

Unleash the Potential of CLIP for Video Highlight Detection

Donghoon Han^{1*} Seunghyeon Seo^{2*} Eunhwan Park¹ Seong-Uk Nam¹ Nojun Kwak^{2†}

¹Buzzni AI Lab ²Seoul National University

{owen, jude, zaid}@buzzni.com {zzzlssh, nojunk}@snu.ac.kr

Abstract

Multimodal and large language models (LLMs) have revolutionized the utilization of open-world knowledge, unlocking novel potentials across various tasks and applications. Among these domains, the video domain has notably benefited from their capabilities. In this paper, we present Highlight-CLIP (HL-CLIP), a method designed to excel in the video highlight detection task by leveraging the pre-trained knowledge embedded in multimodal models. By simply fine-tuning the multimodal encoder in combination with our innovative saliency pooling technique, we have achieved the state-of-the-art performance in the highlight detection task, the QVHighlight Benchmark, to the best of our knowledge.

1. Introduction

Recent advancements in the Natural Language Processing (NLP) area have been significantly revolutionized by the adoption of the pre-training method on large-scale text corpus. Based on the results that have convincingly demonstrated their superiority of training on large-scale datasets, recent pioneering works on multimodal models have unveiled significant capabilities for zero-shot text-image matching. Thanks to the zero-shot text-image matching ability, multimodal models have shown promising results across various downstream tasks, such as visual question answering [20], image captioning [10], text-image retrieval [2], etc.

Despite these improvements, the pre-training method mainly relies on “image–text matching”, meaning that the multimodal models have often suffered from the lack of spatial and temporal knowledge. For example, tasks such as detecting highlights in videos not only require recognizing objects and their descriptions but also understanding how these elements interact over time. Our underlying assumption is that integrating both temporal and spatial knowledge would enhance their performance on tasks that need tempo-

ral awareness.

To this end, we introduce a finetuning strategy for video highlight detection tasks, namely Highlight-CLIP (HL-CLIP). Different from prior studies, our aim is exploiting the full potential of multimodal models. Different from prior studies [9, 11, 14, 16, 17], we only utilize a pre-trained multimodal encoder to achieve better performance in the video highlight detection task on the QVHighlight Benchmark [9], thereby emphasizing the importance of utilizing the capability of pre-trained multimodal models.

The contributions of our paper can be briefly summarized as follows:

- We propose HL-CLIP, a simple finetuning framework for video highlight detection tasks, which leads to enhanced training efficiency.
- We further boost the highlight detection performance at inference time without additional training by our proposed saliency pooling technique.
- Our HL-CLIP achieves the state-of-the-art performance on the QVHighlight benchmark, a highlight detection task.

2. Related Works

2.1. QVHighlight Dataset

QVHighlight dataset [9] consists of over 10,000 videos with annotations, including human-written queries and corresponding segments marked with saliency ratings. The video samples range from lifestyle vlogs to news videos from YouTube, aiming to secure a wide range of content suitable for user annotation. Each video segment spans 150 seconds and features human-annotated free-form queries selected using targeted search queries such as ‘daily vlog’ and ‘news hurricane.’ This annotation process, conducted via Amazon Mechanical Turk, ensures robust data quality characterized by high inter-annotator agreement. Additionally, we assessed the saliency of selected video segments using a Likert-scale rating system, ensuring comprehensive evaluation by involving multiple workers for each clip.

*Equal contribution. †Corresponding author.

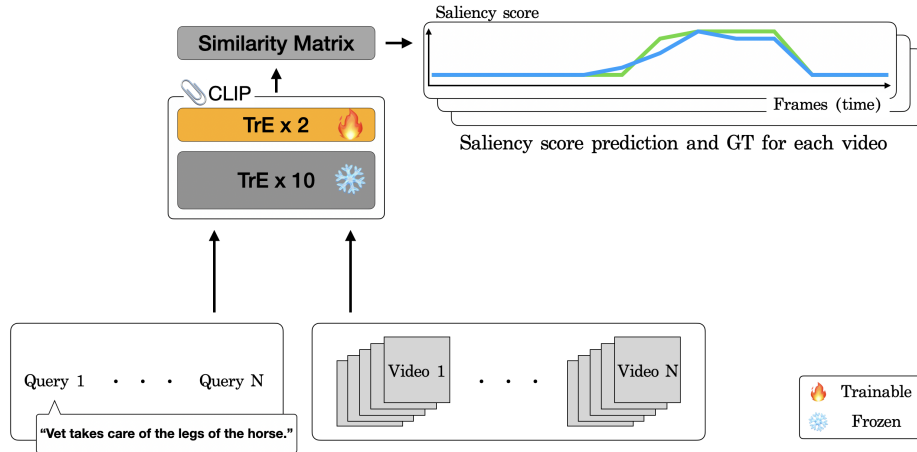


Figure 1. The overall architecture of Highlight-CLIP (HL-CLIP). The last few layers of transformer encoders are finetuned to estimate the saliency score between a frame and a query. The green and blue lines denote the predicted and the ground truth saliency score of a video frame given a query, respectively.

2.2. DETR-Based Highlight Detection Approaches

DETR (Detection Transformer) [3] has provided a novel foundation for object detection by framing the task as a direct set prediction problem. Building upon this architecture, several models like Moment-DETR, QD-DETR, CG-DETR, and others [9, 11, 14, 16, 17] have been proposed in the context of video highlight detection. These approaches leverage the self-attention mechanism inherent in the transformer architecture to process frame features, enabling the model to consider the entire sequence of video frames holistically.

In addition to self-attention across frame features, these models employ cross-modal attention to integrate the contextual information from natural language queries. The frame features are typically obtained by fusing the visual representations from a CLIP (ViT-B/32) [18] encoder and temporal features from a SlowFast [5] model for each frame. The feature representation of the natural language queries is also extracted using the CLIP encoder, ensuring compatibility between the modalities.

These DETR-based approaches are trained with the dual objective of detecting video highlights and retrieving specific moments. By doing so, they are able to localize the most significant segments within a video stream that correspond to a given textual query, demonstrating advanced capabilities in both understanding and indexing video content.

2.3. Efficient Training for Highlight Detection

There exists a line of research to enhance the training efficiency of video highlight detection frameworks. Among them, the prompt tuning methods [21, 22] were suggested as a way to efficiently exploit the pre-trained knowledge out of a model. In the video highlight domain, the Vi-

sual Context Learner (VCL) [6] was suggested to adapt the prompt tuning strategy. The VCL trains the trainable soft prompts prepended at the text input representation. The VCL is a detector-free architecture that applied the Context Optimization (CoOp) [22] framework while freezing the whole pre-trained model and using a simple network to localize the salient moments. This parameter-efficient approach showed improved performance while only using approximately 2,000 trainable parameters. This work showed the possibility that CLIP itself can work as a highlight detector. From this line of work, we suggest HL-CLIP a finetuning strategy for CLIP on the video highlight detection task that aims to maximize the potential of the pre-trained model.

3. Method

3.1. Motivation

Since the pre-trained multimodal encoders contain *implicitly*-trained temporal knowledge derived from the large-scale datasets, they achieve promising results in video highlight detection, which requires understanding of temporal information. Based on this underlying assumption, we can expect that *explicitly* providing temporal information by additional training will lead to both better generalization and better performance in time-related tasks. To this end, we introduce finetuning strategy to enhance temporal knowledge, allowing us to enrich the specialized temporal knowledge for video highlight detection tasks. The overall architecture is shown in Fig. 1.

3.2. Task Definition

The task of our study aims to identify the most salient and relevant moments aligned with a given user query in a

video, namely the highlight detection task. The QVHighlight dataset consists of the diverse collection of 10,000 videos with human-written free-form natural language queries [9]. Each query is associated with specific moments of various length, labeled with their relevance and saliency. Note that the main difference between moment retrieval and highlight detection is that highlight detection deals with identifying the specific moment related to a query which is annotated with high saliency by human annotators while moment retrieval finds all relevant segments in a video.

Now, suppose that we have a set of videos \mathcal{V} , a set of queries \mathcal{Q} and a set of saliency scores \mathcal{Y} each consisting of N elements as follows:

$$\begin{aligned} \mathcal{V} &= \{V_i\}_{i=1}^N, & V_i &= \{x_{i,j}\}_{j=1}^K, \\ \mathcal{Q} &= \{q_i\}_{i=1}^N, & & \\ \mathcal{Y} &= \{Y_i\}_{i=1}^N, & Y_i &= \{y_{i,j}\}_{j=1}^K. \end{aligned} \quad (1)$$

Here, \mathcal{V} contains N videos, each represented by V_i , which consists of K frames $x_{i,j}$. The index i ranges from 1 to N , identifying each video in the dataset, and j ranges from 1 to K , indicating the frame index within a video. Additionally, \mathcal{Q} is the set of query sentences associated with the videos, with q_i being the query corresponding to V_i . \mathcal{Y} denotes the set of saliency scores for N videos, each represented by Y_i is the saliency scores of each video labeled by human annotators. $y_{i,j}$ is the ground-truth saliency score of the j -th frame of the i -th video. The saliency score is originally annotated in a five-point Likert-scale. We normalized this scale to 0 to 1. However, not all frames have a saliency score, and only salient moments are labeled. For training, we set the saliency score of unlabeled frames to ‘Very Bad’.

3.3. CLIP as a Highlight Detector

Our approach to highlight detection is unique in that it operates without a dedicated highlight detector module. Instead, we finetuned a multimodal encoder to process both visual and textual input, thus allowing the model to infer the saliency of different video segments directly from the query and the video content itself. Through this lens, we aim to advance the field of video analysis by developing a model that can intuitively and effectively pinpoint highlights in a vast array of video content.

Unlike previous methods using both CLIP and SlowFast [5] features to train separate highlight detectors, we use only CLIP (ViT-B/32)’s visual encoder and text encoder to perform the highlight detection task as follows:

$$\begin{aligned} h_{i,j}^v &= \mathcal{F}_{\text{vision}}(x_{i,j}; \theta_{\text{vision}}) \\ h_i^t &= \mathcal{F}_{\text{text}}(q_i; \theta_{\text{text}}), \end{aligned} \quad (2)$$

where $h_{i,j}^v$ denotes the visual feature vector for the j -th frame of the i -th video, $\mathcal{F}_{\text{vision}}(\cdot)$ is the visual encoder parameterized by θ_{vision} , $x_{i,j}$ represents the j -th frame of the

i -th video, h_i^t denotes the text feature vector for the i -th video, $\mathcal{F}_{\text{text}}(\cdot)$ is the text encoder, and q_i is the query sentence associated with the i -th video. As shown in Fig. 1, we finetuned the last two layers of both encoders.

Then the cosine similarity between $h_{i,j}^v$ and h_i^t is used as the predicted saliency score for estimating relevancy between the video frame and the given query:

$$\hat{y}_{i,j} = \text{sim}(h_{i,j}^v, h_i^t), \quad (3)$$

where the cosine similarity $\hat{y}_{i,j}$ indicates the saliency of $x_{i,j}$ given query q_i . Note that h_i^t is broadcast across the K frames of V_i .

However, although we achieve a highly competitive performance in highlight detection, the challenge remains in that our approach has difficulty in moment retrieval due to its structural limitation in localizing the moment with a given query in the start and end time.

3.4. Training and Inference

Specifically, our approach involves finetuning the last few layers of a pre-trained CLIP model with a video highlight dataset. Given a video sample consisting of K frames, we observe that adjacent frames exhibit high similarity due to their sequential nature. To capitalize on this characteristic, we arrange the frame features in a batch-wise stack of dimensions $K \times N$, where N is the number of video samples. This configuration aids in distinguishing the subtle differences between similar-looking frames, effectively allowing the model to discern the salient moments that constitute a video highlight.

Additionally, to ensure alignment between the video frames and the textual queries, we replicate the query feature vector K times to match the temporal dimension of the video sample, resulting in a similar $K \times N$ structure. This stacking strategy enables the model to learn in a manner that sharpens its focus on the informative terms within the query that are most relevant for highlight detection.

Loss function. To train our HL-CLIP, we used mean squared error (MSE) between predicted saliency score and ground truth as follows:

$$\mathcal{L}_{\text{saliency}} = \frac{1}{N \cdot K} \sum_{i=1}^N \sum_{j=1}^K (\hat{y}_{i,j} - y_{i,j})^2. \quad (4)$$

Inference. At the inference stage, we use the temporal-aggregated saliency score derived by the **saliency pooling**, *i.e.* simply average-pooling a set of saliency scores from neighboring frames. By the simple pooling technique, we consider the semantic similarities among the adjacent frames and estimate the saliency score more robustly.

Method	Split		test						val					
	MR			HD			MR			HD				
	R1	mAP		>= Very Good			R1	mAP		>= Very Good				
@0.5	@0.7	@0.5	@0.75	Avg.	mAP	HIT@1	@0.5	@0.7	@0.5	@0.75	Avg.	mAP	HIT@1	
BeautyThumb [19]	-	-	-	-	-	14.36	20.88	-	-	-	-	-	-	-
DVSE [13]	-	-	-	-	-	18.75	21.79	-	-	-	-	-	-	-
MCN [1]	11.41	2.72	24.94	8.22	10.67	-	-	-	-	-	-	-	-	-
CAL [4]	25.49	11.54	23.40	7.65	9.89	-	-	-	-	-	-	-	-	-
XML [8]	41.83	30.35	44.63	31.73	32.14	34.49	55.25	-	-	-	-	-	-	-
XML+[8]	46.69	33.46	47.89	34.67	34.90	35.38	55.06	-	-	-	-	-	-	-
Moment-DETR [9]	52.89	33.02	54.82	29.40	30.73	35.69	55.60	53.94	34.84	-	-	32.20	35.65	55.55
UMT [15]	56.23	41.18	53.38	37.01	36.12	38.18	59.99	60.26	44.26	-	-	38.59	39.85	64.19
QD-DETR [17]	62.40	44.98	62.52	39.88	39.86	38.94	62.40	62.68	46.66	62.23	41.82	41.22	39.13	63.03
UniVGT [12]	58.86	40.86	57.6	35.59	35.47	38.20	60.96	59.74	-	-	-	36.13	38.83	61.81
EaTR [7]	-	-	-	-	-	-	-	61.36	45.79	61.86	41.91	41.74	37.15	58.65
VCL [6]	-	-	-	-	-	-	-	43.33	25.75	39.23	24.95	21.13	38.84	65.74
CG-DETR [16]	65.43	48.38	64.51	42.77	42.86	40.33	66.21	67.35	52.06	65.57	45.73	44.93	40.79	66.71
HL-CLIP (Ours)	-	-	-	-	-	41.94	70.60	-	-	-	-	-	42.37	72.40

Table 1. Performance comparison on QVHighlights *test* and *val* splits. We report the average highlight detection score of 3 runs with different seeds on validation and test split. We evaluated with HL-CLIP-2 (SP) our best configuration referred on Sec. 4.1. Although our HL-CLIP is not capable of performing moment retrieval as mentioned in Sec. 3.3, we conjecture that our HL-CLIP has much potential in the moment retrieval task as well considering its superior performance in highlight detection. We leave its application to the moment retrieval as a future work.

Method	mAP (Std.)	HIT@1 (Std.)
CLIP	37.01	63.03
HL-CLIP-1	39.70 (0.03)	67.69 (0.03)
HL-CLIP-1 (SP)	41.97 (0.2)	71.64 (0.1)
HL-CLIP-2	40.00 (0.06)	68.99 (0.28)
HL-CLIP-2 (SP)	42.40 (0.05)	72.42 (0.25)

Table 2. Comparison with multiple finetuning variants. Highlight Detection (>=Very Good) on QVHighlights *val* split. We reported standard deviation of 4 runs with different seeds. SP denotes for Saliency Pooling.

4. Experiment

The experimental section of our work presents an initial exploration into the adaptation of pre-trained multimodal encoders for the task of video highlight detection. Given the brevity of this short paper, our focus is to provide a snapshot of ongoing research efforts rather than an exhaustive evaluation. Each subsection offers insights into different facets of our approach and reflects the potential of these methods to enhance video analysis tasks.

4.1. CLIP Finetuning

In Tab. 2, the quantitative results demonstrate that HL-CLIP-2 outperforms HL-CLIP-1 with a mean Average Precision (mAP) of 40.00 compared to 39.70, and a HIT@1 score of 68.99 and 67.69, respectively. Notably, the integration of saliency pooling further enhances performance, as seen with HL-CLIP-2-SP achieving the highest mAP of 42.40 and HIT@1 of 72.42, confirming its superior efficacy in the highlight detection task. Our HL-CLIP achieves promising results within a detector-free framework, using fewer parameters.

4.2. Comparison with Baselines

In Tab. 1, HL-CLIP achieved the state-of-the-art performance among the previously suggested methods on highlight detection task for both *test* and *val* splits. Due to the structural characteristic, our HL-CLIP cannot directly localize the specific moment with a given queries (Moment Retrieval task). However, since the moment retrieval task is highly relevant to selecting the most salient moment, *i.e.* highlight detection, there exists a possibility of adapting HL-CLIP to the moment retrieval task with further refinement and auxiliary mechanisms.

5. Conclusion

In this work, we explore the potential of CLIP as a video highlight detector. Thanks to its general knowledge from large-scale dataset, our proposed HL-CLIP is equipped with outstanding ability for highlight detection by simply finetuning part of the CLIP and saliency pooling technique. We achieve the state-of-the-art performance on the representative benchmark QVHighlight, showing the competitiveness of CLIP-only framework.

Future work. Our current framework can be extended to advance the model’s capabilities in complex video understanding tasks, *e.g.* leveraging the finetuned encoder to target precise moment retrieval. Based on our work, we plan to suggest score-based efficient network that can localize the moment with a saliency score and finetuned features.

Acknowledgement

S. Seo and N. Kwak were supported by NRF grants (2021R1A2C3006659, 2022R1A5A7026673) funded by MSIT of the Korean Government.

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. 4
- [2] Alberto Baldrati, Lorenzo Agnolucci, Marco Bertini, and Alberto Del Bimbo. Zero-shot composed image retrieval with textual inversion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15338–15347, 2023. 1
- [3] 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
- [4] 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. 4
- [5] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6202–6211, 2019. 2, 3
- [6] Donghoon Han, Seong-Uk Nam, Eunhwan Park, and Nojun Kwak. Parameter-efficient multi-modal highlight detection via prompting. In *Proceedings of the 35th Annual Conference on Human and Cognitive Language Technology*, 2023. 2, 4
- [7] Jinhyun Jang, Jungin Park, Jin Kim, Hyeongjun Kwon, and Kwanghoon Sohn. Knowing where to focus: Event-aware transformer for video grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 4
- [8] 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. 4
- [9] 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, 4
- [10] Wei Li, Linchao Zhu, Longyin Wen, and Yi Yang. Decap: Decoding CLIP latents for zero-shot captioning via text-only training. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 1
- [11] Kevin Qinghong Lin, Pengchuan Zhang, Joya Chen, Shraman Pramanick, Difei Gao, Alex Jinpeng Wang, Rui Yan, and Mike Zheng Shou. Univtg: Towards unified video-language temporal grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2794–2804, 2023. 1, 2
- [12] Kevin Qinghong Lin, Pengchuan Zhang, Joya Chen, Shraman Pramanick, Difei Gao, Alex Jinpeng Wang, Rui Yan, and Mike Zheng Shou. Univtg: Towards unified video-language temporal grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 4
- [13] 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. 4
- [14] 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. 1, 2
- [15] 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. 4
- [16] WonJun Moon, Sangeek Hyun, SuBeen Lee, and Jae-Pil Heo. Correlation-guided query-dependency calibration in video representation learning for temporal grounding. *arXiv preprint arXiv:2311.08835*, 2023. 1, 2, 4
- [17] WonJun Moon, Sangeek Hyun, SangUk Park, Dongchan Park, and Jae-Pil Heo. Query-dependent video representation for moment retrieval and highlight detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23023–23033, 2023. 1, 2, 4
- [18] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [19] 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 conference on information and knowledge management*, pages 659–668, 2016. 4
- [20] Anthony Meng Huat Tiong, Junnan Li, Boyang Li, Silvio Savarese, and Steven CH Hoi. Plug-and-play vqa: Zero-shot vqa by conjoining large pretrained models with zero training. *arXiv preprint arXiv:2210.08773*, 2022. 1
- [21] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16816–16825, 2022. 2
- [22] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. 2