

# Multi-event Video-Text Retrieval

Gengyuan Zhang<sup>1,2</sup> Jisen Ren<sup>1</sup> Jindong Gu<sup>3\*</sup> Volker Tresp<sup>1,2</sup>

<sup>1</sup> LMU Munich, Munich, Germany

<sup>2</sup> Munich Center for Machine Learning, Munich, Germany

<sup>3</sup> University of Oxford, Oxford, United Kingdom

zhang@dbs.ifi.lmu.de jindong.gu@outlook.com

## Abstract

*Video-Text Retrieval (VTR) is a crucial multi-modal task in an era of massive video-text data on the Internet. A plethora of work characterized by using a two-stream Vision-Language model architecture that learns a joint representation of video-text pairs has become a prominent approach for the VTR task. However, these models operate under the assumption of bijective video-text correspondences and neglect a more practical scenario where video content usually encompasses multiple events, while texts like user queries or webpage metadata tend to be specific and correspond to single events. This establishes a gap between the previous training objective and real-world applications, leading to the potential performance degradation of earlier models during inference. In this study, we introduce the Multi-event Video-Text Retrieval (MeVTR) task, addressing scenarios in which each video contains multiple different events, as a niche scenario of the conventional Video-Text Retrieval Task. We present a simple model, MeRetriever, which incorporates key event video representation and a new MeVTR loss for the MeVTR task. Comprehensive experiments show that this straightforward framework outperforms other models in the Video-to-Text and Text-to-Video tasks, effectively establishing a robust baseline for the MeVTR task. We believe this work serves as a strong foundation for future studies. Code is available at <https://github.com/gengyuanmax/MeVTR>.*

## 1. Introduction

With the proliferation of multimedia data on the Internet every day, Video-Text Retrieval (VTR) task gains increasing importance in searching for desired items from the opposite modality given a video or text query. Over the past years, numerous efforts [49, 33, 9, 14, 31, 3, 7, 30, 16, 12, 51, 17] have been undertaken to improve the

multi-modal retrieval performance on both Video-to-Text and Text-to-Video retrieval tasks. Along with the popularity of powerful vision foundation models like CLIP [36] and ALIGN [23] that align the image-text pairs in a joint feature space, two-stream video-text models [30, 16] targeting to learn video-text correspondences have further become a mainstream method for the Video-Text Retrieval task.

However, this paradigm neglects that videos usually contain more than a single event, while a single textual caption can only capture a fragment of the entire video content. As depicted in Fig. 1, the example video spans a sequence of unrelated and discontinuous events, and each single textual caption merely corresponds to a video segment. This congruity undoubtedly contradicts the fundamental of the traditional VTR task, which aims to establish a mapping between video-text pair.

This fact presents a prevalent scenario in Video-Text Retrieval: videos are intricate, often containing multiple events, whereas texts (typically like search queries or webpage metadata) tend to be specific and fragmentary. Consequently, this dichotomy leads to divergent semantic meanings for a given video-text pair. Such divergence contradicts the conventional training objective of aligning representations for video-text pairs in VTR models, that video features are aligned with the corresponding text feature.

In this study, we formally introduce a new and realistic VTR setting involving multi-event videos, dubbed as *Multi-event Video-Text Retrieval (MeVTR)*. In the introduced MeVTR task, a text describes a single event within a video item, whereas a video corresponds to multiple relevant textual captions, each detailing different events.

We assess the performance of previous VTR models trained under standard VTR settings that learn a bijective video-text correspondence on the MeVTR task without re-training. Our evaluation reveals distinct levels of performance degradation, as illustrated in Table 1. This observation confirms the presence of a performance gap between training within the conventional VTR framework and real-world MeVTR inference scenarios.

\*Corresponding author



Figure 1: An example case of multi-event videos from ActivityNet [4]. The video depicts a sequence of unrelated and discontinuous events, including the progression “a girl is sitting on the beach” → “a young man is practicing tightrope walking” → “a scene of sunset by the beach.” Each textual caption only corresponds to a fragment of the video. Such short and specific textual captions are prevalent in our everyday video data and constitute a common video-text retrieval scenario.

Besides, we retrain a series of previous models to adapt to the MeVTR task. However, we find that they cannot perform equally well on the Video-to-Text and Text-to-Video. We speculate that this is caused by non-injective video-text correspondences in the MeVTR setting, where previous VTR models embed a video and multiple distinct caption texts into a joint representation space and aim at aligning the video feature with the features of all caption texts. Consequently, distinct texts corresponding to the same video are mapped to the same feature, potentially resulting in suboptimal retrieval performance.

Under this circumstance, we introduce a new CLIP-based model dubbed Me-Retriever tailored for the MeVTR task. Considering the dichotomy of multi-event video-text pairs, we formulate our approach as follows: (1) representing videos by a bag of key event features and (2) employing a multi-event video-text retrieval loss. This strategy serves the dual purpose of mitigating the collapse of textual features and harmonizing the learning objectives between the Text-to-Video and Video-to-Text tasks.

Our contributions can be summarized as follows:

1. We introduce a new task, Multi-event Video-Text Retrieval for retrieving multi-event video-text pairs, and define new evaluation metrics in MeVTR;
2. We propose a new model called Me-Retriever that represents videos with selected key events and a MeVTR loss for training on MeVTR;
3. We conduct comprehensive experiments to showcase Me-Retriever’s effectiveness and to establish it as a simple yet robust baseline for future research.

## 2. Related Work

**Video Language Model** A series of work on Vision-Language Models (VLM) has been made like [26, 29, 5, 57, 53] for multi-modal representation learning. Given that videos are a structured sequence of images and typically entail higher computational demands than other visual data formats, the domain of Video-Language models has garnered significant attention. Notable contributions within this category include works by [40, 24, 38, 49, 29, 23, 43, 19, 13] which aim at video foundation models for video tasks. These models represent a subcategory of VLMs, adept at processing video and text inputs.

**Video Representation** In a line of work [42, 52, 18, 25], convolutions are performed on videos to fuse spatial and temporal information. Meanwhile, as transformers [44] have achieved success firstly in text [8], and then vision fields [11, 27, 41, 6], attempts to introduce transformers to extract video representations have been made in [39, 2, 47, 56, 32, 28]. The success of pre-trained Vision-Language models has significantly impacted video representation learning. Vision-Language models [29, 38, 31, 36] have demonstrated that textual supervision is conducive to visual representation and to a uniform video-language representation. Among them, the success of representation models [22, 16, 45, 46] has bridged the gap between video and text representations.

**Video-Text Retrieval** Video-Text Retrieval is a fundamental task of multi-modal learning, posing a good proxy task for learning joint representations of videos and texts. Two-stream Video-Text models [10, 9, 51, 34] concentrate on processing video-text pairs as inputs and aligning vi-

Model	Video-to-Text				Text-to-Video			
	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑
FROZEN-in-time	7.0 ↗ 109.5	15.3 ↘ 0.8	43.3 ↘ 3.1	43.3 ↘ 5.6	7.0 ↗ 111.0	15.9 ↘ 1.8	43.0 ↘ 6.8	59.01 ↘ 11.7
CLIPBERT	4.0 ↗ 449.0	22.1 ↘ 1.4	58.2 ↘ 4.7	69.6 ↘ 7.6	5.0 ↗ 153.0	25.9 ↘ 2.6	55.4 ↘ 9.5	69.8 ↘ 15.1
CLIP4Clip	2.0 ↗ 53.0	42.5 ↘ 7.5	74.1 ↘ 19.8	85.8 ↘ 28.4	13.0 ↗ 53.0	14.8 ↘ 7.5	34.2 ↘ 19.8	45.6 ↘ 28.4
CenterCLIP	2.0 ↗ 51.0	42.8 ↘ 7.8	73.8 ↘ 20.8	85.3 ↘ 29.2	2.0 ↗ 13.0	41.8 ↘ 15.1	73.9 ↘ 34.6	84.7 ↘ 46.4

Table 1: Comparison of model performance on VTR and MeVTR tasks on ActivityNet Captions. We train each model on the ActivityNet Captions with its original codes and evaluate them on MeVTR. We show the performance of VTR on the lefthand of ↘(↗) and the performance of MeVTR on the right. ↑ means higher the better, and ↓ means lower the better. It is shown that model performance deteriorates to a large extent on both Video-to-Text and Text-to-Video tasks.

sual and textual representations within a joint representation space. For instance, Frozen [3] employs a visual encoder that can be adaptively trained on images and videos, incorporating temporal contexts via curriculum learning. TACo [51], in contrast, focuses on aligning content words in texts with corresponding visual content, adding an extra learning objective. With the surge of large Vision-Language foundation models representative of CLIP [36], models [30, 16, 12, 35, 3, 50], leverage the potent pre-trained model CLIP and endeavor to extend its capabilities from static images to the domain of videos.

### 3. Problem

#### 3.1. Problem Formulation

Given a video-text dataset with multi-event videos  $\mathcal{V}$  and texts  $\mathcal{T}$ , where a video  $v_i \in \mathcal{V}$  is described by a set of different texts  $\mathcal{T}_i \subset \mathcal{T}$  and a text is denoted as  $t_j$ , our task consists in retrieving a video given a text query  $t^q$  and retrieving all relevant texts given a video query  $v^q$ . A generic Video-Text Retrieval model is to learn a similarity function  $sim(v_i, t_j)$  that scores the similarity between any relevant video-text pair higher and ranks all the candidate items.

#### 3.2. Evaluation Metrics

Video-Text retrieval is a recall-focused retrieval task. A single item is retrieved in traditional settings, and the average Recall at the top rank  $k$  is calculated. Nevertheless, in the MeVTR setting, we also want to evaluate retrieving multiple different text items given one video. Hence we devise the following metrics: (1) **Recall@k-Average** as the fraction of the number of correspondences ranked at top  $k$ ; (2) **Recall@k-One-Hit** as an indication of whether any of the correspondences will be ranked in the top  $k$ ; (3) **Recall@k-All-Hit** as an indication of whether all items can be retrieved within the top  $k$  rank. All metrics are averaged on the whole test set. In the follow-up experiments, we report the Recall for the Text-to-Video tasks and the metrics mentioned above for the Video-to-Text tasks.

### 3.3. Textual Feature Collapse

To investigate the textual feature collapse in previous models to elucidate the performance drop by calculating the textual feature’s similarity after retraining, we define the average text similarity of a video  $v_i$  with  $n_{v_i}$  captions as:

$$sim_t(v_i) = \sum_{m=1}^{N_{v_i}} \sum_{n=1}^{N_{v_i}} \frac{1}{N_{v_i}^2} sim(t_m, t_n) \quad (1)$$

In Fig. 3, we attempt to visualize the average text similarity of all videos with a different number of events. Different texts are expected to have lower sentence similarity to retain their meanings due to retraining on MeVTR. We chose CLIP4Clip for comparison in our experiments, as it still performs relatively well in MeVTR tasks compared to other models. As a comparison, we show that Me-Retriever, which will be introduced in the next section, can generate more diverse textual features than CLIP4Clip and effectively avoid textual feature collapse.

## 4. Method

We propose Me-Retriever consisting of a Key Event Selection module and MeVTR loss as shown in Fig. 2.

### 4.1. Key Event Video Representation

Intuitively, human brains do not remember experiences with all continuous scenes and consolidate all the details. Instead, brains retain memories of certain discrete events [55]. This insight motivates us to represent a video with a bag of frames as key events. For leveraging frame embeddings extracted from the CLIP visual encoder  $\{\mathbf{f}_i^1, \mathbf{f}_i^2, \dots, \mathbf{f}_i^{|v_i|}\}$ , we define the Key Event Selection module that selects the most distinctive keyframes.

We employ a classic clustering method, K-Medoids [1]. Compared to other clustering methods, the representative of each cluster given by the K-Medoids algorithm is from the sample points, which corresponds to a particular frame and has *explicit* interpretation. The criterion is minimizing the distance of all frames in the cluster to its medoid.  $K$  frames

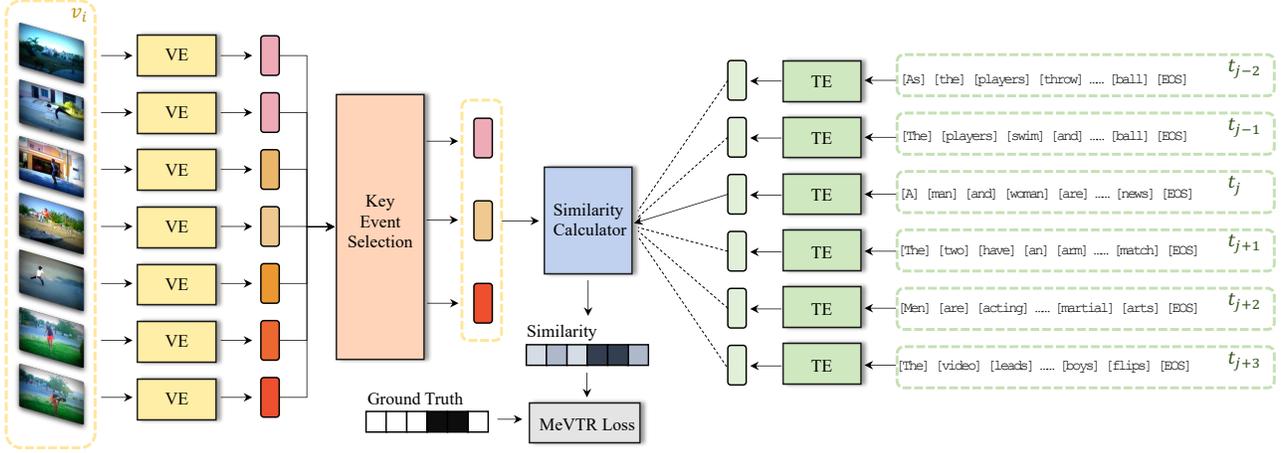


Figure 2: The overall framework of Me-Retriever. The model adopts CLIP [36]’s Visual Encoder(VE) and Text Encoder(TE). After the Visual Encoder, [CLASS] tokens in the last hidden layer are taken as frame embeddings. We use a clustering-based Key Event Selection module to aggregate similar frames and extract key events. Each textual caption is fed into Text Encoder, and [EOS] will be used as text embedding. The similarity between these key events of any video  $v_i$  and any textual caption  $t_j$  is measured in the Similarity Calculator. For each video, there are multiple text correspondences as positive samples.

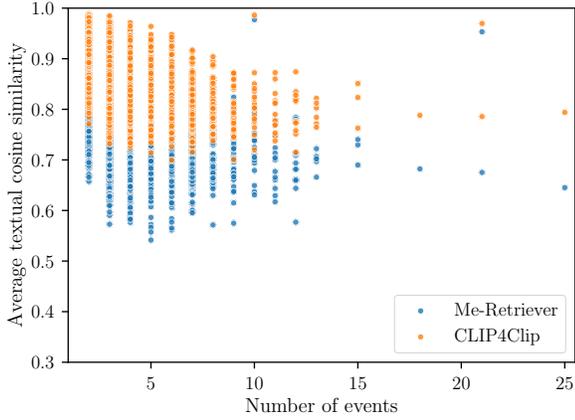


Figure 3: We compare the average cosine similarity between all text pairs of videos with a different number of events. Me-Retriever can generate more diverse text features than CLIP4Clip and refrain from text features collapsing, as we discuss in the main part.

are chosen as initial medoids; the remaining frames are then assigned to their nearest medoids, and within each cluster, new medoids are again selected based on cosine similarity as a metric function. Video  $v_i$  can be represented as a sequence of key event embeddings  $\mathbf{K}_i = \{\mathbf{k}_i^1, \mathbf{k}_i^2, \dots, \mathbf{k}_i^K\}$ . In the loss calculation, we use the pre-computed assignment of all frames as key event video representation.

## 4.2. MeVTR Loss

Following the dual learning paradigm of Text-to-Video and Video-to-Text on VTR tasks, we propose MeVTR loss for MeVTR learning. As previously discussed, there is an issue of textual feature collapse during MeVTR training. To address this, we propose disentangling various positive instances within video samples by excluding non-self positive instances from the softmax function. Moreover, considering the dissimilar granularity between multi-event videos and single-event text, we propose adopting a dynamic weighting strategy. This strategy balances the loss scale of both learning objectives during the training process.

MeVTR loss is defined as a weighted sum of  $\mathcal{L}_{t2v}$  and  $\mathcal{L}_{v2t}$  as follows:

$$\mathcal{L}_{MeVTR} = \mathcal{L}_{v2t} + \alpha \mathcal{L}_{t2v} \quad (2)$$

where  $\mathcal{L}_{v2t}$  and  $\mathcal{L}_{t2v}$  are respectively formulated as:

$$\mathcal{L}_{v2t} = -\frac{1}{|\mathcal{B}|} \sum_i^{|\mathcal{B}|} \sum_k^{|\mathcal{T}_i|} \frac{1}{|\mathcal{T}_i|} \log \frac{\exp \text{sim}(\mathbf{K}_i, \mathbf{w}_k) / \tau}{\sum_{j=1}^{\mathcal{B} \cup \{k\}} \exp \text{sim}(\mathbf{K}_i, \mathbf{w}_j) / \tau}, \quad (3)$$

$$\mathcal{L}_{t2v} = -\frac{1}{|\mathcal{B}|} \sum_j^{\mathcal{B}} \log \frac{\exp \text{sim}(\mathbf{K}_i, \mathbf{w}_j) / \tau}{\sum_{i=1}^{\mathcal{B}} \exp \text{sim}(\mathbf{K}_j, \mathbf{w}_j) / \tau}, \quad (4)$$

Video-to-Text					
Model	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑	k=50 ↑
CLIP4Clip(mean)	39.0	6.60 / 20.74 / -	20.32 / 47.02 / 3.46	29.37 / 61.26 / 6.71	58.41 / 89.77 / 26.24
CLIP4Clip(tight)	49.5	5.23 / 16.45 / -	16.74 / 41.94 / 1.77	25.50 / 57.37 / 4.23	54.22 / 88.00 / 21.52
CenterCLIP	38.0	6.93 / 21.62 / -	20.58 / 47.98 / 3.66	29.90 / 61.72 / 6.96	58.88 / 90.34 / 26.28
FROZEN-in-time	330.0	0.30 / 0.67 / -	0.83 / 2.10 / 0.22	1.60 / 3.37 / 0.67	5.89 / 12.97 / 2.55
CLIPBERT	174.5	1.80 / 6.10 / -	6.48 / 18.97 / 0.18	10.93 / 29.00 / 0.87	31.34 / 61.42 / 8.72
Singularity	64.5	4.83 / 14.68 / -	15.1 / 37.79 / 1.95	22.23 / 49.50 / 4.41	48.96 / 81.94 / 19.38
Ours(avg)	<b>35.0</b>	<b>8.52 / 26.91 / -</b>	<b>23.56 / 54.18 / 3.86</b>	<b>32.95 / 67.13 / 7.79</b>	<b>61.26 / 92.50 / 27.80</b>
Ours(max)	49.0	6.98 / 22.72 / -	18.72 / 46.37 / 2.30	27.04 / 60.62 / 5.76	54.19 / 88.37 / 20.90

Table 2: Results of the Video-to-Text retrieval task on ActivityNet Captions dataset. In each column, we report the Recall@k-Average/One-Hit/All-Hit for each  $k = 1, 5, 10, 50$  respectively. The best result in each column is emphasized in **bold**. We ignore the Recall@1-All-Hit since all models can only achieve a nearly zero result. Ours(avg) maintains its advantage in the text retrieval task and surpasses all other models by a large margin.

Video-to-Text					
Model	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑	k=50 ↑
CLIP4Clip(mean)	441.5	1.32 / 2.92 / -	4.07 / 8.32 / 1.80	5.65 / 11.84 / 2.55	16.84 / 32.01 / 7.65
CLIP4Clip(tight)	391.0	0.80 / 1.80 / -	2.45 / 6.30 / 0.75	4.19 / 10.27 / 1.35	15.26 / 30.43 / 6.37
CenterCLIP	408.0	1.32 / 3.00 / -	3.71 / 7.87 / 1.57	6.33 / 13.42 / 2.62	17.37 / 33.43 / 7.72
FROZEN-in-time	810.0	0.30 / 0.67 / -	0.83 / 2.10 / 0.22	1.60 / 3.37 / 0.67	5.89 / 12.97 / 2.55
ClipBERT	1339.5	0.15 / 0.45 / -	0.60 / 1.50 / 0.15	1.19 / 3.30 / 0.22	5.31 / 11.92 / 1.87
Singularity	621.5	0.95 / 2.17 / -	2.84 / 5.92 / 1.42	4.35 / 9.52 / 1.95	12.22 / 25.94 / 4.87
Ours(avg)	406.0	<b>1.39 / 3.22 / -</b>	<b>5.04 / 9.90 / 2.47</b>	<b>7.02 / 13.72 / 3.37</b>	17.62 / 32.08 / <b>8.70</b>
Ours(max)	<b>365.6</b>	1.01 / 2.62 / -	4.03 / 9.00 / 1.72	6.39 / 13.72 / 2.70	<b>19.08 / 36.66 / 8.10</b>

Table 3: Results of the Video-to-Text retrieval task on Charades-Event dataset. We use the same notation as in Tab. 2. Me-Retriever(avg) and Me-Retriever(max) can achieve best performance.

$sim(\cdot)$  is the similarity function,  $\mathcal{B}$  is the set of samples in the training batch,  $\mathcal{B}^-$  is the set of negative samples in the batch,  $\tau$  is the temperature coefficient, and  $\mathcal{T}_i$  is the set of positive instances of one sample video in  $\mathcal{L}_{t2v}$ . Notably, given a video item  $v_i$ , we discard all positive instances in  $\mathcal{T}_i$  other than the current training sample  $t_k$ . This encourages similarities between a video and its all positive instances to be increased. In  $\mathcal{L}_{t2v}$ , it degenerates to a normal softmax function.

We denote  $\alpha$  as the weighting coefficient that balances  $\mathcal{L}_{t2v}$  and  $\mathcal{L}_{v2t}$ . Given that the two losses exhibit differing scales during training, which owes to the positive/negative sample ratio asymmetry, determining the value of  $\alpha$  requires employing certain hyperparameter-tuning techniques to ensure balanced performance on both tasks. We suggest employing a dynamic weighting strategy as opposed to a fixed weight. A dynamic weighting coefficient  $\alpha$  is computed as the ratio of  $\mathcal{L}_{v2t}$  to  $\mathcal{L}_{t2v}$  during training. In Section 5.3, we conduct experimental studies to compare the impact of various weighting strategies.

**Similarity function** For video-text pairs in bijective correspondences, the similarity between video and text is measured by cosine similarity in the feature space. In MeVTR, we use two approaches are used to measure the similarity between a video  $v_i$  consisting of key events  $\{\mathbf{k}_i^1, \mathbf{k}_i^2, \dots, \mathbf{k}_i^K\}$  and a text  $t_j$ : (1) we average the cosine similarity scores of all key event embeddings and the text as in Eq. 5; (2) we select the maximum cosine similarity of any key event-text pair to represent the similarity between the video to which these medoids belong as in Eq. 6;

$$s_{avg}(v_i, t_j) = \text{avg}(\cos(\mathbf{k}_i^1, \mathbf{w}_j), \dots, \cos(\mathbf{k}_i^K, \mathbf{w}_j)) \quad (5)$$

$$s_{max}(v_i, t_j) = \max(\cos(\mathbf{k}_i^1, \mathbf{w}_j), \dots, \cos(\mathbf{k}_i^K, \mathbf{w}_j)) \quad (6)$$

We denote Me-Retriever with different similarity functions as Me-Retriever(avg) and Me-Retriever(max).

Text-to-Video					
Model	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑	k=50 ↑
CLIP4Clip(mean)	<b>11.0</b>	15.25	36.43	49.11	76.64
CLIP4Clip(tight)	15.0	10.14	29.07	41.93	72.50
CenterCLIP	<b>11.0</b>	<u>15.53</u>	<b>37.21</b>	<b>49.78</b>	<b>77.19</b>
FROZEN-in-time	111.0	1.82	6.80	11.76	33.86
ClipBERT	48.0	4.20	14.57	23.35	51.02
Singularity(FT)	22.0	9.29	24.96	36.26	65.55
Ours(avg)	<b>11.0</b>	<b>15.99</b>	<u>37.13</u>	<u>49.60</u>	<u>76.66</u>
Ours(max)	13.0	14.93	34.44	46.36	74.65

Table 4: Results of the Text-to-Video retrieval task on ActivityNet Captions dataset. We compare CLIP4Clip [30], CenterCLIP [54], FROZEN-in-time [3], ClipBERT [22], and [21]. The best result in each column is emphasized in **bold**, and the second best result is underlined. Our Me-Retriever(avg) achieves comparably good performance.

## 5. Experiments

### 5.1. Experiment Details

**Datasets** We perform experiments on two datasets containing multi-event videos: ActivityNet Captions and Charades-Event. It’s worth noting that datasets such as MSR-VTT [48], even though videos are accompanied by multiple similar captions, do not align with our multi-event scenario.

**ActivityNet Captions** [20] is a large-scale dataset with 20,000 video-text pairs extracted by web crawling on YouTube based on ActivityNet [4] videos. It covers a wide range of daily scenarios. The training set contains 10,009 videos and 4,917 videos/17,505 texts in the publicly accessible evaluation set ‘val1’. Each video has multiple textual descriptions with time-interval annotation. This helps us construct a multi-event scenario for retrieval by taking all textual captions of a video as its associated events.

**Charades-Event** is a video dataset of indoor activities with extensive action classes that we adopt from [37]. To create a MeVTR dataset, we filter out all videos with text annotation provided by [15] and interpret each text as an *event* in our MeVTR setting. There are 5,338 videos/12,408 texts in the training set and 1,334 videos/3,720 texts in the test set.

**Experiment Settings** In the key event selection module, we use 60 as the maximum number of iterations and  $10^{-5}$  as the distance threshold for clustering. Our model chooses 16 as a fixed number of key events and does not use a smaller or larger number of key events because shorter sequences cannot cover the number of events in part of the videos (the maximum event number in ActivityNet Captions is 27), and longer sequences contradict our intention of clustering. For each experiment, we set the epoch number as 5.

Text-to-Video					
Model	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑	k=50 ↑
CLIP4Clip(mean)	136.0	1.53	6.48	10.54	29.38
CLIP4Clip(tight)	127.0	1.45	5.54	9.35	30.11
CenterCLIP	131.0	<u>1.93</u>	6.61	<u>10.67</u>	<u>30.13</u>
FRONZEN-in-Time	277.0	0.35	1.99	3.82	14.52
ClipBERT	442.0	0.35	1.77	3.31	11.10
Singularity	179.5	1.53	5.22	8.52	23.87
Ours(avg)	<u>126.0</u>	1.88	<u>6.69</u>	10.59	29.97
Ours(max)	<b>112.0</b>	<b>2.02</b>	<b>7.61</b>	<b>12.74</b>	<b>34.49</b>

Table 5: Results of Text-to-Video retrieval task on Charades-Event dataset. We use the same notation as Tab. 2. Me-Retriever(max) and Me-Retriever(avg) can achieve the best performance compared to other models.

### 5.2. Results

The experiments show that Me-Retriever performs competitively within the context of MeVTR. On the Video-to-Text task, as presented in Tables 2 and 3, Me-Retriever(avg) surpasses its counterparts. This result underscores our model’s capability to address both tasks across different datasets effectively. Additionally, we observe that CLIP4Clip also attains a commendable Recall@1-Average performance on the Video-to-Text task compared to other baseline models. This observation sheds light on the robust transferability of the pre-trained encoder to videos. This indicates that the primary challenge in MeVTR often resides within representation alignment.

Our model also demonstrates comparable results on both datasets on the Text-to-Video task, as shown in Tab. 4 and Tab. 5. Me-Retriever(avg) demonstrates comparable performance to baseline models like CenterCLIP across numerous metrics on both datasets. Interestingly, we also notice that Me-Retriever(max) performs best on the Charades-Event dataset but not on the ActivityNet Captions dataset. This disparity can be attributed to overfitting, warranting further investigation. The noticeable performance of certain baseline models aligns with our conjecture that prior models can be effectively adapted to the Text-to-Video retrieval task. However, this adaptation might come at the expense of suboptimal results in the Video-to-Text task, as evident in Table 2. This phenomenon can be attributed to the embeddings of different texts collapsing to become more proximate, leading to an unintentional competition among positive instances that ultimately dampens their significance.

**Performance on videos with different subsets** In order to reveal the model performance on videos of different properties like video duration and the number of events, we partition the complete ActivityNet Caption test set into two groups of subsets. These subsets are partitioned based on the video duration and the count of events within the videos, respectively:

Text to Video					
Model	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑	k=50 ↑
Ours (Best)	<b>11.0</b>	<b>15.17</b>	<b>35.94</b>	<b>48.39</b>	<b>75.85</b>
w/o Key Event	12.0	15.03	35.67	47.94	75.06
w/o MeVTR loss	12.0	14.45	34.90	47.25	75.24

Table 6: Ablation studies on key modules of Me-Retriever training on the Text-to-Video task for ActivityNet Captions. Only the Recall@ $k$ -Average is reported in the main part. We find that model performance drops with removing any modules, proving its efficacy.

Video to Text					
Model	Median Rank ↓	k=1 ↑	k=5 ↑	k=10 ↑	k=50 ↑
Ours (Best)	<b>35.0</b>	<b>8.52</b>	<b>23.56</b>	<b>32.79</b>	<b>61.26</b>
w/o Key Event	41.5	7.40	21.00	30.27	57.71
w/o MeVTR loss	41.0	6.70	20.05	28.96	57.45

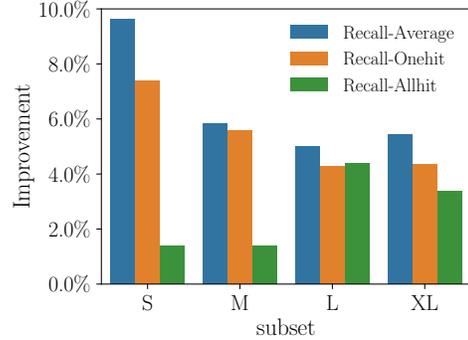
Table 7: Ablation studies on key modules of Me-Retriever training on the Video-to-Text task for ActivityNet Captions. Model performance declines to different extents without our modules. We conclude that our model can benefit Video-to-Text retrieval prominently.

1. Video duration: test-**S** (Short) with videos shorter than 1 minute, test-**M** (Medium) with videos between 1-2 minutes, test-**L** (Long video) with videos between 2-3 minutes, and test-**XL** (eXtra Long video) with even longer videos than 3 minutes;
2. Number of events: test-**E1** with videos containing fewer than 4 events, test-**E2** with 5-12 events, and test-**E3** with more than 12 events.

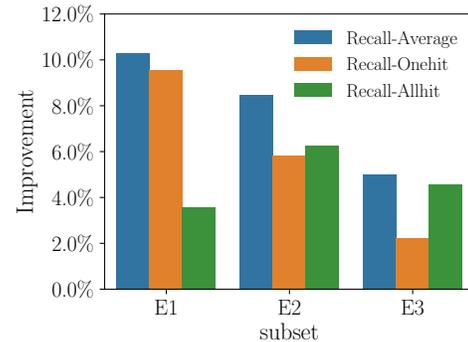
Comparing performance across datasets of different sizes isn't equitable for a retrieval task, as larger candidate sets exacerbate retrieval difficulties. Therefore, we gauge the average improvement in the model performance of Me-Retriever(avg) to the baseline CLIP4Clip(mean) and identify the subsets in which our model showcases notable enhancements. From Fig. 4, we also see that our model usually performs better on test-**E1** and test-**E2**, and improvement slightly decreases for even longer videos. We assume this problem is inevitable since we use a fixed number of clusters and videos containing many events in test-**E3**, and larger numbers of events increase the retrieval difficulties. Whereas for videos of different duration, model performance slightly varies.

### 5.3. Ablation Studies

In this section, we conduct various experiments to investigate the effectiveness of different weighting strategies in MeVTR loss and ablate the proposed components of our architecture.



(a) On subsets of different duration



(b) On subsets of different event number

Figure 4: We compare average performance improvement in percentage(%) for Video-to-Text task on subsets of ActivityNet Captions. It shows how much percentage Me-Retriever(avg) is better than CLIP4Clip(mean) on different subsets. Fig. 4a: Video-to-Text results for test-**S/M/L/XL**; Fig. 4b: Video-to-Text results for test-**E1/E2/E3**.

**Impact of weighting strategy in MeVTR Loss** To study the effect of different weighting strategies of MeVTR loss, i.e., dynamic weighting and fixed weighting, we test our model trained with a set of different weighting coefficients  $\alpha$  as illustrated in Fig. 5 and Fig. 6. On the Video-to-Text task and Text-to-Video Retrieval tasks, we notice that dynamic weighting  $\alpha$  can guarantee us a comparable performance in general and less sensitivity than any fixed coefficient. It is also illustrated that when we choose a smaller coefficient, the Video-to-Text retrieval is favored, while a bigger coefficient benefits the other way around. A cutoff point is empirically between 1.0 and 2.0. The reason is that  $\mathcal{L}_{v2t}$  loss always has a comparably larger scale than  $\mathcal{L}_{t2v}$  loss and dominates training due to the unbalanced number of video-text instances, as we have discussed in Section 4.2. This ensures us that a dynamic weighting strategy can improve the Text-to-Video task by amplifying the influence of  $\mathcal{L}_{t2v}$  and guarantee stable improvements on both tasks.

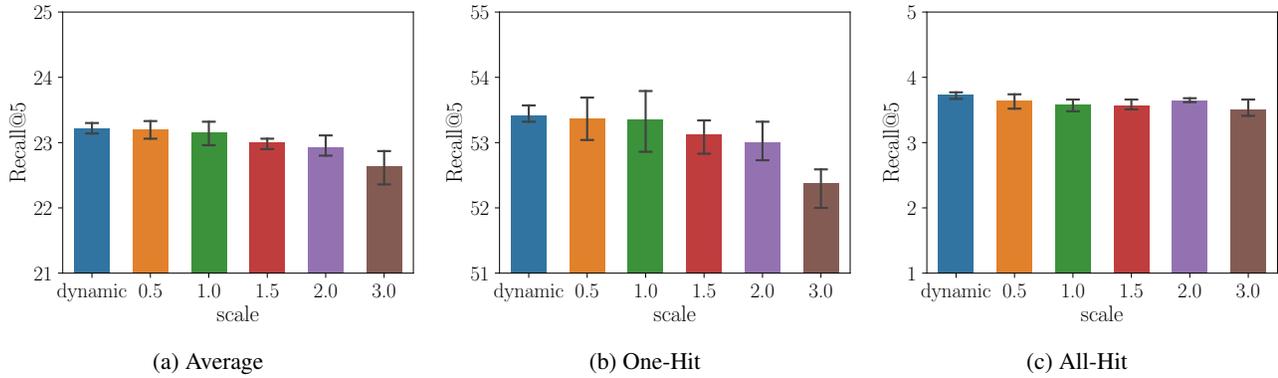


Figure 5: We compare the model performance with different weighting strategies, a dynamic weight  $\alpha$  and different choices of fixed weights from a choice of  $\{0.5, 1.0, 1.5, 2.0, 3.0\}$ , in the MeVTR loss on Recall@5 on the Video-to-Text task for ActivityNet Captions. Fig. 5a-5c shows the results of Recall@5-Average/One-Hit/All-Hit respectively. We can find that compared to a fixed weighting coefficient, a dynamic weight  $\alpha$  guarantees more stable and good results in different metrics.

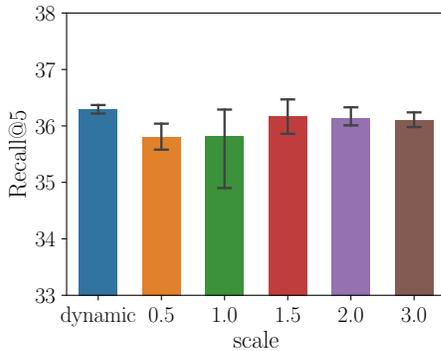


Figure 6: We compare the model performance with different weighting coefficients in the MeVTR loss on Recall@5 on the Text-to-Video task for ActivityNet Captions. The same graphical representations are used as Fig. 5. The dynamic weighting coefficient performs better than any other fixed coefficient considering performance and variance.

**Effects of different modules** We choose our best model Me-Retriever(avg) and ablate the key event selection module by removing the key event selection module and the MeVTR loss by replacing it with the cross-entropy loss, respectively. Results are shown in Tab. 6 and Tab. 7. We display the Recall@ $k$ -Average on both tasks in the table and observe that the model performances drop after removing any modules. In particular, our method demonstrates a larger improvement by 1.12% without the key event selection module and 1.82% without the MeVTR loss on the Video-to-Text task. This agrees with our assumption that our Me-Retriever model can deal with textual feature collapse in the MeVTR training. On the Text-to-Video task, both modules also show benefits to the performance on all metrics.

**Textual features of multi-event videos** To show whether Me-Retriever can prevent heterogeneous text features of a multi-event video from collapsing, we calculate the average cosine similarity. If the text features of one video are too similar, then the average cosine similarity among texts is expected to be closer to 1. We find that in Me-Retriever(avg), the cosine similarity averaged on the whole test set is 0.789 with a variance of 0.0059 while in CLIP4Clip(mean), the value is 0.864 with a variance of 0.0028. And as in Fig. 3, Me-Retriever has lower average cosine similarity than CLIP4Clip. This shows that text features of the same video generated by Me-Retriever are more diverse than those by CLIP4Clip, which agrees with our motivation.

## 6. Limitations

Constrained by time and labor expenses, our experimental studies on the MeVTR task are based on the pre-existing video-text datasets including ActivityNet Captions and Charades. Nonetheless, it is important to recognize the potential bias in these datasets due to their emphasis on human activities and limitations in representing a wider range of videos. We aim to create a more comprehensive MeVTR benchmark encompassing more general videos.

## 7. Conclusion

We study a practical Video-Text retrieval scenario, Multi-event Video-Text Retrieval, in this paper. We find previous models suffer notable performance degradation in inference on MeVTR and textual feature collapse when re-trained on MeVTR. We present a new CLIP-based model, Me-Retriever, to conquer this challenge. In our experimental studies, we demonstrate the efficacy of these techniques on both Text-to-Video and Video-to-Text tasks for ActivityNet Captions and Charades-Event datasets. We believe this can be a strong baseline for future studies.

## Acknowledgements

This work has been funded by the German Federal Ministry of Education and Research and the Bavarian State Ministry for Science and the Arts. The authors of this work take full responsibility for its content.

## References

- [1] Partitioning Around Medoids (Program PAM). In Leonard Kaufman and Peter J. Rousseeuw, editors, *Wiley Series in Probability and Statistics*, pages 68–125. John Wiley & Sons, Inc., Hoboken, NJ, USA, Mar. 1990. 3
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6836–6846, 2021. 2
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1728–1738, 2021. 1, 3, 6
- [4] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. 2, 6
- [5] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer, 2020. 2
- [6] Zhengsu Chen, Lingxi Xie, Jianwei Niu, Xuefeng Liu, Longhui Wei, and Qi Tian. Visformer: The vision-friendly transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 589–598, 2021. 2
- [7] Xing Cheng, Hezheng Lin, Xiangyu Wu, Fan Yang, and Dong Shen. Improving video-text retrieval by multi-stream corpus alignment and dual softmax loss. *arXiv preprint arXiv:2109.04290*, 2021. 1
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [9] Jianfeng Dong, Xirong Li, and Cees GM Snoek. Predicting visual features from text for image and video caption retrieval. *IEEE Transactions on Multimedia*, 20(12):3377–3388, 2018. 1, 2
- [10] Jianfeng Dong, Xirong Li, Chaoxi Xu, Xun Yang, Gang Yang, Xun Wang, and Meng Wang. Dual encoding for video retrieval by text. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):4065–4080, 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
- [12] Han Fang, Pengfei Xiong, Luhui Xu, and Yu Chen. Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097*, 2021. 1, 3
- [13] Tsu-Jui Fu, Linjie Li, Zhe Gan, Kevin Lin, William Yang Wang, Lijuan Wang, and Zicheng Liu. Violet: End-to-end video-language transformers with masked visual-token modeling. *arXiv preprint arXiv:2111.12681*, 2021. 2
- [14] Valentin Gabeur, Chen Sun, Karteek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IV 16*, pages 214–229. Springer, 2020. 1
- [15] 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. 6
- [16] Zijian Gao, Jingyu Liu, Sheng Chen, Dedan Chang, Hao Zhang, and Jinwei Yuan. Clip2tv: An empirical study on transformer-based methods for video-text retrieval. *arXiv preprint arXiv:2111.05610*, 1(2):6, 2021. 1, 2, 3
- [17] Jindong Gu, Zhen Han, Shuo Chen, Ahmad Beirami, Bailan He, Gengyuan Zhang, Ruotong Liao, Yao Qin, Volker Tresp, and Philip Torr. A systematic survey of prompt engineering on vision-language foundation models. *arXiv preprint arXiv:2307.12980*, 2023. 1
- [18] Yukun Huang, Yongcai Guo, and Chao Gao. Efficient parallel inflated 3d convolution architecture for action recognition. *IEEE Access*, 8:45753–45765, 2020. 2
- [19] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting visual-language models for efficient video understanding. In *European Conference on Computer Vision*, pages 105–124. Springer, 2022. 2
- [20] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pages 706–715, 2017. 6
- [21] Jie Lei, Tamara L Berg, and Mohit Bansal. Revealing single frame bias for video-and-language learning. *arXiv preprint arXiv:2206.03428*, 2022. 6
- [22] Jie Lei, Linjie Li, Luowei Zhou, Zhe Gan, Tamara L Berg, Mohit Bansal, and Jingjing Liu. Less is more: Clipbert for video-and-language learning via sparse sampling. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7331–7341, 2021. 2, 6
- [23] Dongxu Li, Junnan Li, Hongdong Li, Juan Carlos Niebles, and Steven CH Hoi. Align and prompt: Video-and-language pre-training with entity prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4953–4963, 2022. 1, 2
- [24] Gen Li, Nan Duan, Yuejian Fang, Ming Gong, and Daxin Jiang. Unicoder-vl: A universal encoder for vision and language by cross-modal pre-training. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 11336–11344, 2020. 2
- [25] Hengduo Li, Zuxuan Wu, Abhinav Shrivastava, and Larry S Davis. 2d or not 2d? adaptive 3d convolution selection for

- efficient video recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6155–6164, 2021. 2
- [26] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*, 2019. 2
- [27] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 2
- [28] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3202–3211, 2022. 2
- [29] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32, 2019. 2
- [30] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022. 1, 3, 6
- [31] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9879–9889, 2020. 1, 2
- [32] Daniel Neimark, Omri Bar, Maya Zohar, and Dotan Asseilmann. Video transformer network. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3163–3172, 2021. 2
- [33] Mayu Otani, Yuta Nakashima, Esa Rahtu, Janne Heikkilä, and Naokazu Yokoya. Learning joint representations of videos and sentences with web image search. In *Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part I 14*, pages 651–667. Springer, 2016. 1
- [34] Mandela Patrick, Po-Yao Huang, Yuki Asano, Florian Metze, Alexander Hauptmann, Joao Henriques, and Andrea Vedaldi. Support-set bottlenecks for video-text representation learning. *arXiv preprint arXiv:2010.02824*, 2020. 2
- [35] Jesús Andrés Portillo-Quintero, José Carlos Ortiz-Bayliss, and Hugo Terashima-Marín. A straightforward framework for video retrieval using clip. In *Mexican Conference on Pattern Recognition*, pages 3–12. Springer, 2021. 3
- [36] 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. 1, 2, 3, 4
- [37] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 510–526. Springer, 2016. 6
- [38] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *arXiv preprint arXiv:1908.08530*, 2019. 2
- [39] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019. 2
- [40] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7464–7473, 2019. 2
- [41] 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
- [42] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 2
- [43] Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. Multimodal few-shot learning with frozen language models. *Advances in Neural Information Processing Systems*, 34:200–212, 2021. 2
- [44] 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. 2
- [45] Jinpeng Wang, Yixiao Ge, Rui Yan, Yuying Ge, Kevin Qinghong Lin, Satoshi Tsutsui, Xudong Lin, Guanyu Cai, Jianping Wu, Ying Shan, et al. All in one: Exploring unified video-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6598–6608, 2023. 2
- [46] Yi Wang, Kunchang Li, Yizhuo Li, Yanan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, et al. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. 2
- [47] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 22–31, 2021. 2
- [48] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5288–5296, 2016. 6
- [49] Ran Xu, Caiming Xiong, Wei Chen, and Jason Corso. Jointly modeling deep video and compositional text to bridge vision

- and language in a unified framework. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29, 2015. 1, 2
- [50] Hongwei Xue, Yuchong Sun, Bei Liu, Jianlong Fu, Ruihua Song, Houqiang Li, and Jiebo Luo. Clip-vip: Adapting pre-trained image-text model to video-language representation alignment. *arXiv preprint arXiv:2209.06430*, 2022. 3
- [51] Jianwei Yang, Yonatan Bisk, and Jianfeng Gao. Taco: Token-aware cascade contrastive learning for video-text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11562–11572, 2021. 1, 2, 3
- [52] Xinyi Ying, Longguang Wang, Yingqian Wang, Weidong Sheng, Wei An, and Yulan Guo. Deformable 3d convolution for video super-resolution. *IEEE Signal Processing Letters*, 27:1500–1504, 2020. 2
- [53] Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5579–5588, 2021. 2
- [54] Shuai Zhao, Linchao Zhu, Xiaohan Wang, and Yi Yang. Centerclip: Token clustering for efficient text-video retrieval. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 970–981, 2022. 6
- [55] Jie Zheng, Andrea GP Schjetnan, Mar Yebra, Bernard A Gomes, Clayton P Mosher, Suneil K Kalia, Taufik A Valiante, Adam N Mamelak, Gabriel Kreiman, and Ueli Rutishauser. Neurons detect cognitive boundaries to structure episodic memories in humans. *Nature neuroscience*, 25(3):358–368, 2022. 3
- [56] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021. 2
- [57] Luwei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13041–13049, 2020. 2