–24, 2018. 6
- [29] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022. 2, 4, 5
- [30] Wu Liu, Tao Mei, Yongdong Zhang, Cherry Che, and Jiebo Luo. Multi-task deep visual-semantic embedding for video thumbnail selection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3707–3715, 2015. 1, 2, 6
- [31] Ye Liu, Siyuan Li, Yang Wu, Chang-Wen Chen, Ying Shan, and Xiaohu Qie. Umt: Unified multi-modal transformers for joint video moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3042–3051, 2022. 2, 3, 5, 6
- [32] Chujie Lu, Long Chen, Chilie Tan, Xiaolin Li, and Jun Xiao. Debug: A dense bottom-up grounding approach for natural language video localization. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5144–5153, 2019. 2, 6
- [33] Behrooz Mahasseni, Michael Lam, and Sinisa Todorovic. Unsupervised video summarization with adversarial lstm networks. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 202–211, 2017. 1, 2, 6
- [34] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Conditional detr for fast training convergence. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3651–3660, 2021. 2, 4
- [35] WonJun Moon, Junho Park, Hyun Seok Seong, Cheol-Ho Cho, and Jae-Pil Heo. Difficulty-aware simulator for open set recognition. In *European Conference on Computer Vision*, pages 365–381. Springer, 2022. 4
- [36] Rameswar Panda, Abir Das, Ziyang Wu, Jan Ernst, and Amit K Roy-Chowdhury. Weakly supervised summarization of web videos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3657–3666, 2017. 2
- [37] Hamid Rezaatofghi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 658–666, 2019. 5
- [38] Mrigank Rochan, Linwei Ye, and Yang Wang. Video summarization using fully convolutional sequence networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 347–363, 2018. 2
- [39] Yale Song, Miriam Redi, Jordi Vallmitjana, and Alejandro Jaimes. To click or not to click: Automatic selection of beautiful thumbnails from videos. In *Proceedings of the 25th ACM international on conference on information and knowledge management*, pages 659–668, 2016. 6
- [40] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. Tvsum: Summarizing web videos using titles. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5179–5187, 2015. 5
- [41] Min Sun, Ali Farhadi, and Steve Seitz. Ranking domain-specific highlights by analyzing edited videos. In *European conference on computer vision*, pages 787–802. Springer, 2014. 2
- [42] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 2
- [43] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [44] Jingwen Wang, Lin Ma, and Wenhao Jiang. Temporally grounding language queries in videos by contextual boundary-aware prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12168–12175, 2020. 1
- [45] Lezi Wang, Dong Liu, Rohit Puri, and Dimitris N Metaxas. Learning trailer moments in full-length movies with co-contrastive attention. In *European Conference on Computer Vision*, pages 300–316. Springer, 2020. 6
- [46] Weining Wang, Yan Huang, and Liang Wang. Language-driven temporal activity localization: A semantic matching reinforcement learning model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 334–343, 2019. 6
- [47] Fanyue Wei, Biao Wang, Tiezheng Ge, Yuning Jiang, Wen Li, and Lixin Duan. Learning pixel-level distinctions for video highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3073–3082, 2022. 1
- [48] Zuxuan Wu, Ting Yao, Yanwei Fu, and Yu-Gang Jiang. Deep learning for video classification and captioning. In *Frontiers of multimedia research*, pages 3–29. 2017. 1
- [49] Bo Xiong, Yannis Kalantidis, Deepti Ghadiyaram, and Kristen Grauman. Less is more: Learning highlight detection from video duration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1258–1267, 2019. 2, 6
- [50] Minghao Xu, Hang Wang, Bingbing Ni, Riheng Zhu, Zhenbang Sun, and Changhu Wang. Cross-category video highlight detection via set-based learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7970–7979, 2021. 2, 6
- [51] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Tubedetr: Spatio-temporal video grounding

- with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16442–16453, 2022. 2
- [52] Ting Yao, Tao Mei, and Yong Rui. Highlight detection with pairwise deep ranking for first-person video summarization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 982–990, 2016. 2
- [53] Qinghao Ye, Xiyue Shen, Yuan Gao, Zirui Wang, Qi Bi, Ping Li, and Guang Yang. Temporal cue guided video highlight detection with low-rank audio-visual fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7950–7959, 2021. 2, 6
- [54] Yitian Yuan, Lin Ma, Jingwen Wang, Wei Liu, and Wenwu Zhu. Semantic conditioned dynamic modulation for temporal sentence grounding in videos. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [55] Da Zhang, Xiyang Dai, Xin Wang, Yuan-Fang Wang, and Larry S Davis. 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*, pages 1247–1257, 2019. 6
- [56] Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Span-based localizing network for natural language video localization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6543–6554, Online, July 2020. Association for Computational Linguistics. 2, 6
- [57] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. Video summarization with long short-term memory. In *European conference on computer vision*, pages 766–782. Springer, 2016. 1, 6
- [58] Renrui Zhang, Han Qiu, Tai Wang, Ziyu Guo, Xuanzhuo Xu, Yu Qiao, Peng Gao, and Hongsheng Li. Monodetr: Depth-guided transformer for monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [59] Songyang Zhang, Houwen Peng, Jianlong Fu, and Jiebo Luo. Learning 2d temporal adjacent networks for moment localization with natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12870–12877, 2020. 1, 2, 6
- [60] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *International Conference on Learning Representations, ICLR 2021*, 2021. 2