
Zero-shot Video Moment Retrieval With Off-the-Shelf Models

Anuj Diwan*, Puyuan Peng*, Raymond J. Mooney
Department of Computer Science, The University of Texas at Austin
{anuj.diwan, pyp}@utexas.edu, mooney@cs.utexas.edu

Abstract

For the majority of the machine learning community, the expensive nature of collecting high-quality human-annotated data and the inability to efficiently finetune very large state-of-the-art pretrained models on limited compute are major bottlenecks for building models for new tasks. We propose a zero-shot simple approach for one such task, Video Moment Retrieval (VMR), that does not perform any additional finetuning and simply repurposes off-the-shelf models trained on other tasks. Our three-step approach consists of moment proposal, moment-query matching and postprocessing, all using only off-the-shelf models. On the QVHighlights (Lei et al., 2021) benchmark for VMR, we vastly improve performance of previous zero-shot approaches by at least 2.5x on all metrics and reduce the gap between zero-shot and state-of-the-art supervised by over 74%. Further, we also show that our zero-shot approach beats non-pretrained supervised models on the Recall metrics and comes very close on mAP metrics; and that it also performs better than the best pretrained supervised model on shorter moments. Finally, we ablate and analyze our results and propose interesting future directions.

1 Introduction

Given a video and a natural language query, the task of Video Moment Retrieval (VMR) involves temporally localizing moments (i.e. video segments) within the given video that are relevant to the query (Hendricks et al., 2017; Escorcia et al., 2019; Lei et al., 2021). This task has practical applications in many domains; for example, it can be used for efficiently skipping to parts of the video that are most relevant; this is especially useful with long videos (sports events, television/news broadcasts, and so on).

It is becoming increasingly clear that data collection is a major bottleneck for progress in machine learning (Roh et al., 2019); it is quite expensive to collect high-quality human-annotated data for a given task of interest. Furthermore, even assuming the existence of high-quality downstream task data, the ability of machine learning practitioners to finetune the best pretrained models on this data is dwindling due to the prohibitively large size of state-of-the-art pretrained Transformer models like CLIP (Radford et al., 2021), GPT-3 (Brown et al., 2020) and PaLM (Chowdhery et al., 2022); additionally it is often impossible to finetune models if their weights are locked behind an API (e.g. for GPT-3).

We choose the task of Video Moment Retrieval as our task of interest for this case study. This task is related to other tasks like Video Captioning (Xu et al., 2016), Dense Video Captioning (Wang et al., 2021) and Image-Text Matching (Radford et al., 2021) in the field of vision-language research and although data and models from these tasks can benefit VMR, there has been no prior work, to our knowledge, on using these models to solve VMR. We ask whether it is possible to bypass these

*Co-first authors contributed equally. Order was determined by a coin flip.

overarching bottlenecks by a) developing **Zero-shot** methods that require no downstream VMR task data at all and require no finetuning, only inference; and b) using only **Off-the-shelf** models whose weights are publicly available.

In this work, we explore a three-step approach: a) generating moment proposals, b) performing moment-query matching and c) post-processing, all using off-the-shelf models in a novel zero-shot manner. In particular, we use a simple shot transition detector (Castellano, 2022) for the first step, an off-the-shelf CLIP (Radford et al., 2021) model and a Video Captioning model (UniVL: Luo et al. (2020)) for the second step, and post-processing using a simplified version of the Watershed (Roerdink and Meijster, 2000) algorithm. Our overall approach is depicted in Figure 2. We evaluate on the QVHighlights (Lei et al., 2021) benchmark and show that using our simple pipeline, we vastly improve performance on a previous zero-shot approach (between 2.5x-5x on all metrics), and surprisingly beat the best non-pretrained models that are explicitly trained on VMR QVHighlights data on the Recall metrics while coming very close on the mAP metrics. Although our approach still lags behind the current state-of-the-art model on QVHighlights, we hope that our proposed future improvements described in Section 7 can help bridge the gap; we conduct an oracle experiment to check how far we can potentially go with shot-detection-based moment proposals. In addition, we run a length-wise analysis of performance that shows that while all models (including ours) struggle to perform well on short moments, our models perform much better than state-of-the-art models on short moments.

In summary, our contributions include: a) proposing novel zero-shot approaches for Video Moment Retrieval that only use off-the-shelf public models b) demonstrating that zero-shot approaches are extremely strong baselines for VMR and can even beat models supervised on in-domain task data on some metrics and c) showing a length-based discrepancy in performance and improving performance for short moments beyond the best approach.

2 Related Work

Video Moment Retrieval: Datasets. Datasets for VMR include DiDeMo (Hendricks et al., 2017), Charades-STA (Gao et al., 2017), ActivityNet Captions (Krishna et al., 2017), TVR (Lei et al., 2020) and QVHighlights (Lei et al., 2021). Each example in most of these datasets consists of a video, a natural language query, and ground truth moment timestamps (multiple for QVHighlights and one for all other datasets) for localizing the query. QVHighlights also provides additional *saliency scores* that are frame-level annotations on a five-point Likert scale for how well the frame matches the query.

Video Moment Retrieval: Approaches. Recent state-of-the-art approaches for Video Moment Retrieval including Moment-DETR (Lei et al., 2021) and UMT (Liu et al., 2022)) involve designing an end-to-end deep neural network that produces segment proposals and confidence scores simultaneously and training this network end-to-end using a moment retrieval dataset. Older approaches like MCN (Hendricks et al., 2017), CAL (Escorcia et al., 2019) and XML (Lei et al., 2020) are not end-to-end; they consist of a multi-step pipeline that usually consists of generating moment proposals followed by either regressing/predicting moment indices or scoring moment proposals. Often these methods also have a post-processing non-maximal suppression step. Each component is trained separately using a moment retrieval dataset.

Zero-shot transfer with CLIP. The rise of large pretrained models trained on large quantities of data like CLIP (Radford et al., 2021), BERT (Devlin et al., 2019), GPT-3 (Brown et al., 2020) and other so-called Foundation Models has unlocked the ability to transfer knowledge to new tasks without either ever training on the new task (‘zero-shot’) or with very few examples of the task (‘few-shot’). CLIP, in particular, has been used in innovative zero-shot ways for out-of-distribution detection (Esmailpour et al., 2022), text-to-shape generation (Sanghi et al., 2022) and referring-expression comprehension (Subramanian et al., 2022) in addition to zero-shot image classification as studied in the original CLIP paper (Radford et al., 2021).

3 Video Moment Retrieval

The task of Video Moment Retrieval (VMR) involves temporally localizing moments (i.e. video segments) within a video given a natural language query. Formally, given a video V and a query Q , a VMR model must be able to output K segments v_1, \dots, v_K (where each segment v_i consists of a start and end timestamp start_i and end_i) with corresponding scores s_1, \dots, s_K that denote how

similar the model believes the segment is to the query. We describe commonly used datasets and approaches for this task in Section 2 and the evaluation metrics in Section 5.2. Figure 1 shows an example from a recently released dataset for this task, QVHighlights (Lei et al., 2021); provided the query ‘A shark is swimming underwater’ and a nature documentary video, the model must be able to output the predicted relevant video moments and the confidence score.



Figure 1: Moment retrieval example from QVHighlights (Lei et al., 2021)

4 Our Approach

As discussed in Section 2, state-of-the-art VMR models devise an end-to-end neural network to produce proposals and scores simultaneously while older methods separately train proposal and scoring models. In contrast, we are interested in developing an **untuned** method for VMR i.e. a **zero-shot** method that does not require any VMR training data at all. In addition, we would like to build it using publicly available **off-the-shelf** models. With that spirit, we develop a three-step approach consisting of a) generating moment proposals b) performing moment-query matching and c) postprocessing. This overall procedure is depicted in Figure 2. This closely resembles older VMR approaches discussed in Section 2, but with the important distinction that unlike those approaches we use off-the-shelf models for each step without any further training required.

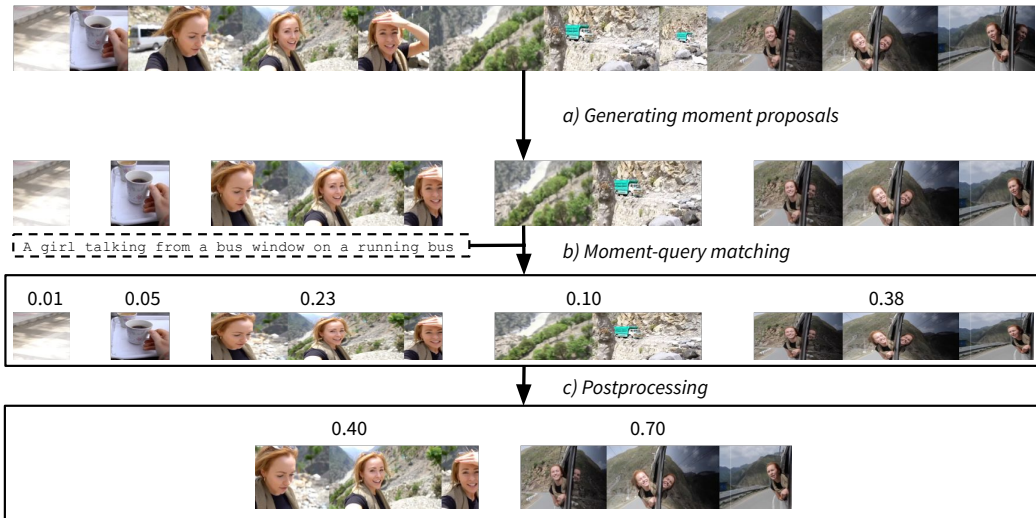


Figure 2: Our overall three-step procedure.

4.1 Generating Moment Proposals.

Our first step is generating segment or moment proposals from the provided video i.e. given the video V of length T seconds and a natural language query Q , we generate K moment proposals with v_1, v_2, \dots, v_K with lengths t_1, t_2, \dots, t_K . To do this, we propose a simple method **ShotDetect**; using preexisting shot transition detection¹ software (Castellano, 2022), we obtain several disjoint

¹https://en.wikipedia.org/wiki/Shot_transition_detection

segments of the video, where, informally, each segment consists of contiguous frames with smooth camera motion. There is a rich body of literature on improving shot transition detection (a few recent papers include Gruzman and Kostenkova (2014), Reddy and Jadhav (2015), Souček and Lokoč (2020), Hassanien et al. (2017) and Singh et al. (2020)) and we choose to use the content-aware detector from the open-source PySceneDetect² toolkit which has been used for other video understanding tasks as well (Wu and Krahenbuhl, 2021). This detector works by triggering scene changes when adjacent frames have a significant color change that is greater than a provided threshold λ ; we treat λ as a hyperparameter.

In order to ablate the choice of moment proposal method, we also devise a baseline moment proposal method called **SlidingWindow**; here, we simply generate 15-second long proposal windows with a stride of 10 seconds.

4.2 Moment-Query matching

In this step, we would like to score every moment proposal using a similarity score $s \in [0.0, 1.0]$ i.e. a similarity score s_k between every segment v_k and the provided query Q . We explore two different methods, **VideoCaptioning** and **CLIP** to obtain this similarity score.

VideoCaptioning. We first run an off-the-shelf video captioning model UniVL (Luo et al., 2020) on each video segment $\{v_1, v_2, \dots, v_K\}$ producing natural language captions $\{l_1, l_2, \dots, l_K\}$. Next, we use an off-the-shelf sentence embedding model MPNet (Song et al., 2020) denoted as E (note that the sentence embedding is obtained by mean-pooling individual token embeddings) to embed both the video segment captions and the natural language query into $\{E(l_1), E(l_2), \dots, E(l_K)\}$ and $E(Q)$. Finally, we compute the cosine similarity between the query embedding and each caption embedding to obtain the final similarity scores $\{s_1, s_2, \dots, s_K\}$. Figure 3a illustrates the similarity score computation with VideoCaptioning.

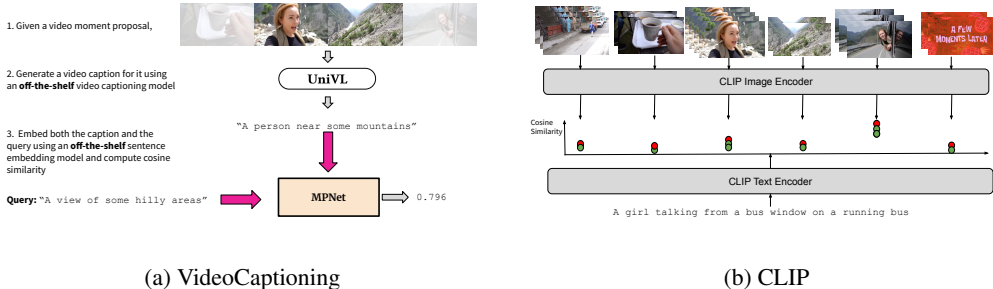


Figure 3: Moment-Query Matching

CLIP. For every video segment $v_k, k \in \{1, \dots, K\}$ with a length of t_k seconds, we first sample $M_k = \lceil t_k \rceil$ frames from v_k at a rate of 1 fps; denote these frames as $f_1^{(k)}, \dots, f_{M_k}^{(k)}$. We then embed every sampled frame using the CLIP (Radford et al., 2021) image encoder. Similarly, we use the CLIP text encoder to embed the query Q . Taking the cosine similarity between the embeddings of each frame and query, we get $\{s_1^{(k)}, s_2^{(k)}, \dots, s_{M_k}^{(k)}\}$. We then calculate the similarity score between segment v_k and query Q as $s_k = f(s_1^{(k)}, s_2^{(k)}, \dots, s_{M_k}^{(k)})$ where f is an aggregation function. We only consider a simple aggregation function f , namely max. Applying this procedure to every video segments, we get similarity scores $\{s_1, s_2, \dots, s_K\}$. Figure 3b illustrates this process with max as the aggregation function f .

4.3 Post-processing

We can use the computed video segments and corresponding similarity scores from the previous step as-is without post-processing. Alternatively, we propose a post-processing step called **Simple Watershed** which is a simplified version of the Watershed algorithm (Roerdink and Meijster, 2000)

²<http://scenedetect.com/en/latest/>

that is a popular technique in image segmentation (Beucher, 1992) and is used to post-process the CLIP baseline for VMR in Lei et al. (2021). We use a simplified variant of Watershed as it is easier to describe and work with.

In SimpleWatershed, given a sequence of segments v_1, \dots, v_k and similarity scores s_1, \dots, s_k and a similarity score threshold γ (which is a hyperparameter), for every consecutive run of segments v_i, \dots, v_j where $s_k \geq \gamma \forall i \leq k \leq j$, we merge the run into a single segment and set its similarity score to $\max(s_i, \dots, s_j)$.

5 Experiments

5.1 Experimental Setup

For Video Captioning, we use the UniVL model (Luo et al., 2020) pretrained on HowTo100M (Miech et al., 2019) finetuned for video captioning on the MSRVTT (Xu et al., 2016) using the UniVL codebase³. For CLIP (Radford et al., 2021), we use the CLIP-ViT-B-32⁴ pretrained model. We obtained the best value of λ , the PySceneDetect color change threshold as 53 for both moment-query matching methods without post-processing and $\lambda = 32, \gamma = 0.7$ for SimpleWatershed. We ablate λ in Appendix A. We run experiments on an RTX 8000 with 45 GB of GPU RAM.

5.2 Datasets and Evaluation

We evaluate our model on the QVHighlights (Lei et al., 2021) dataset. Each example consists of a video, a natural language query, and multiple ground truth moment timestamps for localizing the query. In addition, QVHighlights also provides additional saliency scores that are frame-level annotations on a five-point Likert scale for how well the frame matches the query. We note that collecting frame-level annotation data is extremely expensive and not done in previous VMR datasets; therefore, we place models that are trained on this additional data in a different category. We recap the QVHighlights dataset by borrowing figures and statistics from Lei et al. (2021) in Appendix B.

The QVHighlights dataset has a train-val-test split of 70% – 15% – 15%. However, the test set annotations are private. Further, while the validation set has 1550 videos, we were able to obtain only a subset consisting of 1434 videos; the other videos were either no longer on YouTube or failed to download. Thus, we evaluate our models on this subset that we call the *filtered validation* set or **val-filt** for short. Table 1 shows the same model Moment-DETR from Lei et al. (2021) evaluated on these three sets (test, val, val-filt). We find that all three sets result in approximately the same evaluation metric results and thus could be fairly interchangeable with respect to analysis. However, to ensure completely fair comparison, we report metrics on the val-filt subset (which we have access to) for as many models as possible and the test subset for other models (whose metrics are directly taken from Lei et al. (2021)) with the understanding that metrics on val-filt and test can be likely fairly compared to each other.

To evaluate our models, we use the same evaluation code and metrics as Lei et al. (2021); a) Mean Average Precision (mAP) with IoU thresholds 0.5 and 0.75, as well as the average mAP over multiple IoU thresholds and b) Recall@1 with IoU thresholds 0.5 and 0.7. Please refer to Lei et al. (2021) for a detailed description of these evaluation metrics.

Table 1: Comparing results on Moment-DETR w/ PT (Lei et al., 2021) on different evaluation sets. The test and val set results are taken directly from Lei et al. (2021) while we evaluate ourselves on the val-filt set.

Evaluation Set	R1		mAP		
	@0.5	@0.7	@0.5	@0.75	avg
test	59.78	40.33	60.51	35.36	36.14
val	59.68	40.84	-	-	36.30
val-filt	59.74	41.10	59.90	35.42	36.19

³<https://github.com/microsoft/UnivL>

⁴<https://huggingface.co/sentence-transformers/clip-ViT-B-32>

5.3 Baselines

We compare our model with three types of baselines: a) VMR-Supervised, b) VMR-Supervised + Saliency and c) Zero-Shot approaches.

VMR-Supervised. These models are trained on the QVHighlights training dataset excluding the saliency score annotations. We select the same baselines as the ones considered in the QVHighlights paper (Lei et al., 2021); two proposal based models MCN (Hendricks et al., 2017) and CAL (Escorcia et al., 2019) and two span prediction models, XML (Lei et al., 2020) and Moment-DETR w/o saliency (Lei et al., 2021). We reproduced Moment-DETR w/o saliency ourselves since the checkpoint for this setting is not publicly available.

VMR-Supervised + Saliency. These models are trained on the QVHighlights training set, including saliency score annotations. We compare with XML+, Moment-DETR and Moment-DETR w/ Pretraining (pretrained on ASR Captions) all from Lei et al. (2021) and UMT (Liu et al., 2022), the current SoTA on QVHighlights that is trained on an additional audio modality. Note that unlike zero-shot approaches, the architecture, loss and training scheme of these models have been specifically developed for Video Moment Retrieval and have been trained on in-domain supervised data.

Zero-shot. These models (including ours) have not been explicitly trained on QVHighlights. We compare against the only baseline known to us: the CLIP+Watershed baseline provided in Lei et al. (2021). Our CLIP-based approach is different from this baseline. We use segment proposals obtained from shot detection software and provide results with and without SimpleWatershed postprocessing. On the other hand, CLIP+Watershed computes image-text similarities on frames sampled at 0.5 Hz and then runs the full Watershed algorithm on them; essentially, it just uses 2-second moment proposals instead of obtaining them from shot detectors.

5.4 Oracular Bounds

Our approach is based on the video segments provided by PySceneDetect, which introduces a possible weak link into our approach: if the segments predicted by this tool are of low quality, it will be impossible to ever succeed on QVHighlights, even given a perfect Moment-Query Matcher. To ensure that this is not the case, we compute oracular bounds to gauge whether high scores are achievable. We design our ‘perfect moment-query matcher’ (which has access to the ground truth moments) to be such that predicted similarity scores nearly perfectly reflect the match between video segments and the query.

Non-Postprocessing Oracular Bound. We define the similarity score between a given segment and the query to be the highest IoU⁵ between the segment and the nearest ground truth moment.

Postprocessing Oracular Bound. We merge adjacent segments (for the SimpleWatershed post-processing step) as long as doing so will increase the IoU of the combined segment with the nearest ground truth moment.

6 Results

6.1 Main Results

We present our main results in Table 2. As described in Section 5.3, we categorize the approaches we compare into VMR-Supervised, VMR-Supervised + Saliency and Zero-Shot. Further, as discussed in Section 5.2, we report our approaches and some baselines on the val-filt validation set; for others (marked with an asterisk *) we report on the test set. Note that from Table 1 and by comparing M-DETR w/ PT* with M-DETR w/ PT in Table 2 one can conclude that results on the test set and the val-filt set are fairly comparable.

Comparison with Zero-shot. Comparing zero-shot approaches, we see that all of our approaches significantly outperform the sole zero-shot baseline (CLIP+Watershed) from Lei et al. (2021); our best approach, ShotDetect+CLIP+SimpleWatershed raises the zero-shot performance at least by 2.5x on all metrics and by as much as 5x on R1@0.7 without ever having seen data from VMR,

⁵There can be oracular similarity scores that are better than IoU and thus our oracular bounds are not upper bounds.

Table 2: Results using all evaluation metrics on the QVHighlights val-filt set (for rows without an asterisk *) or the test set (for rows with an asterisk *). Bold values denote the best-performing approaches within each category.

Category	Method	R1		mAP		
		@0.5	@0.7	@0.5	@0.75	avg
Zero Shot	CLIP+Watershed(Lei et al., 2021))	16.88	5.19	18.11	7.00	7.67
	SlidingWindow + VideoCaptioning (Ours)	19.60	6.00	25.94	6.00	9.58
	ShotDetect+VideoCaptioning (Ours)	22.25	14.71	28.90	17.30	18.06
	SlidingWindow + CLIP (Ours)	29.71	8.86	35.26	8.31	13.42
	ShotDetect+CLIP (Ours)	40.24	25.94	41.74	24.11	24.82
	ShotDetect+CLIP+SimpleWatershed (Ours)	48.33	30.96	46.94	25.75	27.96
VMR-Sup	MCN*(Hendricks et al. (2017))	11.41	2.72	24.94	8.22	10.67
	CAL*(Escorcio et al. (2019))	25.49	11.54	23.40	7.65	9.89
	XML*(Lei et al. (2020))	41.83	30.35	44.63	31.73	32.14
	M-DETR w/o saliency loss(Lei et al. (2021))	45.03	25.81	48.42	21.91	24.68
VMR-Sup+Saliency	XML+*(Lei et al. (2021))	46.69	33.46	47.89	34.67	34.90
	M-DETR*(Lei et al. (2021))	52.89	33.02	54.82	29.40	30.73
	M-DETR w/ PT(Lei et al. (2021))	59.74	41.10	59.90	35.42	36.19
	M-DETR w/ PT*(Lei et al. (2021))	59.78	40.33	60.51	35.36	36.14
	UMT*(Liu et al. (2022))	56.23	41.18	53.83	37.01	36.12
	UMT w/ PT*(Liu et al. (2022))	60.83	43.26	57.33	39.12	38.08

or indeed, the video modality (since CLIP was trained solely on images and text). We note that our ShotDetect approach performs significantly better than our SlidingWindow approach (for both VideoCaptioning and CLIP), showing that ShotDetect is indeed helpful for VMR as compared to the baseline. Furthermore, CLIP works much better than UniVL VideoCaptioning; we hypothesize that this is because the generated video captions (trained on MSRVT Xu et al. (2016)) are not in-domain for QVHighlights queries, which may make matching captions to queries harder.

Comparison with VMR-Supervised. Surprisingly, our best performing zero-shot approach ShotDetect+CLIP+Watershed outperforms all VMR-Supervised models (MCN, CAL, XML and Moment-DETR w/o saliency loss) on both Recall metrics. It outperforms Moment-DETR w/o saliency loss on four out of five metrics. This shows that a simple approach like ours that uses no supervised data can perform competitively or even beat supervised approaches.

Comparison with VMR-Supervised + Saliency (+ Pretraining). Here we can see that our approach performs worse than the worst method (XML+) in this category and our approach is significantly worse than the best approach, UMT w/ PT. We first note that adding saliency score supervision (an expensive annotation) provides a huge jump in performance (compare M-DETR to M-DETR w/o saliency loss). Although we bridged the gap between zero-shot and normal supervised approaches, there is still a significant gap between zero-shot and pretrained models or supervision with frame-level annotations like saliency scores. In order to bridge this gap, it will be essential to understand how to obtain saliency score-like information in a zero-shot manner to incorporate into our models, and we leave this to future work.

6.2 Oracular bounds

Table 3 shows the two oracular bounds introduced in section 5.4 for our two best zero shot approaches (with and without SimpleWatershed post-processing). The Non-Postprocessing Oracular Bound, which bounds our ShotDetect+CLIP approach, is comparable to the best approaches from Table 2 in terms of Recall but is much worse at the mAP score, indicating that, without post-processing, we are likely upper bounded by a poor mAP score. The Postprocessing Oracular Bound, which bounds our ShotDetect+CLIP+SimpleWatershed approach outperforms the best supervised UMT w/ PT by a very large margin. However, note that this Oracular Bound will be very hard to attain, even compared to the Non-Preprocessing Oracular Bound, as one needs both perfect similarity scores and segment merging algorithms.

Table 3: Oracular Bounds

Category	Method	R1		mAP		
		@0.5	@0.7	@0.5	@0.75	avg
Zero Shot	ShotDetect + CLIP(Ours)	40.24	25.94	41.74	24.11	24.82
Oracular Bound	Non-Postprocessing Bound ($\lambda = 53$)	63.94	41.49	54.29	29.41	30.98
Zero Shot	ShotDetect+CLIP+SimpleWatershed(Ours)	48.33	30.96	46.94	25.75	27.96
Oracular Bound	Postprocessing Bound ($\lambda = 47$)	87.94	79.01	81.02	67.23	63.99

6.3 Length-wise analysis of performance

Next, we perform a more fine-grained comparison of our approaches with supervised models by evaluating on different subsets of the val-filt set determined by the lengths of the ground truth moments, in order to understand how models perform on short vs. middle vs. long ground truth moments. **Short** moments have less than 10 seconds, **Medium** moments are between 10 and 30 seconds and **Long** moments are greater than 30 seconds, which are, respectively, 27.63%, 62.41% and 36.26% of the total set of moments. Using the avg. mAP metric as our metric of comparison, we present these results in Table 4. We first note that for all approaches, the short mAP is very low in absolute terms and these models score well on the overall avg. mAP metric (that ignores length) only due to long and medium moments. This indicates a systematic failure of all approaches on a specific subset of the QVHighlights dataset and should be taken into consideration when building or using models for this task. Our zero-shot approach ShotDetect+CLIP is much better at predicting short moments, while supervised methods are better at predicting long moments. We do not have a good explanation for this phenomenon, but it could be a good hint to build better inductive bias like ours (implicitly from the ShotDetection step) into the design of supervised models. Note that using the SimpleWatershed post-processing still beats the supervised baselines on the short avg mAP, which is encouraging, but has lower performance as compared to without post-processing (4.90 vs 7.08). Instead, the post-processing helps the long (33.19 vs 27.49) and middle (28.44 vs 26.15) mAPs. This makes intuitive sense as the SimpleWatershed algorithm combines segments of high similarities into longer segments.

Table 4: Length-wise analysis of performance using the avg. mAP metric

Category	Method	Long	Medium	Short
VMR-Supervised	M-DETR w/o Saliency loss	31.38	25.29	2.63
VMR-Supervised w/ Saliency	M-DETR	41.11	32.3	3.28
	M-DETR w/ PT	45.18	37.53	3.50
Zero Shot	ShotDetect+CLIP	27.49	26.15	7.08
	ShotDetect+CLIP+SimpleWatershed	33.19	28.44	4.90

6.4 Relaxing the zero-shot assumption

Although this falls outside our main setting of untuned zero-shot approaches, we also investigate the possibility of fine-tuning CLIP on the QVHighlights training set. This investigation is largely preliminary.

We use the original CLIP image encoder to encode RGB frames of the video, the original CLIP text encoder to encode the query, and use the same CLIP contrastive loss as the training objective. Negative frames are either sampled from a) the set of frames in the same video that have low (original) CLIP similarity with the text query (‘hard’ negative examples as they are from the same video), b) other videos within the same batch (‘easy’ negative examples). We tune the following hyperparameters: the number of early CLIP layers to freeze, the number of hard and easy negative examples, the number of positive examples, and the learning rate. We found that freezing early layers during fine-tuning is helpful, and increasing the number of negative examples is also helpful, although we couldn’t scale it up as much as CLIP’s original training recipe due to computational limitations. However, in all experiments, we did not see a significant improvement over the original zero-shot

Table 5: CLIP fine-tuning performance

Category	R1		mAP		
	@0.5	@0.7	@0.5	@0.75	avg
Zero-shot	40.24	25.94	41.74	24.11	24.82
Fine-tuning	42.12	27.89	43.00	24.68	25.50

result, and the validation performance usually starts to decrease after a few hundred steps. In Table 5, we list our best fine-tuning performance and compare it with zero-shot performance. We obtained this result with the following hyperparameters: freezing the first 8 layers of both CLIP encoders, 30 hard examples, and 200 negative examples, 1 positive example (i.e. we only sample one frame from the interval that the query corresponds to), with a learning rate schedule involving a linear ramp-up to 10^{-6} for 10% steps and a linear decrease to 0 for the remaining 90% steps.

7 Future Work

In future work, we would like to improve our use of the CLIP embeddings (e.g. by increasing our framerate) or use better, modality-matched embeddings to encode the video (e.g. using video-text pretrained models such as VideoCLIP (Xu et al., 2021)). We would also like to dive deeper into the choice of shot transition detector; there is a rich body of literature with many different techniques for handling videos in different domains and settings (Cotsaces et al., 2006). Finally, it would be interesting to improve our finetuning results and investigate how to adapt pretrained models in a sample-efficient manner for VMR.

8 Conclusion

We proposed a three-step zero-shot approach for Video Moment Retrieval and demonstrated how we vastly improve zero-shot performance, reducing the gap between zero-shot and state-of-the-art supervised models by a large margin. Furthermore, we also showed that our zero-shot approach beats non-pretrained supervised models on some metrics and comes close on other metrics and performs better than the best pretrained supervised model on shorter moments. We also ablated our hyperparameters and design choices and proposed interesting future directions.

Acknowledgments and Disclosure of Funding

We thank the students in the UT Austin CS398T “Grounded Natural Language Processing” course for their helpful feedback.

References

- Serge Beucher. 1992. The watershed transformation applied to image segmentation. *Scanning Microscopy*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *NeurIPS 2020*.
- Brandon Castellano. 2022. Pyscenedetect: Intelligent scene cut detection and video splitting tool.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James

- Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. Palm: Scaling language modeling with pathways. *arxiv.2204.02311*.
- Costas Cotsaces, Nikos Nikolaidis, and Ioannis Pitas. 2006. Video shot detection and condensed representation. a review. *IEEE signal processing magazine*, 23(2):28–37.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL 2019*.
- Victor Escorcia, Mattia Soldan, Josef Sivic, Bernard Ghanem, and Bryan C. Russell. 2019. Temporal localization of moments in video collections with natural language. *ArXiv*, abs/1907.12763.
- Sepideh Esmailpour, Bing Liu, Eric Robertson, and Lei Shu. 2022. Zero-shot out-of-distribution detection based on the pre-trained model clip. In *AAAI 2022*.
- Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. 2017. Tall: Temporal activity localization via language query. In *ICCV 2017*.
- Igor S. Gruzman and Anna S. Kostenkova. 2014. Algorithm of scene change detection in a video sequence based on the threedimensional histogram of color images. In *2014 12th International Conference on Actual Problems of Electronics Instrument Engineering (APEIE)*, pages 1–1.
- Ahmed Hassanien, Mohamed Elgharib, Ahmed Selim, Sung-Ho Bae, Mohamed Hefeeda, and Wojciech Matusik. 2017. Large-scale, fast and accurate shot boundary detection through spatio-temporal convolutional neural networks. *arXiv preprint arXiv:1705.03281*.
- Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. 2017. Localizing moments in video with natural language. In *ICCV 2017*.
- Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. 2017. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pages 706–715.
- Jie Lei, Tamara L. Berg, and Mohit Bansal. 2021. Qvhighlights: Detecting moments and highlights in videos via natural language queries. In *NeurIPS 2021*.
- Jie Lei, Licheng Yu, Tamara L Berg, and Mohit Bansal. 2020. Tvr: A large-scale dataset for video-subtitle moment retrieval. In *European Conference on Computer Vision*, pages 447–463. Springer.
- Ye Liu, Siyuan Li, Yang Wu, Chang Wen Chen, Ying Shan, and Xiaohu Qie. 2022. Umt: Unified multi-modal transformers for joint video moment retrieval and highlight detection. In *CVPR 2022*.
- Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Jason Li, Taroon Bharti, and Ming Zhou. 2020. Univl: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2002.06353*.
- Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2630–2640.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision.

- Bindu Reddy and Anita Jadhav. 2015. Comparison of scene change detection algorithms for videos. In *2015 Fifth International Conference on Advanced Computing Communication Technologies*, pages 84–89.
- Jos B.T.M. Roerdink and Arnold Meijster. 2000. The watershed transform: Definitions, algorithms and parallelization strategies. *Fundam. Inf.*, 41(1,2):187–228.
- Yuji Roh, Geon Heo, and Steven Euijong Whang. 2019. A survey on data collection for machine learning: a big data-ai integration perspective. *IEEE Transactions on Knowledge and Data Engineering*, 33(4):1328–1347.
- Aditya Sanghi, Hang Chu, Joseph G Lambourne, Ye Wang, Chin-Yi Cheng, Marco Fumero, and Kamal Rahimi Malekshan. 2022. Clip-forge: Towards zero-shot text-to-shape generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18603–18613.
- Alok Singh, Dalton Meitei Thounaojam, and Saptarshi Chakraborty. 2020. A novel automatic shot boundary detection algorithm: robust to illumination and motion effect. *Signal Image Video Process.*, 14(4):645–653.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. Mpnet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems*, 33:16857–16867.
- Tomáš Souček and Jakub Lokoč. 2020. Transnet v2: an effective deep network architecture for fast shot transition detection. *arXiv preprint arXiv:2008.04838*.
- Sanjay Subramanian, Will Merrill, Trevor Darrell, Matt Gardner, Sameer Singh, and Anna Rohrbach. 2022. Reclip: A strong zero-shot baseline for referring expression comprehension. *arXiv preprint arXiv:2204.05991*.
- Teng Wang, Ruimao Zhang, Zhichao Lu, Feng Zheng, Ran Cheng, and Ping Luo. 2021. End-to-end dense video captioning with parallel decoding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6847–6857.
- Chao-Yuan Wu and Philipp Krahenbuhl. 2021. Towards long-form video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1884–1894.
- Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. 2021. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *ArXiv*, abs/2109.14084.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. Msr-vtt: A large video description dataset for bridging video and language. *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*.

A Ablation for the λ hyperparameter

We show the effect of the PySceneDetect threshold hyperparameter λ by computing the avg mAP with for both ShotDetect+CLIP and ShotDetect+CLIP+SimpleWatershed (with the best postprocessing hyperparameter γ pre-selected) in Figure 4, showing how we obtained optimal λ values of 53 and 32 for the two settings.

B The QVHighlights dataset.

This dataset contains videos from user-created vlogs (categorized into ‘Daily Vlog’ and ‘Travel Vlog’) and news videos that have lots of raw footage. See Table 6 for video statistics by categories, and Figure 5 for statistics of ground truth moment. Both the figures and tables are borrowed from the QVHighlights paper (Lei et al., 2021).

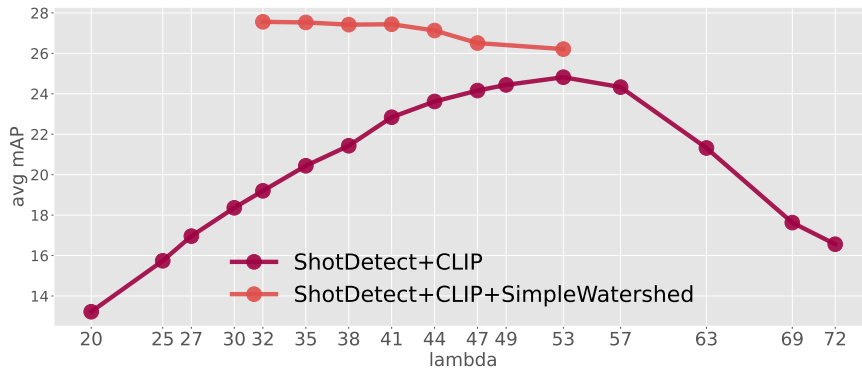


Figure 4: Ablating the threshold λ hyperparameter

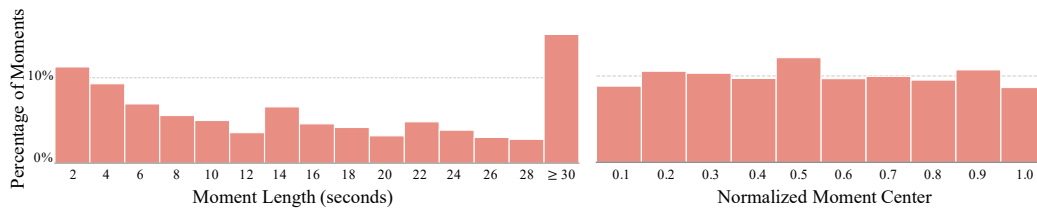


Figure 5: Distribution of moment lengths (*left*) and normalized (by video duration) center timestamps (*right*). The moments vary greatly in length, and they distribute almost evenly along the videos. Borrowed from Lei et al. (2021)

Table 6: Top unique verbs and nouns in queries, in each video category. Borrowed from Lei et al. (2021)

Category	#Queries	Top Unique Verbs	Top Unique Nouns
Daily Vlog	4,473	cook, apply, cut, clean	dog, kitchen, baby, floor
Travel Vlog	4,694	swim, visit, order, travel	beach, hotel, tour, plane
News	1,143	report, gather, protest, discuss	news, interview, weather, police