

# Tagging before Alignment: Integrating Multi-Modal Tags for Video-Text Retrieval

Yizhen Chen<sup>1</sup>, Jie Wang<sup>1</sup>, Lijian Lin<sup>2</sup>, Zhongang Qi<sup>2\*</sup>, Jin Ma<sup>1</sup>, Ying Shan<sup>2</sup>

<sup>1</sup>IPS Search, Tencent PCG

<sup>2</sup>ARC Lab, Tencent PCG

{grayyzchen, zhongangqi, yingsshan}@tencent.com, photonicsjay@163.com, ljlin@stu.xmu.edu.cn, majin01@mail.ustc.edu.cn

## Abstract

Vision-language alignment learning for video-text retrieval arouses a lot of attention in recent years. Most of the existing methods either transfer the knowledge of image-text pretraining model to video-text retrieval task without fully exploring the multi-modal information of videos, or simply fuse multi-modal features in a brute force manner without explicit guidance. In this paper, we integrate multi-modal information in an explicit manner by tagging, and use the tags as the anchors for better video-text alignment. Various pretrained experts are utilized for extracting the information of multiple modalities, including object, person, motion, audio, etc. To take full advantage of these information, we propose the TABLE (TAGging Before aLignmEnt) network, which consists of a visual encoder, a tag encoder, a text encoder, and a tag-guiding cross-modal encoder for jointly encoding multi-frame visual features and multi-modal tags information. Furthermore, to strengthen the interaction between video and text, we build a joint cross-modal encoder with the triplet input of [vision, tag, text] and perform two additional supervised tasks, Video Text Matching (VTM) and Masked Language Modeling (MLM). Extensive experimental results demonstrate that the TABLE model is capable of achieving State-Of-The-Art (SOTA) performance on various video-text retrieval benchmarks, including MSR-VTT, MSVD, LSMDC and DiDeMo.

## Introduction

Vision-language alignment learning for video-text retrieval becomes an emerging requirement with the increasing of videos and short videos uploaded online, and has attracted great attention in recent years. In the research field, it returns the most relevant videos for a given text query to facilitate large-scale videos searching and management. In the recommendation field, it recommends the most relevant text queries for the video that the user is watching to prompt more searching and browsing. Although a lot of recent works (Zhu and Yang 2020; Gabeur et al. 2020; Dzabraev et al. 2021; Lei et al. 2021; Liu et al. 2021; Luo et al. 2021; Fang et al. 2021; Cheng et al. 2021; Chen, Liu, and Albanie 2021; Cao et al. 2022) have made remarkable progress, the cross-modal alignment between video and text remains a challenging task.

\*Corresponding author

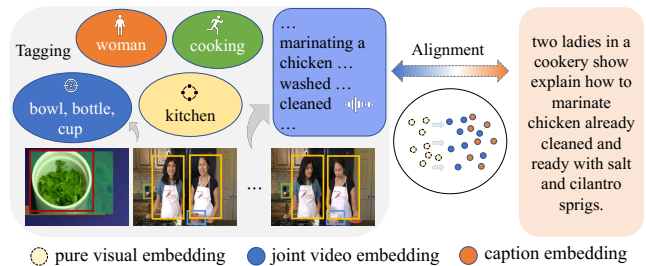


Figure 1: The multi-modal information is able to motivate the video-text alignment in an explicit manner, i.e., tagging.

Most of the existing video-text retrieval models use multi-layer transformers to learn generic representations from massive video-text pairs, which can be roughly divided into two categories. The first category uses only frame features to transfer the knowledge of image-text pretrained model to video-text retrieval task without fully exploring the multi-modal information of videos (Lei et al. 2021; Luo et al. 2021; Fang et al. 2021; Cheng et al. 2021; Cao et al. 2022). A representative method is CLIP4Clip (Luo et al. 2021), which utilizes the knowledge of the CLIP (Contrastive Language-Image Pretraining) (Radford et al. 2021) model to visually encode multi-frame information. The model achieves good performance on multiple benchmarks, demonstrating the effectiveness of the knowledge transferring. The second category is to jointly encode the information of various modalities in the video, such as objects, actions, scenes, audio, etc. (Gabeur et al. 2020; Dzabraev et al. 2021; Wang, Zhu, and Yang 2021; Hao et al. 2021). A typical representative method is MMT (Multi-Modal Transformer) (Gabeur et al. 2020), which uses multiple pretrained experts to obtain embeddings of different modalities, and builds a multi-modal transformer for feature fusion. The performance of this model demonstrates the advantage of using multi-modal information over using single visual information of the video. However, most of these methods simply fuse multi-modal features in a brute force manner without explicit guidance, which increases the difficulty of the learning process.

In this paper, we introduce a novel approach which not only transfers image-text knowledge, but also fully exploits multi-modal information in an explicit manner. A set of pre-

trained experts are utilized to extract features from multiple modalities, including object, person, scene, motion, and audio. Object and person focus on instance features or local features of the video; scene focuses on background or global features. The motion features reveal the behavior or action happening in the video, which extract temporal information from consecutive frames. The audio provides supplementary descriptions, which may not appear in visual features. The most challenging part of the proposed approach is how to fuse the features of different modalities effectively and efficiently. In this work, we utilize tagging as bridge, whose advantages lie in the following two aspects. First, embeddings from various modalities generated by different experts are not compatible from each other; while tagging generates unified representations for different modalities. Second, multi-modal tags can be treated as anchor points to make the visual-language alignment learn the local feature, the global feature, the temporal feature, and other supplementary features comprehensively, which eases the learning process significantly. For example, in Fig. 1, by using multiple pre-trained experts, we can extract object, person, motion, scene and audio tags from video. These tags provide abundant information from multi-modality of the video, which bridge the gap between vision and language, making the alignment learning more efficiently and precisely.

Specifically, we construct a TABLE (TAGging Before aLignmEnt) network, where a cross-modal encoder is constructed to jointly learn multi-frame visual features and multi-modal tags features of the video. It uses multi-modal tags as anchors to guide the visual-text alignment, and learns the temporal information between frames. Further, we introduce two auxiliary tasks, namely Video Text Matching (VTM) and Masked Language Modeling (MLM). The goal of VTM is to judge whether the [visual, tag] matches the [text], and the objective of MLM is to recover the masked word according to [vision, tag, unmasked text].

In summary, the contributions of this work lie in four-fold:

- We propose a novel method for video-text retrieval, which not only transfers image-text knowledge, but also fully exploits the multi-modal information of video, including object, person, scene, motion and audio.
- We integrate multi-modal information by an explicit and unified method, tagging, and use them as anchors for guiding visual-text alignment.
- We build a TABLE network to jointly encode multi-frame visual features and multi-modal tag information. And we introduce VTM and MLM as auxiliary supervisions to strengthen video-text interaction.
- TABLE achieves SOTA performance on several benchmarks, i.e., MSR-VTT, MSVD, LSMDC and DiDeMo.

## Related Work

### Video Representation Learning

A lot of approaches have been proposed for video representation learning, which can be classified into convolution-based methods and transformer-based methods. Many previous works (Tran et al. 2015; Xie et al. 2018; Feichten-

hofer et al. 2019) employed 2D or 3D convolutional network to encode spatial and temporal information of the videos. Recently, the Vision Transformer (ViT) was proposed and attained excellent results on image classification compared to convolutional methods. Since then, many recent works applied ViT to encode the visual features of videos, such as ViViT (Arnab et al. 2021) and TimeSformer (Bertasius, Wang, and Torresani 2021), which generally decoupled the spatial and temporal information by designing two-stage transformer structures. Most of the above approaches exploited temporal information by improving image representation, however, the abundant multi-modal information of videos were not fully utilized. Our work performs temporal encoding with cross-modal interaction and exploits comprehensive information from multiple modalities, which are extracted by various pretrained experts.

### Video-Text Retrieval

Recent video-text retrieval approaches can be divided into two types according to whether the multi-modal information of videos is utilized or not. The first type simply adopted the visual features for video-text alignment, where CLIP-based methods (Luo et al. 2021; Fang et al. 2021; Cheng et al. 2021; Cao et al. 2022) show obvious advantages in recent years. CLIP4Clip (Luo et al. 2021) firstly transferred the knowledge of large-scale image-text pretraining to the task of video-text retrieval with fine-tuning. Based on it, CLIP2VIDEO (Fang et al. 2021) proposed a temporal different block to capture video’s motion feature, and a temporal alignment block to re-align the tokens of video clips and phrases. The second type used multi-modal information contained in videos for enhancing video-text alignment (Gabeur et al. 2020; Dzabraev et al. 2021; Wang, Zhu, and Yang 2021; Hao et al. 2021). MMT (Gabeur et al. 2020) exploited multi-modal information extracted by seven pre-trained experts but only fused them in a brute force manner without explicit guidance. Wang et al. (Wang, Zhu, and Yang 2021) and Hao et al. (Hao et al. 2021) performed local alignment between multi-modal features and text features, but the gap between the embeddings of different modalities is too huge to bridge. To solve this problem, we propose to integrate multi-modal information by an explicit method, i.e., tagging. Tagging generates unified representations to erase the gap between different modalities, and can be treated as anchors to guide the visual-language alignment more explicitly. Some image-text retrieval methods (Zhen et al. 2019; Qian et al. 2021) also used tag to boost performance, but our work is different in many ways. Firstly, these work only used the unimodal tag of image, while our work fully extracts the multi-modal tags of video, which is more comprehensive. Secondly, these work used classification loss as extra supervision, while our method does not directly use tag as training target, but regards it as anchor for video-text alignment.

### Methodology

Given a set of videos and texts, our goal is to learn accurate representations and calculate their similarities. The video (or text) candidates are then ranked by their similarities with the

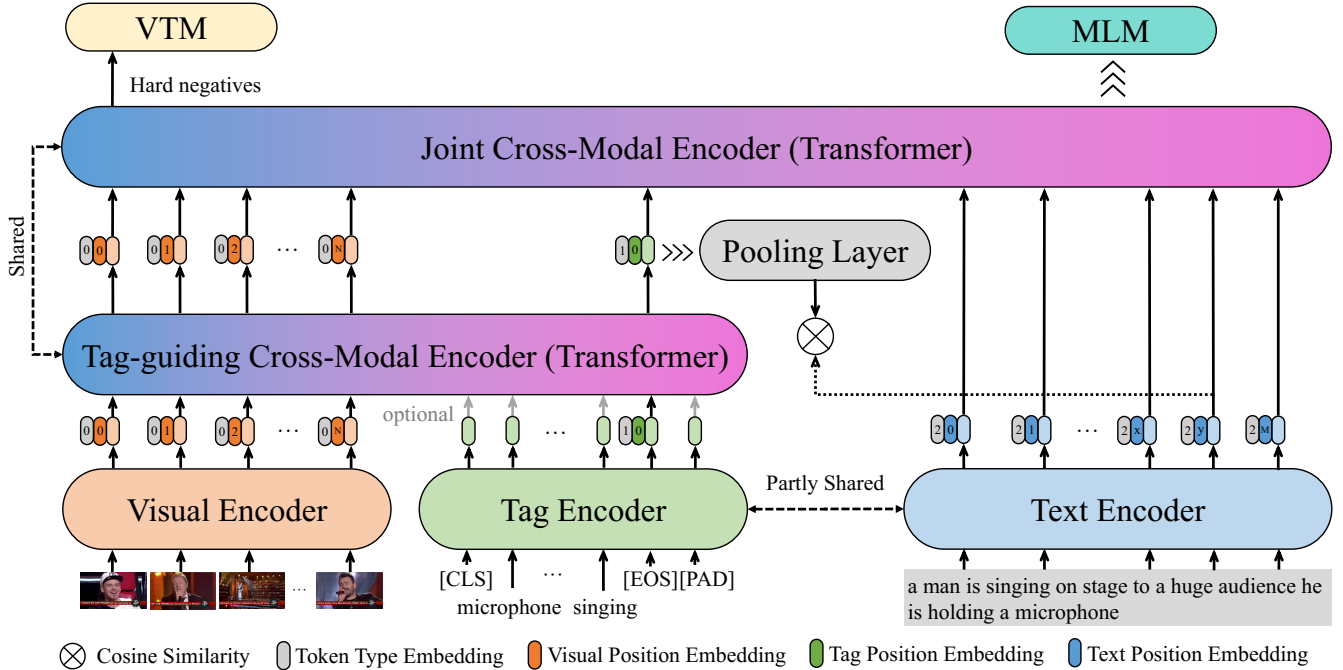


Figure 2: Overall framework of our proposed TABLE (TAGging Before aLignmEnt) model.

query text (or video) in the field of text-to-video (or video-to-text) retrieval. To achieve this goal, video and text need to be aligned in a joint embedding space, which is difficult due to the independence of the two feature extraction processes. In this paper, we propose to fully employ the multi-modal tags of the video by using them as anchors to motivate the visual-text alignment. Specifically, we construct a TABLE (TAGging Before aLignmEnt) network to jointly encode multi-frame visual information and multi-modal tag information. As shown in Fig. 2, our model consists of a visual encoder, a tag encoder, a text encoder, a tag-guiding (TG) cross-modal encoder and a joint cross-modal encoder. The two cross-modal encoders share parameters with each other and performs multi-modal information integration.

### Multi-Modal Tag Mining

In this paper, multiple pretrained experts are used to extract related tag information of individual modality. Firstly, we adopt the yolov5s model pretrained on the COCO (Lin et al. 2014) dataset for video object detection. Main objects appearing in multiple frames with high confidence score are selected with object tags  $t_{obj}$ . In particular, we choose the detection model pretrained on Open Image dataset (Kuznetsova et al. 2020) to distinguish person into man, woman, boy and girl, obtaining specific person tags  $t_{per}$ . Secondly, we employ the COCO pretrained ViT-B-16 (Dosovitskiy et al. 2020) model to perform image classification of video frames, where predicted categories with high confidence score are regarded as the scene tags  $t_{sce}$ . Object and person tags focus on instance or local features, while scene tags put attention to global information of video

frames. Thirdly, we adopt the S3D (Xie et al. 2018) network pretrained on Kinetics-400 dataset (Carreira and Zisserman 2017) to obtain the motion tags  $t_{mot}$ , which is important for video understanding. Furthermore, we apply the Automatic Speech Recognition (ASR) API of iFLYTEK to get video transcripts and then employ the KeyBERT (Grootendorst 2020) to extract transcript keywords as audio tags  $t_{aud}$ . The audio modality usually provides complementary information which is always not included in visual features. At last, we concatenate the above tags and obtain the whole multi-modal tag information of the video,  $t_{mul} = [t_{obj}, t_{per}, t_{sce}, t_{mot}, t_{aud}]$ .

### Visual and Text Encoder

**Visual Encoder.** We uniformly sample  $N$  frames to form a sequence as the video representation,  $f_i = \{f_i^1, f_i^2, \dots, f_i^N\}$ . Inspired by previous image-text pretraining work, we adopt the ViT model in CLIP (Radford et al. 2021) to extract visual features, which processes non-overlapping image patches of individual frame and linearly project them to 1D token sequence. The patch tokens are then passed through a 12-layer transformer to realize self-attention procedure. The [CLS] embedding is projected into a normalized lower-dimensional embedding space to obtain the overall representation of each frame. In Fig. 2, the output of the visual encoder represents sequential representations of multiple frames,  $v_i = \{v_i^1, v_i^2, \dots, v_i^N\}$ .

**Text Encoder.** We apply the BERT-base encoder in CLIP (Radford et al. 2021) to obtain tag and caption embedding, which is a 12-layer transformer with 512 dimensional width and 8 attention heads. The transformer output of tag

and caption can be expressed as  $t_i = \{t_i^1, t_i^2, \dots, t_i^K\}$  and  $c_i = \{c_i^1, c_i^2, \dots, c_i^M\}$  respectively, where  $K$  and  $M$  represents the token length of tag and caption. The [EOS] embedding of the last layer is passed through a linear projection layer and regarded as the overall representation of the text (e.g., tag and caption). Particularly, the parameters of transformer blocks are shared between tag and caption encoders, except the linear projection layer.

### Tag-Guiding Cross-Modal Encoder

To bridge the visual-text semantic gap for alignment, we introduce multi-modal tags as anchors. As shown in Fig. 2, we concatenate multi-frame visual features and multi-modal tag embeddings as the input of the cross-modal encoder. Specifically, the concatenated input can be represented by  $\{v_i^1, v_i^2, \dots, v_i^N; t_i^e\}$ , where  $t_i^e$  represents the [EOS] embedding of the tag encoder output and the overall information of the multi-modal tags. The **TG** cross-modal encoder is composed of 4-layer transformer with 512 dimensional width and 8 cross-attention heads. The position embedding and transformer parameters are initialized by the weights of first four layers of the CLIP’s text encoder. The **TG** cross-modal encoder conducts deep fusion of multi-frame visual features with multi-modal tag embeddings, meanwhile modelling the temporal information of the video. The fused output can be expressed as  $\{g_i^1, g_i^2, \dots, g_i^N; g_i^e\}$ . We then apply an average pooling layer to acquire the overall representation, which can be formulated as  $g_i^o = \rho(\{g_i^1, g_i^2, \dots, g_i^N; g_i^e\})$ . In addition, to fully exploit the pretrained image-text knowledge, we adopt residual connection between the pooled visual features and the overall cross-modal representation, which can be represented by  $\hat{g}_i^o = \lambda g_i^o + \rho(\{v_i^1, v_i^2, \dots, v_i^N\})$ .  $\lambda$  is a learnable weight factor.

The output of text encoder can be denoted by  $c_i = \{c_i^1, c_i^2, \dots, c_i^N\}$ , where the [EOS] embedding  $c_i^e$  is selected as the overall representation of caption. We define a similarity function as  $s(V_i, T_i) = \phi(\hat{g}_i^o)^T \psi(c_i^e)$ , where  $\phi(\cdot)$  and  $\psi(\cdot)$  are the linear projection functions that project visual and caption embedding into the shared semantic space. Then, we construct in-batch similarity between video and text:

$$L_{t2v} = -\frac{1}{B} \sum_i \log \frac{\exp(s(V_i, T_i)/\tau)}{\sum_{j=1}^B \exp(s(V_j, T_i)/\tau)}, \quad (1)$$

$$L_{v2t} = -\frac{1}{B} \sum_i \log \frac{\exp(s(V_i, T_i)/\tau)}{\sum_{j=1}^B \exp(s(V_i, T_j)/\tau)}, \quad (2)$$

$$L_{con} = \frac{1}{2}(L_{t2v} + L_{v2t}), \quad (3)$$

where  $\tau$  is a learnable temperature parameter,  $B$  is batch size,  $L_{t2v}$  and  $L_{v2t}$  represents the text-to-video loss and video-to-text loss.  $L_{con}$  is the total contrastive loss.

### Video Text Matching and Masked Language Modeling

**VTM** determines whether the pair of [visual, tag] and [text] is matched, where [visual, tag] is the **TG** cross-modal fu-

sion output of multi-frame visual features and multi-modal tag information, [text] is the overall representation of text. This pair is then fed into a joint cross-modal encoder, which shares parameters with the **TG** cross-modal encoder. The first output token embedding is regarded as the joint video-text representation  $V_i$ , and passed through a fully connected layer for binary prediction (match or not match with the text representation  $T_i$ ). Defining the precision as  $p^{vtm}(V_i, T_i)$ , the VTM loss is:

$$L_{vtm} = -\frac{1}{O} \sum_i \sum_{t=0}^1 y_{it}^{vtm} \log_2(p_t^{vtm}(V_i, T_i)), \quad (4)$$

where  $y_{it}^{vtm}$  is a sign function which has the value of 1 if  $t=1$  else 0 if  $t=0$ .  $(V_i, T_i)$  is a positive pair when  $t=1$ , and a negative pair otherwise.  $p_t^{vtm}(V_i, T_i)$  denotes the prediction probability of  $t$ .  $O$  is the total number of video-text pairs for VTM task, which is composed of positive and negative pairs. We perform hard negative mining strategy when constructing the triplet sample. For each [visual, tag] in mini-batch, we sample a negative [text] according to the in-batch similarity matrix calculated in Equation 1. Likewise, for each [text], we sample a negative [visual, tag] in the mini-batch.

**MLM** predicts the masked word based on the [visual, tag] representations and the unmasked context, which is also a classification task. To be specific, we randomly mask the text tokens with a probability of 15%, and the replacement is the [MASK] token with an 80% probability, a random token with a 10% probability, original token with a 10% probability. Supposing  $\hat{T}_i$  denotes the masked text and  $p^{mlm}(V_i, \hat{T}_i)$  denotes the prediction result of the masked word, the loss function of MLM can be expressed as:

$$L_{mlm} = -\frac{1}{Q} \sum_i \sum_{v=1}^V y_{iv}^{mlm} \log_2(p_v^{mlm}(V_i, \hat{T}_i)), \quad (5)$$

where  $y_{iv}^{mlm}$  is a sign function which has the value of 1 if the masked word of sample  $i$  is  $v$ .  $V$  is the vocabulary size and  $Q$  is the total number of video-text pairs for MTM task. Finally, the overall loss function of the TABLE model is:

$$L = L_{con} + L_{vtm} + L_{mlm} \quad (6)$$

### Inference Strategy

During the inference stage, the joint cross-modal encoder with VTM and MLM loss is discarded, and the similarity confidences between video and text are only calculated by the **TG** encoder and the text encoder. To further improve the performance, we also adopt the inference strategy proposed in CAMoE (Cheng et al. 2021), which conduct SoftMax operation to revise the similarity matrix.

## Experiments

### Datasets

**MSR-VTT** (Xu et al. 2016) dataset is widely studied in video-text retrieval task, which contains 10,000 videos of 10-32 seconds, each corresponding to 20 captions. We follow the data split in previous method (Gabeur et al. 2020), i.e., 9000 videos for training and 1,000 videos for testing.

		Text-to-Video (T2V)					Video-to-Text (V2T)				
Type	Method	R@1	R@5	R@10	MdR	MnR	R@1	R@5	R@10	MdR	MnR
NO-CLIP	JSFusion (Yu, Kim, and Kim 2018)	10.2	31.2	43.2	13.0	-	-	-	-	-	-
NO-CLIP	HT-Pretrained (Miech et al. 2019)	14.9	40.2	52.8	9.0	-	-	-	-	-	-
NO-CLIP	CE (Liu et al. 2019)	20.9	48.8	62.4	6.0	28.2	20.6	50.3	64.0	5.3	25.1
NO-CLIP	MMT-Pretrained (Gabeur et al. 2020)	26.6	57.1	69.6	4.0	24.0	27.0	57.5	69.7	3.7	21.3
NO-CLIP	TACo (Yang, Bisk, and Gao 2021)	26.7	54.5	68.2	4.0	-	-	-	-	-	-
NO-CLIP	SUPPORT-SET (Patrick et al. 2020)	27.4	56.3	67.7	3.0	-	26.6	55.1	67.5	3.0	-
NO-CLIP	FROZEN (Bain et al. 2021)	31.0	59.5	70.5	3.0	-	-	-	-	-	-
NO-CLIP	HIT-Pretrained (Liu et al. 2021)	30.7	60.9	73.2	2.6	-	32.1	62.7	74.1	3.0	-
CLIP-based	CLIP (Portillo-Quintero and Ortiz-Bayliss 2021)	31.2	53.7	64.2	4.0	-	27.2	51.7	62.6	5.0	-
CLIP-based	MDMMT (Dzabraev et al. 2021)	38.9	69.0	79.7	2.0	16.5	-	-	-	-	-
CLIP-based	CLIP4Clip-meanP (Luo et al. 2021)	43.1	70.4	80.8	2.0	16.2	43.1	70.5	81.2	2.0	12.4
CLIP-based	CLIP4Clip-seqTransf (Luo et al. 2021)	44.5	71.4	81.6	2.0	15.3	42.7	70.9	80.6	2.0	11.6
CLIP-based	VCM (Cao et al. 2022)	43.8	71.0	80.9	2.0	14.3	45.1	72.3	82.3	2.0	10.7
CLIP-based	CLIP2VIDEO (Fang et al. 2021)	45.6	72.6	81.7	2.0	14.6	43.5	72.3	82.1	2.0	10.2
CLIP-based	CAMoE (Cheng et al. 2021)	44.6	72.6	81.8	2.0	13.3	45.1	72.4	83.1	2.0	10.0
CLIP-based	<b>TABLE (ours)</b>	47.1	74.3	82.9	2.0	13.4	47.2	74.2	84.2	2.0	11.0
CLIP-based	CAMoE* (Cheng et al. 2021)	47.3	74.2	84.5	2.0	11.9	49.1	74.3	84.3	2.0	9.9
CLIP-based	CAMoE-online* (Cheng et al. 2021)	48.8	75.6	<b>85.3</b>	2.0	12.4	50.3	74.6	83.8	2.0	<b>9.9</b>
CLIP-based	<b>TABLE* (ours)</b>	<b>52.3</b>	<b>78.4</b>	85.2	<b>1.0</b>	<b>11.4</b>	<b>51.8</b>	<b>77.5</b>	<b>85.1</b>	<b>1.0</b>	10.0

Table 1: Retrieval results on the MSR-VTT dataset. \* indicates methods with inference strategy (as in subsequent tables).

**MSVD** (Chen and Dolan 2011) dataset contains 1,970 videos, each video has nearly 40 titles. Train, validation and test sets have 1,200, 100 and 670 videos respectively. Following previous works (Luo et al. 2021), we report the result on testing set with multiple captions per video. The limited training data makes the learning on this dataset challenging.

**LSMDC** (Rohrbach, Rohrbach, and Schiele 2015) dataset consists of 118,081 short videos clips extracted from 202 movies. Similar to previous work (Luo et al. 2021), we validate the performance on the test set containing 1,000 videos.

**DiDeMo** (Anne Hendricks et al. 2017) dataset involves about 10,000 videos and each video has about 4 annotated sentences. Following previous works (Luo et al. 2021), we conduct video-paragraph retrieval task on the test split (with 1,004 videos) by concatenating all sentences of each video into a single query. The challenge of this dataset lies in the alignment of long videos and long texts.

## Metrics

We evaluate the model performance with standard retrieval metrics, i.e., Recall at rank K (R@K, K=1,5,10), Median Rank (MdR) and Mean Rank (MnR). R@K represents the proportion of the ground-truth result included in the top-K recalled results. MdR and MnR represent the median and mean rank of correct results respectively. Therefore, higher R@K, lower MdR and MnR indicates better performance.

## Implementation Details

The visual encoder, tag and text encoder of our TABLE model are initialized by the pretrained CLIP (ViT-B/32) (Radford et al. 2021) model. The parameters of tag and text encoder are shared in the transformer blocks but individually learned in the linear projection layer. The position embedding and the blocks of the cross-modal encoders are initial-

ized by the first four layers of the CLIP’s text encoder. The initial learning rate is  $1e-7$  for visual and text encoder, and  $1e-4$  for cross-modal encoders. The TABLE model is trained for 5 epoches by Adam optimizer and warmup scheduler. The batch size is 128, except on DiDeMo is 48. For MSR-VTT, MSVD and LSMDC, the max token length of caption and tag, and the frame length are set to 32, 32, 12, respectively. For DiDeMo, the value of the above parameters are 64, 32, 32, respectively. To prevent the tag information from being overwhelmed by a larger number of video frames on DiDeMo, we adopt all token output of the tag encoder instead of the [EOS] output as the input of TG cross-modal encoder. We adopt 5 pretrained experts (object, person, scene, motion and audio) to extract multi-modal tags on MSR-VTT and DiDeMo dataset, and 4 experts (without audio) on MSVD and LSMDC dataset.

## Comparisons to the State of the Art

In this subsection, we compare our model with state-of-the-art methods on four representative benchmarks.

As shown in Table 1, without inference strategy, our TABLE model achieves great performance on the MSR-VTT dataset. For example, our method surpasses CAMoE (Cheng et al. 2021) by a large margin of 2.5 and 2.1 on R@1 in text-to-video (T2V) and video-to-text (V2T) task respectively. It not only demonstrates the importance of multi-modal information, but also proves the superiority of our tagging method for enabling cross-modal alignment. Furthermore, with the inference strategy, our method achieves SOTA performance, i.e., 52.3 R@1 in T2V and 51.8 R@1 in V2T task. For the MSVD dataset, as illustrated in Table 2, our TABLE method surpasses CAMoE by 3.0 on R@1 in T2V task without inference strategy, and finally achieves SOTA performance of 52.3 R@1. We believe that on scale-limited datasets, the explicit guidance of multi-modal tags is more

Method	Text-to-Video (T2V)				
	R@1	R@5	R@10	MdR	MnR
VSE (Kiros 2014)	12.3	30.1	42.3	14.0	-
CE	19.8	49.0	63.8	6.0	23.1
SSML (Amrani et al. 2020)	20.3	49.0	63.3	6.0	-
SUPPORT-SET	28.4	60.0	72.9	4.0	-
FROZEN	33.7	64.7	76.3	3.0	-
CLIP	37.0	64.1	73.8	3.0	-
CLIP4Clip-seqTransf	45.2	75.5	84.3	2.0	10.3
CLIP4Clip-meanP	46.2	76.1	84.6	2.0	10.0
CLIP2VIDEO	47.0	76.8	85.9	2.0	9.6
CAMoE	46.9	76.1	85.5	-	9.8
<b>TABLE</b>	49.9	79.3	87.4	2.0	<b>9.1</b>
CAMoE*	49.8	79.2	87.0	-	9.4
<b>TABLE* (ours)</b>	<b>52.3</b>	<b>80.5</b>	<b>87.9</b>	<b>1.0</b>	9.8

Table 2: Retrieval results on the MSVD dataset.

Method	Text-to-Video (T2V)				
	R@1	R@5	R@10	MdR	MnR
JSFusion	9.1	21.2	34.1	36.0	-
CE	11.2	26.9	34.8	25.3	96.8
HIT-Pretrained	14.0	31.2	41.6	18.5	-
FROZEN	15.0	30.8	39.8	20.0	-
CLIP	11.3	22.7	29.2	56.5	-
MDMMT	18.8	38.5	47.9	12.3	58.0
CLIP4Clip-meanP	20.7	38.9	47.2	13.0	65.3
CLIP4Clip-seqTransf	22.6	41.0	49.1	11.0	61.0
CAMoE	22.5	42.6	50.9	-	56.5
<b>TABLE (ours)</b>	24.3	44.9	53.7	8.0	52.7
CAMoE*	25.9	<b>46.1</b>	53.7	-	54.4
<b>TABLE* (ours)</b>	<b>26.2</b>	45.9	<b>55.0</b>	<b>7.0</b>	<b>51.0</b>

Table 3: Retrieval results on the LSMDC dataset.

important for visual-text alignment. LSMDC dataset contains the most videos and each video matches only one caption, and thus most methods perform poorly. But our method still improves 1.8 on R@1 and achieves SOTA performance (26.2 R@1) in T2V task, as shown in Table 3. Finally, for the video-paragraph retrieval task on DiDeMo, our method obtains 5.3 R@1 gains over SOTA method CAMoE, as presented in Table 4. The alignment of long videos and long text meets greater challenge compared to short video-text pairs, and thus the introduction of multi-modal tags shows greater advantages on this dataset. In conclusion, our TABLE model proves to be effective and superior by achieving SOTA performance on various benchmarks.

## Abalation Studies

**Effects of Multi-Modal Tags.** In this subsection, we delve into the impact of multi-modal tags on model’s performance.

As shown in Table 5, our backbone model achieves R@1 of 45.2 for text-to-video (T2V) retrieval, which is already a comparable performance on the MSR-VTT dataset. The R@1 of T2V is improved to 45.6 after the object tags are introduced by our TABLE model, and the R@1 of V2T is further improved to 45.9 with object and person tags. It proves that the visual-text alignment benefits from additional tag in-

Method	Text-to-Video (T2V)				
	R@1	R@5	R@10	MdR	MnR
S2VT (Venugopalan et al. 2014)	11.9	33.6	-	13.0	-
FSE (Zhang and Hu 2018)	13.9	36.0	-	11.0	-
CE	16.1	41.1	-	8.3	43.7
TT-CE (Croitoru et al. 2021)	21.6	48.6	62.9	6.0	-
FROZEN	34.6	65.0	74.7	3.0	-
ClipBERT (Lei et al. 2021)	20.4	48.0	60.8	6.0	-
CLIP4Clip-seqTransf	42.8	68.5	79.2	2.0	18.9
CLIP4Clip-meanP	43.4	70.2	80.6	2.0	17.5
<b>TABLE (ours)</b>	47.9	74.0	82.1	2.0	<b>14.3</b>
CAMoE*	43.8	71.4	79.9	2.0	16.3
<b>TABLE* (ours)</b>	<b>49.1</b>	<b>75.6</b>	<b>82.9</b>	<b>2.0</b>	14.8

Table 4: Retrieval results on the DiDeMo dataset.

formation. The object and person tags contain local information, which can guide the local attention of visual features. After adding motion tags, the performance is obviously improved, i.e., the T2V R@1 reaches 46.2. It is further promoted to 46.5 with scene tags, which usually describe the background information. Finally, the T2V R@1 is improved to 46.8 with audio tags, which contain information not included in visual features, which is helpful for the retrieval of videos like news report. To sum up, introducing multi-modal tags is beneficial for video-text retrieval task.

**Effects of VTM and MLM.** In this subsection, we explore the effect of VTM and MLM on model’s performance.

As shown in Table 5, with the help of VTM, the video-to-text retrieval performance is improved from 45.9 to 46.6 at R@1. Using both VTM and MLM, the model performance is further promoted. The R@1 of text-to-video reaches 47.1, and the R@1 of video-to-text achieves 47.2. Although VTM is a binary classification task, the model can only make well judgment on hard negative with profound understanding of video and text content. MLM is a more complex task which requires fine-grained alignment between the visual features and text tokens for inferring masked words. In general, these two tasks strengthen the interaction between video and text, and thus effectively improves the model performance.

## Visual Analysis

To clearly reveal the anchor role of multi-modal tags, we present some visualization results of the **TG** cross-modal encoder and the joint cross-modal encoder using the Attention Rollout (Abnar and Zuidema 2020) method.

In Fig. 3(a) and Fig. 3(b), the multi-modal tags are treated as an entirety for calculating cross attention with visual features. As shown, the cross attention focuses on frames that highly correlated with multi-modal tags. For example, in Fig. 3(a), the 1st frame (highly correlated with “man” and “ballplayer”) and the 10th frame (highly correlated with “catching or throwing baseball” and “baseball glove”) are highlighted, while the 2rd-4th frames are suppressed. In each frame, important spatial regions like head, body, baseball glove are getting more attention. Although the multi-modal tags might contain noise, e.g., “bird” in Fig. 3(a), they can still guide the model to pay more attention to key frames and key regions. We also find that the motion modality show



Method	Tag Modalities					Task		Text-to-Video (T2V)					Video-to-Text (V2T)				
	Object	Person	Motion	Scene	Audio	VTM	MLM	R@1	R@5	R@10	MdR	MnR	R@1	R@5	R@10	MdR	MnR
Baseline	-	-	-	-	-	-	-	45.2	72.0	81.3	2.0	14.3	45.4	73.8	83.8	2.0	11.5
TABLE	✓	-	-	-	-	-	-	45.6	72.8	81.7	2.0	13.8	45.3	73.6	83.9	2.0	11.1
TABLE	✓	✓	-	-	-	-	-	45.6	72.8	81.9	2.0	14.1	45.9	74.9	83.4	2.0	11.1
TABLE	✓	✓	✓	-	-	-	-	46.2	73.0	82.2	2.0	<b>13.8</b>	45.9	74.7	83.8	2.0	<b>11.0</b>
TABLE	✓	✓	✓	✓	-	-	-	46.5	73.4	82.2	2.0	13.9	45.6	74.8	83.6	2.0	11.2
TABLE	✓	✓	✓	✓	✓	-	-	<b>46.8</b>	<b>73.5</b>	<b>82.2</b>	<b>2.0</b>	13.9	<b>45.9</b>	<b>75.1</b>	<b>84.1</b>	<b>2.0</b>	11.1
TABLE	✓	✓	✓	✓	✓	✓	-	46.7	73.9	82.1	2.0	13.6	46.6	<b>75.0</b>	<b>84.7</b>	2.0	<b>10.6</b>
TABLE	✓	✓	✓	✓	✓	✓	✓	<b>47.1</b>	<b>74.3</b>	<b>82.9</b>	<b>2.0</b>	<b>13.4</b>	<b>47.2</b>	74.2	84.2	<b>2.0</b>	11.0

Table 5: Ablation Results on MSR-VTT dataset.

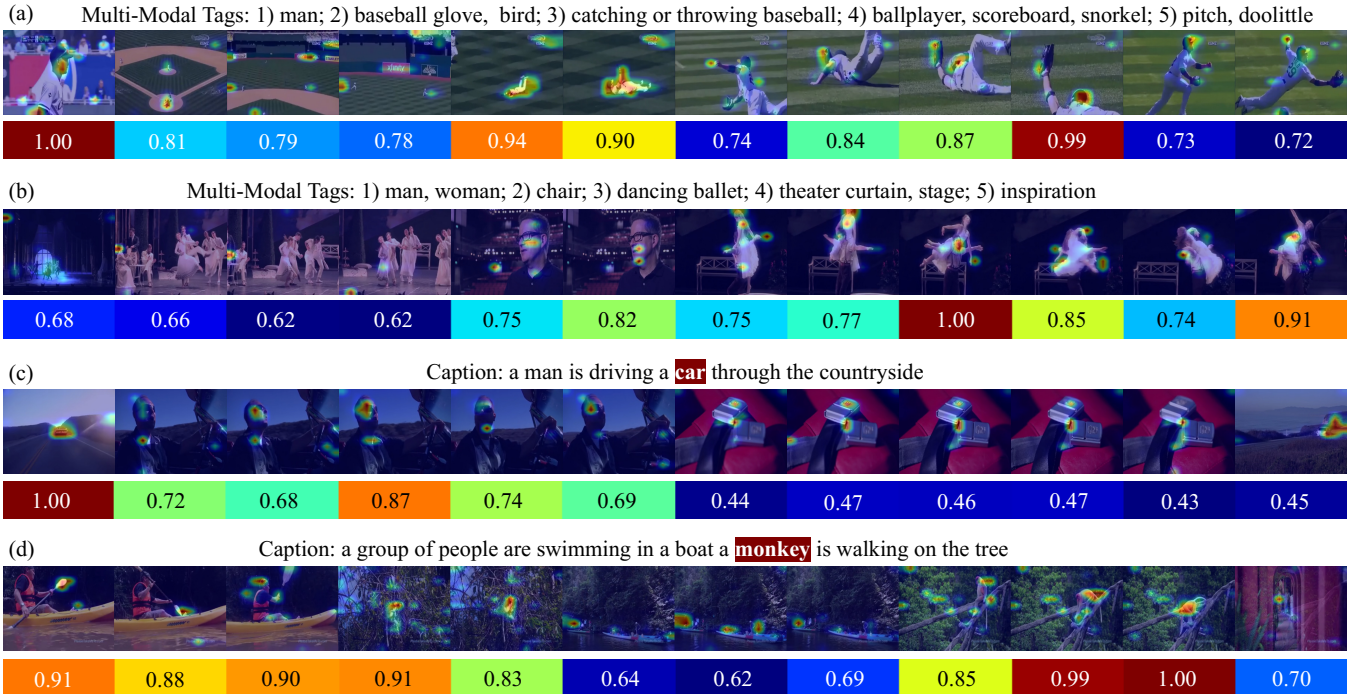


Figure 3: We visualize the spatial and temporal attention of TABLE, where brighter regions in each frame represent higher spatial attention, and redder color in the color bar represents higher temporal attention.

dominant role in the cross-modal attention. For example, in Fig. 3(b), the 9th and 12th frames describing standard “dancing ballet” are highlighted, while the first few frames about scene and person are suppressed. To sum up, the **TG** cross-modal encoder is capable of selecting key frames and key regions from redundant visual features with the guidance of multi-modal tags, which is beneficial for video-text retrieval.

In Fig. 3(c) and Fig. 3(d), we visualize the cross attention in the joint cross-modal encoder. Note that the cross attention is calculated between visual features and specific word token of the caption, not the entire text. In Fig. 3(c), given the “car” word, the 1st frame is highlighted as it gives the whole picture of the car, while the last few frames about seat belts are suppressed. In Fig. 3(d), the 10th and 11th frames are getting more temporal attention because they describe the monkey with a close view. The other frames describing about “people” or “swimming” are suppressed. It can be concluded that the proposed joint cross-modal encoder is capa-

ble of modeling the fine-grained correlation between video and a single word of text, which is inherited from VTM and MLM tasks. Moreover, it reveals the strong capability of **TG** cross-modal encoder as their parameters are shared.

## Conclusion

In this paper, we propose to integrate multi-modal tags as anchors to motivate the video-text alignment. Specifically, we construct a TABLE model to jointly encode multi-frame visual features and multi-modal information. To further enhance the video-text interaction, we introduce VTM and MLM tasks with the triplet input of [visual, tag, text] as auxiliary supervisions. TABLE achieves SOTA performance on various benchmarks, including MSR-VTT, MSVD, LSMDC and DiDeMo, which indicates the superiority of introducing multi-modal tags as anchors for video-text retrieval task. The ablation studies and visualization results further reveal the anchor role of multi-modal tags in visual-text alignment.

## References

- Abnar, S.; and Zuidema, W. 2020. Quantifying attention flow in transformers. *arXiv preprint arXiv:2005.00928*.
- Amrani, E.; Ben-Ari, R.; Rotman, D.; and Bronstein, A. 2020. Noise estimation using density estimation for self-supervised multimodal learning. *arXiv preprint arXiv:2003.03186*, 8.
- Anne Hendricks, L.; Wang, O.; Shechtman, E.; Sivic, J.; Darrell, T.; and Russell, B. 2017. Localizing moments in video with natural language. In *Proceedings of the IEEE international conference on computer vision*, 5803–5812.
- Arnab, A.; Dehghani, M.; Heigold, G.; Sun, C.; Lučić, M.; and Schmid, C. 2021. Vivit: A video vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6836–6846.
- Bain, M.; Nagrani, A.; Varol, G.; and Zisserman, A. 2021. 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*, 1728–1738.
- Bertasius, G.; Wang, H.; and Torresani, L. 2021. Is space-time attention all you need for video understanding. *arXiv preprint arXiv:2102.05095*, 2(3): 4.
- Cao, S.; Wang, B.; Zhang, W.; and Ma, L. 2022. Visual Consensus Modeling for Video-Text Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Carreira, J.; and Zisserman, A. 2017. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6299–6308.
- Chen, D.; and Dolan, W. B. 2011. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, 190–200.
- Chen, Q.; Liu, Y.; and Albanie, S. 2021. Mind-the-Gap! Unsupervised Domain Adaptation for Text-Video Retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 1072–1080.
- Cheng, X.; Lin, H.; Wu, X.; Yang, F.; and Shen, D. 2021. Improving video-text retrieval by multi-stream corpus alignment and dual softmax loss. *arXiv preprint arXiv:2109.04290*.
- Croitoru, I.; Bogolin, S.-V.; Leordeanu, M.; Jin, H.; Zisserman, A.; Albanie, S.; and Liu, Y. 2021. Teachtext: Cross-modal generalized distillation for text-video retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11583–11593.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Dzabraev, M.; Kalashnikov, M.; Komkov, S.; and Petiushko, A. 2021. Mdmmt: Multidomain multimodal transformer for video retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3354–3363.
- Fang, H.; Xiong, P.; Xu, L.; and Chen, Y. 2021. Clip2video: Mastering video-text retrieval via image clip. *arXiv preprint arXiv:2106.11097*.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2019. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6202–6211.
- Gabeur, V.; Sun, C.; Alahari, K.; and Schmid, C. 2020. Multi-modal transformer for video retrieval. In *European Conference on Computer Vision*, 214–229. Springer.
- Grootendorst, M. 2020. KeyBERT: Minimal keyword extraction with BERT. <https://doi.org/10.5281/zenodo.4461265>.
- Hao, X.; Zhou, Y.; Wu, D.; Zhang, W.; Li, B.; and Wang, W. 2021. Multi-Feature Graph Attention Network for Cross-Modal Video-Text Retrieval. In *Proceedings of the 2021 International Conference on Multimedia Retrieval*, 135–143.
- Kiros, R. 2014. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*.
- Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Kolesnikov, A.; et al. 2020. The open images dataset v4. *International Journal of Computer Vision*, 128(7): 1956–1981.
- Lei, J.; Li, L.; Zhou, L.; Gan, Z.; Berg, T. L.; Bansal, M.; and Liu, J. 2021. 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*, 7331–7341.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Liu, S.; Fan, H.; Qian, S.; Chen, Y.; Ding, W.; and Wang, Z. 2021. Hit: Hierarchical transformer with momentum contrast for video-text retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11915–11925.
- Liu, Y.; Albanie, S.; Nagrani, A.; and Zisserman, A. 2019. Use what you have: Video retrieval using representations from collaborative experts. *arXiv preprint arXiv:1907.13487*.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2021. Clip4clip: An empirical study of clip for end to end video clip retrieval. *arXiv preprint arXiv:2104.08860*.
- Miech, A.; Zhukov, D.; Alayrac, J.-B.; Tapaswi, M.; Laptev, I.; and Sivic, J. 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*, 2630–2640.
- Patrick, M.; Huang, P.-Y.; Asano, Y.; Metzger, F.; Hauptmann, A.; Henriques, J.; and Vedaldi, A. 2020. Support-set bottlenecks for video-text representation learning. *arXiv preprint arXiv:2010.02824*.
- Portillo-Quintero, J. A.; and Ortiz-Bayliss, J. C. 2021. A straightforward framework for video retrieval using clip.



In *Mexican Conference on Pattern Recognition*, 3–12. Springer.

Qian, S.; Xue, D.; Zhang, H.; Fang, Q.; and Xu, C. 2021. Dual adversarial graph neural networks for multi-label cross-modal retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2440–2448.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 8748–8763. PMLR.

Rohrbach, A.; Rohrbach, M.; and Schiele, B. 2015. The long-short story of movie description. In *German conference on pattern recognition*, 209–221. Springer.

Stewart, D.; Seymour, R.; Pass, A.; and Ming, J. 2013. Robust audio-visual speech recognition under noisy audio-video conditions. *IEEE transactions on cybernetics*, 44(2): 175–184.

Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; and Paluri, M. 2015. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 4489–4497.

Venugopalan, S.; Xu, H.; Donahue, J.; Rohrbach, M.; Mooney, R.; and Saenko, K. 2014. Translating videos to natural language using deep recurrent neural networks. *arXiv preprint arXiv:1412.4729*.

Wang, X.; Zhu, L.; and Yang, Y. 2021. T2vlad: global-local sequence alignment for text-video retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5079–5088.

Xie, S.; Sun, C.; Huang, J.; Tu, Z.; and Murphy, K. 2018. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European conference on computer vision (ECCV)*, 305–321.

Xu, J.; Mei, T.; Yao, T.; and Rui, Y. 2016. Msr-vtt: A large video description dataset for bridging video and language. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5288–5296.

Yang, J.; Bisk, Y.; and Gao, J. 2021. Taco: Token-aware cascade contrastive learning for video-text alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11562–11572.

Yu, Y.; Kim, J.; and Kim, G. 2018. A joint sequence fusion model for video question answering and retrieval. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 471–487.

Zhang, B.; and Hu, H. 2018. Cross-modal and hierarchical modeling of video and text. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 374–390.

Zhen, L.; Hu, P.; Wang, X.; and Peng, D. 2019. Deep supervised cross-modal retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10394–10403.

Zhu, L.; and Yang, Y. 2020. Actbert: Learning global-local video-text representations. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8746–8755.