

STOA-VLP: Spatial-Temporal Modeling of Object and Action for Video-Language Pre-training

Weihong Zhong¹, Mao Zheng², Duyu Tang³, Xuan Luo², Heng Gong¹,
Xiaocheng Feng^{1,4*}, Bing Qin^{1,4}

¹ Harbin Institute of Technology

² Tencent MLPD

³ Independent Researcher

⁴ Peng Cheng Laboratory

{whzhong, hgong, xcfeng, qinb}@ir.hit.edu.cn, {moonzheng, tssxuanluo}@tencent.com,
tangduyu.nlp@hotmail.com

Abstract

Although large-scale video-language pre-training models, which usually build a global alignment between the video and the text, have achieved remarkable progress on various downstream tasks, the idea of adopting fine-grained information during the pre-training stage is not well explored. In this work, we propose STOA-VLP, a pre-training framework that jointly models object and action information across spatial and temporal dimensions. More specifically, the model regards object trajectories across frames and multiple action features from the video as fine-grained features. Besides, We design two auxiliary tasks to better incorporate both kinds of information into the pre-training process of the video-language model. The first is the dynamic object-text alignment task, which builds a better connection between object trajectories and the relevant noun tokens. The second is the spatial-temporal action set prediction, which guides the model to generate consistent action features by predicting actions found in the text. Extensive experiments on three downstream tasks (video captioning, text-video retrieval, and video question answering) demonstrate the effectiveness of our proposed STOA-VLP (e.g. 3.7 Rouge-L improvements on MSR-VTT video captioning benchmark, 2.9% accuracy improvements on MSVD video question answering benchmark, compared to previous approaches).

Introduction

With the introduction of large-scale datasets, video-language pre-training methods learn to build adaptive cross-modal representations for downstream video-text understanding and generation tasks. Previous dominant approaches focus on learning an overall alignment from raw video and text data (Lei et al. 2021; Bain et al. 2021), lacking modeling of fine-grained information. Meanwhile, fine-grained information can be necessary for understanding the video scene in many situations. For example, as Figure 1(a) depicts, generating such captions requires detailed dynamic modeling of relevant objects and movements across time. To answer questions like Figure 1(b), a fine-grained object-text alignment is required, as well as object movements and their dif-

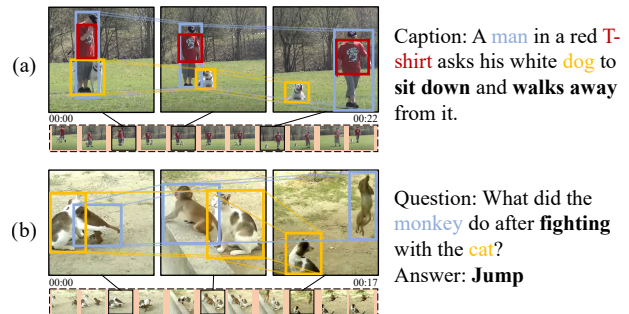


Figure 1: (a): A video captioning example. Recognitions of salient objects and movements of objects are important clues for generating the caption. (b): A video question answering example. To answer this question, a fine-grained connection between salient objects should be built, as well as object movements and their different interactions should be carefully analyzed.

ferent interactions should be carefully analyzed. A coarse-grained method can easily fail at handling such information.

Efforts have been made by previous pre-training approaches to maintain fine-grained information from different aspects. Although achieved promising improvements, these methods still have some limitations: (1) Objects are usually moving and changing across time. Existing methods (Li et al. 2022; Wang et al. 2022) usually utilize fixed bounding boxes to model objects, leading to an insufficient understanding of object movements in the video. (2) There can be multiple overlapping actions at the same time range, triggered by different object interactions in different spatial regions. Existing methods (Zhu and Yang 2020) focus on only one action representation for each video clip, which leads to an information loss. As Figure 1 shows, the main content of the video can be described by the movement of objects and the interaction between them. Therefore, an ideal approach is to incorporate these two kinds of information into the pre-training stage for fine-grained spatial-temporal modeling.

In this paper, we tackle the above problems by proposing a new pre-training framework—STOA-VLP: Spatial-Temporal Modeling of Object and Action for Video-

*Corresponding Author

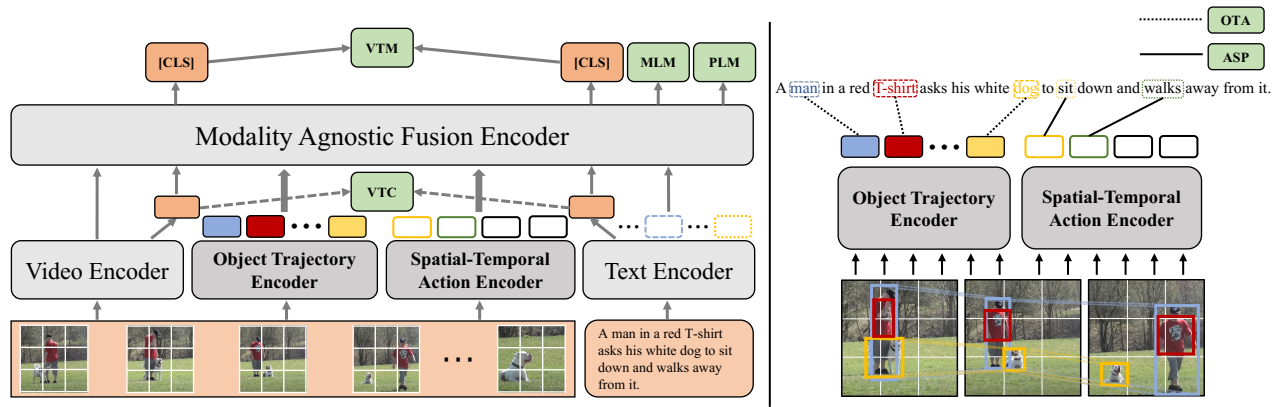


Figure 2: Left: Overview of our STOA-VLP’s architecture. STOA-VLP contains a video encoder, a text encoder, as well as two newly proposed encoders: the object trajectory encoder and the spatial-temporal action encoder, which are used to explicitly model dynamic object trajectories and action information, respectively. On top of them, a modality-agnostic fusion encoder is built to enable further cross-modal interaction. Right: We design two auxiliary tasks to build a fine-grained alignment with the text. In the object-text alignment (OTA) task, we align object trajectories with noun tokens in the paired text. In the action set prediction (ASP) task, we predict actions found in the text with generated action features to guide the training of the action encoder.

Language Pre-training that captures and jointly models object trajectories and multiple actions to better handle movements and interactions of objects in the video. More specifically, we automatically extract bounding boxes from multiple frames. Temporal-aware object trajectories are generated using raw video patches and bounding boxes across frames with identical labels. Next, raw video patches and bounding boxes from different frames are utilized to generate multiple spatial-temporal action features with different action query tokens. In addition, to better incorporate object trajectories and action features, two novel auxiliary tasks are designed to build a fine-grained alignment with the text (see the right of Figure 2). The first one is dynamic object-text alignment (OTA), which calculates an optimal bipartite matching between generated object trajectories and noun tokens in the paired text and maximizes the matching score to dynamically connect the object feature with the paired text. The second one is spatial-temporal action set prediction (ASP), which firstly constructs an action set out of the text, then constructs a prediction matrix with action features to find an optimal bipartite matching between them. The matching score is maximized to enhance their alignment.

Finally, we evaluate our STOA-VLP on three commonly used tasks including video captioning, text-video retrieval, and video question answering. Our framework achieves substantial improvements over other methods on most of the metrics with only 2.5M pre-training data. To sum up, our contributions are:

- We propose STOA-VLP that explicitly models temporal-aware object trajectories and multiple spatial-temporal action features to better handle the dynamic movements and interactions of objects in the video.¹

¹Our code will be available at <https://github.com/whongzhong/STOA-VLP>.

- We design two novel auxiliary pre-training tasks—OTA and ASP to better model object trajectories and action features during the pre-training stage.
- Our STOA-VLP surpasses previous approaches on most of the metrics, and achieves a large performance gain over our base model, for example, achieving 1.9 CIDER improvements on the MSR-VTT video captioning benchmark and 1.2% absolute accuracy improvements on the MSVD video question answering benchmark.

Related Work

Video-Language Pre-training

Earlier works usually adopt a dense sampling way to extract consecutive frames. Due to the limit of computational cost when extracting offline features from videos (Sun et al. 2019b,a; Luo et al. 2020) using ResNet (He et al. 2016), S3D (Xie et al. 2018) or Faster r-cnn (Ren et al. 2015), etc., information from consecutive frames tends to be similar, which inspires the idea of sparse sampling. With this thought, Lei et al. (2021) enable the end-to-end pre-training with image-text data and fine-tune with only a few sampled video frames, demonstrating the potential of knowledge transfer from image to video. Furthermore, Bain et al. (2021) and Luo et al. (2021) adopt a dual encoder structure to align video and text with simple instance-level contrastive objectives. However, different from image data, scenes are changing across video frames, which makes it important to build a temporal alignment. As far as temporal information is concerned, Li et al. (2020) try to model frame-level temporal representations. Zellers et al. (2021) and Tang, Lei, and Bansal (2021) further build an alignment between video clips and temporal marked sentences.

Most of these works focus on learning an overall connection between video and text modality, without explicitly han-

dling fine-grained information. Though achieved advanced results, it can be insufficient when detailed semantic information is needed (e.g. in the video captioning and video question answering tasks).

Object and Action Modeling in Video-Language Pre-training

Since objects and their interactions can describe the major content of a video, fine-grained information considered by most works can fall into two categories, i.e. object and action information. (1) To model object information, Wang et al. (2022) and Li et al. (2022) utilize fixed object regions to build a fine-grained video-text alignment, which can be insufficient to handle moving objects and changing scenes. previous video captioning works have tried to construct a graph to model object interactions based on their spatial locations in consecutive frames (Zhao, Li, and Lu 2018; Pan et al. 2020), which can be hard when sampled frames are sparse or changing fast. Herzig et al. (2022) propose to implicitly refine the video encoder with object bounding box features without further explicitly modeling inter-frame relationships. (2) To model actions, Zhu and Yang (2020) explicitly train an action recognizer to generate a single action feature from a whole video clip, ignoring the multiple actions triggered by different object interactions. Ge et al. (2022) take a step further to consider object and action information concurrently. Meanwhile, they construct verb and noun multiple-choice questions, implicitly answering them with the intermediate video tokens to refine the instance-level alignment for a dual-encoder structure. Without further modeling the cross-modal interaction limits its application.

Since objects are moving and interacting across video frames, as well as multiple actions can be triggered by different interactions, those mentioned above are insufficient to fully address detailed semantic features. Our proposed method jointly considers modeling object and action information at both spatial and temporal dimensions. With our designed OTA and ASP tasks, a coarse-to-fine video-text alignment is built to learn a better multi-modality representation.

Method

In this section, we first introduce the overall model architecture of our STOA-VLP and further discuss the proposed object trajectory encoder and spatial-temporal action encoder. Then, we present the pre-training objectives used to model overall and fine-grained cross-modal alignment, especially our proposed auxiliary tasks OTA and ASP.

STOA-VLP Model Architecture

Overall Architectures The overall model structure of our STOA-VLP is shown on the left of Figure 2, which contains four modality-specific encoders: the video encoder, the text encoder, the object trajectory encoder, and the spatial-temporal action encoder. Firstly, raw video and text data are separately encoded by the video and text encoder. We sparsely extract object bounding box features from the

video. The intermediate encoded video features and extracted bounding box features are utilized to generate object trajectories and action features with the corresponding encoders. A modality-agnostic fusion encoder is built on top of these four encoders, in which the overall video and text information, object trajectories, and spatial-temporal-based action information interact with each other. All five modules adopt the Transformer (Vaswani et al. 2017) structure. We will further introduce each module in order:

Video Encoder Given a video with uniformly sparsely sampled T frames, our video encoder extracts $H \times W$ non-overlapping image patches from each frame and linearly projects them into 1-D tokens. We denote the intermediate patch output as $z \in \mathbb{R}^{T \times HW \times h}$, where h is the dimension of the visual representation. Following (Dosovitskiy et al. 2020), we utilize [CLS] tokens prepended to each frame as the frame level representation and linearly project them into cross-modal representation space. The final output of video encoder is $v = \{v_1, v_2, \dots, v_T\}$. where $v_i \in \mathbb{R}^{1 \times d}$ and d is the dimension of the cross-modal representation.

Text Encoder For the given paired text, we tokenize it and add two special tokens [CLS] and [SEP] to the beginning and end respectively. Assuming that the processed text is of length L , we use a text encoder to encode it and acquire the output $t = \{t_1, t_2, \dots, t_L\}$, where $t_i \in \mathbb{R}^{1 \times d}$.

Object Trajectory Modeling As mentioned in the Introduction, lacking the modeling of objects’ movement between frames might result in sub-optimal performance. Therefore, we propose to explicitly model object trajectories between frames to better capture their temporal changing. Given a sparsely sampled video with T frames, we first extract objects per frame with the off-the-shelf object detector VinVL (Zhang et al. 2021). For each video, We select the top-K distinct objects for frame τ and keep the bounding box feature (i.e. coordinates) $O_{\text{box}}^\tau \in \mathbb{R}^{K \times 4}$ and object class $O_c^\tau \in \mathbb{R}^K$. Instead of directly inserting extracted features, We follow a flexible way from (Herzig et al. 2022) to get object features via a RoIAlign layer (He et al. 2017). With the intermediate video patch output $z \in \mathbb{R}^{T \times HW \times h}$ and bounding box features O_{box}^τ , we can get object region features of K objects in frame τ with the following equation:

$$O^\tau = \text{MLP}(\text{MaxPool}(\text{RoIAlign}(z, O_{\text{box}}^\tau))), \quad (1)$$

where object features of frame τ is $O^\tau \in \mathbb{R}^{K \times h}$. Note that a spatial-temporal positional embedding is added for each object. To get N object trajectory tokens across all video frames, we sum up the classification confidence for each distinct object label in the O_c within all sampled video frames and select the top-N object class according to the summed confidence. We flatten the extracted object features and get $O \in \mathbb{R}^{(TN) \times h}$. For each selected object class c , we generate a class mask $M_c \in \mathbb{R}^{TN}$, the corresponding place of M_c will be set as 0 to be masked if it’s not labeled as c in the O_c . Therefore, with mask M_c , objects labeled as c is visible, while the rest is masked. Despite the fact that this could be a noisy process, We view this mask object sequence as a trajectory for object c , since this sequence represents its movements across frames. Similar to our video encoder, we add

a [CLS] token to the start of each object trajectory sequence. Next, we input the object feature representation O and mask M_c into the object trajectory encoder to model the object trajectory between frames. the output of [CLS] is utilized as a single temporal-aware trajectory representation for each selected object class. We denote the final projected output of temporal-aware object representation as $o \in \mathbb{R}^{N \times d}$ for each video.

Spatial-Temporal Action Modeling The key cues for recognizing actions in videos are: the objects in the scene, their interactions, and their movement (Herzig et al. 2022). Modeling action information naturally requires catching above mentioned clues in the spatial-temporal dimension. Different from (Zhu and Yang 2020) which models a global action for each video clip, we assume there possibly exist multiple actions, and each action can be inferred from different object interactions across frames. Thus, we initialize M action tokens $q = \{q_1, q_2, \dots, q_M\} \in \mathbb{R}^{M \times h}$ as query to extract action information from videos. Similar to (Herzig et al. 2022), we concatenate the intermediate patch features z with the object token features O to better catch action clues with a local information enhancement, getting feature sequence $F \in \mathbb{R}^{T \times (HW+K) \times h}$. Firstly, we utilize each query token to calculate different patch-level spatial clues via the attention mechanism for each frame:

$$x_i = \text{softmax}(q_i \cdot F^T)F, \quad (2)$$

where frame-wise representation $x_i \in \mathbb{R}^{T \times h}$. We conduct a similar process with our object trajectory encoder to add a [CLS] token to the start of x_i . Finally, we utilize the Spatial-Temporal Action Encoder to model the temporal clue across frames. The projected [CLS] output $x \in \mathbb{R}^{M \times d}$ is expected to encode spatial-temporal features needed for modeling the action feature. To better guide the training procedure, we propose a spatial-temporal action set prediction task, which will be illustrated in the following section.

Modality-agnostic Encoder Given encoded frame-level video representation $v = \{v_1, v_2, \dots, v_T\}$, we first do mean-pooling over the entire representation, getting video representation $v_{[\text{CLS}]}$. All video features v are then projected into the same projection dimension d . The final concatenated video representation is $v = \{v_1, v_2, \dots, v_T, v_{[\text{CLS}]}\}$, along with the above mentioned text feature t , temporal-aware object feature o and spatial-temporal action feature x . The final concatenated feature is $[v, o, x, t] \in \mathbb{R}^{(T+1+N+M+L) \times d}$, which will be input to the modality-agnostic encoder to do a further cross-modal interaction. We take the output of $v_{[\text{CLS}]}$ token v^o as the overall feature representation of the video, following (Radford et al. 2021) to take the output of $t_{[\text{SEP}]}$ token t^o as the overall text representation.

Pre-training Objectives

As shown in Figure 2, our STOA-VLP is pre-trained with four types of pre-training tasks, including two canonical types (i.e. overall video-text alignment and conditional language modeling). In overall video-text alignment, two instance-level pre-training tasks: video-text contrastive

(VTC) loss and video-text matching (VTM) are utilized to build a global alignment between paired video and text. In Conditional language modeling, mask language modeling (MLM) and prefix language modeling (PLM) are incorporated to strengthen the modality-agnostic encoder’s ability to generate fluent text conditioned on the given videos and texts. Besides those canonical objectives, we propose two new pre-training tasks—dynamic object-text alignment (OTA) and spatial-temporal action set prediction (ASP) to align temporal-aware objects with text and to learn to extract relevant spatial-temporal action features, respectively. We will further introduce these objectives next.

Overall Video-Text Alignment VTC and VTM loss is used to learn an overall alignment between paired video and text. For modality-specific video and text encoders, we utilize symmetric contrastive mechanism (i.e. VTC) to encourage an early alignment:

$$\mathcal{L}_{t2v} = - \sum_i \log \frac{\exp(\text{sim}(v^i, t^i))}{\sum_j \exp(\text{sim}(v^j, t^i))}, \quad (3)$$

$$\mathcal{L}_{v2t} = - \sum_i \log \frac{\exp(\text{sim}(v^i, t^i))}{\sum_j \exp(\text{sim}(v^i, t^j))}, \quad (4)$$

$$\mathcal{L}_{\text{VTC}} = (\mathcal{L}_{t2v} + \mathcal{L}_{v2t})/2, \quad (5)$$

where v^i is the video representation $v_{[\text{CLS}]}$ for sample i , t^i is the text representation $t_{[\text{SEP}]}$ for sample i . Note that the $\text{sim}(\cdot)$ operation is based on the cosine distance. Then, following (Li et al. 2021), hard negative video-text combinations that are similar with each true video-text pairs is sampled using similarity matrix $\text{sim}(v, t)$ compute for VTC loss. The VTM is applied to distinguish the original video-text pairs from constructed ones. Therefore a late alignment is built after the modality-agnostic encoder.

Conditional Language Modeling We utilize conditional language modeling tasks to further enhance the modality-agnostic encoder’s conditional text generation ability. We randomly mask 15% of input text tokens and predict them over the whole vocabulary.

We follow (Wang et al. 2021) to apply the PLM task, which masks a random length suffix of the given text, and tries to predict the mask place with the prefix text left.

Dynamic Object-Text Alignment To further enhance the alignment between the extracted object trajectory information and the paired text, we propose a dynamic object-text alignment. The idea is to align temporal-aware object trajectories with relevant objects from the text. We use SpaCy² to apply POS-tagging on all texts. Text encoder’s output of text tokens tagged as nouns are denoted as $t_{\text{noun}} \in \mathbb{R}^{L_{\text{noun}} \times d}$. We find a optimal bipartite matching between extracted noun tokens t_{noun} and temporal-aware object trajectory o_t based on the Hungarian algorithm (Kuhn 1955) to find an optimal alignment between object trajectories and relevant noun tokens. We assume the found optimal matching is the contextualized ground truth alignment, and optimize the alignment

²<https://spacy.io/>

	PT data	MSR-VTT				MSVD			
		C	M	R	B-4	C	M	R	B-4
ORG-TRL	-	50.9	28.8	62.1	43.6	95.2	36.4	73.9	54.3
SemSynAN	-	51.9	30.4	64.7	46.4	111.5	41.9	79.5	64.4
APML	-	52.2	30.3	63.6	43.8	108.3	39.2	76.2	58.0
VNS-GRU	-	53.0	29.9	63.4	45.3	121.5	42.1	79.7	66.5
SAM-SS	-	53.2	29.3	63.6	45.8	109.7	39.0	77.0	62.4
SwinBERT	-	53.8	29.9	62.1	41.9	120.6	41.3	77.5	58.2
Clip4Caption	136M	57.7	30.7	63.7	46.1	-	-	-	-
UniPerceiver	45.4M	-	-	-	-	131.0	42.3	79.0	61.5
MV-GPT†	53M	60.0	38.7	64.0	48.9	-	-	-	-
STOA-VLP	2.5M	60.2	31.0	68.4	45.8	131.8	43.4	83.9	64.4

Table 1: Video captioning results on the MSR-VTT and MSVD datasets. †: using subtitles as an additional input.

with the following equation:

$$[o_{\text{ind}}, t_{\text{ind}}] = \text{matching}(\text{Softmax}(\text{sim}(o_t, t_{\text{noun}}))), \quad (6)$$

$$\mathcal{L}_{\text{OTA}} = \text{MSE}(o_{\text{ind}}, t_{\text{ind}}), \quad (7)$$

where $|o_{\text{ind}}| = |t_{\text{ind}}| = \min(K, L_{\text{noun}})$. $o_{\text{ind}}, t_{\text{ind}}$ represents matching tokens from o_t and t_{noun} , respectively. The overall text representation $t_{[\text{CLS}]}$ is included in t_{noun} . Being matched with the $t_{[\text{CLS}]}$ represents the corresponding object trajectory fails to be matched with any specific noun token. Note that the text encoder is frozen during this process. By optimizing this objective, the object trajectory encoder should produce a spatial-temporal aware feature that connects closer to the paired text.

Spatial-Temporal Action Set Prediction Instead of carefully labeling action sequences for the video which might not be relevant to the text due to the transfer of domains, a more flexible way to guide the spatial-temporal action modeling procedure is to predict action out of verb tokens from the paired text. To achieve this goal, we extract tokens tagged as verbs similar to OTA. The action sequence from text usually is not in chronological order, since humans tend to summarize and reconstruct information from the video. Thus, we follow (Carion et al. 2020) to define a action set $s_{\text{act}} = \text{Set}(t_{\text{act}}) \cup \{\emptyset\}$. An optimal matching can be calculated with equation (8) between action feature x and action set s_{act} . The matching result is viewed as the prediction ground truth. The prediction loss is calculated as follows:

$$[x_{\text{ind}}, s_{\text{ind}}] = \text{matching}(\text{Softmax}(\text{sim}(x_t, s_{\text{act}}))), \quad (8)$$

$$\mathcal{L}_{\text{ASP}} = \text{MSE}(x_{\text{ind}}, s_{\text{ind}}). \quad (9)$$

To optimize the \mathcal{L}_{ASP} , encoded action tokens and matched verbs tend to be closer. Note that when the action token is matched with the \emptyset (here we use the overall text representation, i.e. $\emptyset = t_{[\text{CLS}]}$), the predicted action is none.

Overall Objective The final loss function of our proposed STOA-VLP is:

$$\mathcal{L} = \mathcal{L}_{\text{VTC}} + \mathcal{L}_{\text{VTM}} + \mathcal{L}_{\text{MLM}} + \mathcal{L}_{\text{PLM}} + \mathcal{L}_{\text{OTA}} + \mathcal{L}_{\text{ASP}}. \quad (10)$$

By incorporating two spatial-temporal objectives, which are complementary to the overall alignment. Our STOA-VLP learns a better coarse-to-fine cross-modal connection, thus benefiting the downstream tasks that require a more detailed understanding of the video and text.

Experiment

Pre-training Datasets

Instead of pre-training on the commonly used dataset HowTo100M (Miech et al. 2019) with 136M noisy pre-training data which only contains instructional videos, we utilize the WebVid-2M (Bain et al. 2021) to prevent an enormous computation cost. Webvid-2M contains 2.5M web-collected video-text pairs and is an order of magnitude less than the HowTo100M.

Implementation Details

Both the single-modal video and language encoder are 12-layer transformers (Vaswani et al. 2017). We initialize the parameters from (ViT-B/16) (Radford et al. 2021). Similar to (Li et al. 2022), The 6-layer modality-agnostic fusion encoder is initialized with the first 6 layers of the CLIP Text Encoder. The 2-layer object trajectory encoder and spatial-temporal action encoder are initialized from the top 2 layers of the CLIP ViT. We uniformly sample 12 frames from each video. We resize and center crop them into 224x224 to split into patches with size 16x16, getting H=W=14. The maximum length of the text is set to 32. We select K=10 objects per frame. We set the number of object trajectory tokens to 20, and the number of action trajectory tokens is set to 4. Inspired by Cordonnier, Loukas, and Jaggi (2020), the 10th and 11th layer’s patch output of the video encoder is used to generate object trajectories and action features, respectively. All pre-trainings are conducted on 128 Nvidia Tesla V100 GPUs with a batch size of 1024.

Results on Downstream tasks

We evaluate our pre-trained STOA-VLP on three commonly used video-text understanding and generation tasks, i.e. video captioning, text-video retrieval, and video question answering.

Video Captioning We compared our STOA-VLP with recent works including task-specific methods: ORG-TRL (Zhang et al. 2020), SemSynAn (Perez-Martin, Bustos, and Pérez 2021), APML (Lin, Gan, and Wang 2021), VNS-GRU (Chen, Li, and Hu 2020), SAM-SS (Chen et al. 2020), SwinBERT (Lin et al. 2022); and pre-training methods:

	PT data	MSRVTT			LSMDC			DiDeMo		
		R@1	R@5	R@10	R@1	R@5	R@10	R@1	R@5	R@10
Frozen‡	5.5M	32.5	61.5	71.2	15.0	30.8	40.3	31.0	59.8	72.4
ALPRO‡	5.5M	33.9	60.7	73.2	-	-	-	35.9	67.5	78.8
VIOLET	11.5M	34.5	63.0	73.4	16.1	36.6	41.2	32.6	62.8	74.7
HDVILA	103M	35.6	65.3	78.0	17.4	34.1	44.1	28.8	57.4	69.1
OA-Trans‡	5.5M	35.8	63.4	76.5	18.2	34.3	43.7	34.8	64.4	75.1
BridgeFormer‡	5.5M	37.6	64.8	75.1	17.9	35.4	44.5	37.0	62.2	73.9
<i>CLIP initialization</i>										
OA-Trans‡◊	5.5M	40.9	70.4	80.3	-	-	-	-	-	-
CLIP4CLIP	380K	44.5	71.4	81.6	21.6	41.8	49.8	43.4	70.2	80.6
CAMOE*	-	44.6	72.6	81.8	22.5	42.6	50.9	43.8	71.4	-
BridgeFormer‡◊	5.5M	44.9	71.9	80.3	21.8	41.1	50.6	-	-	-
CLIP2Video	-	45.6	72.6	81.7	-	-	-	-	-	-
CLIP2TV*	-	48.3	74.6	82.8	-	-	-	-	-	-
MDMMT-2†	140M	48.5	75.4	83.9	26.9	46.7	55.9	-	-	-
STOA-VLP	2.5M	50.1	75.5	83.8	24.8	46.2	56.0	51.1	76.4	84.0

Table 2: Text-video retrieval comparison results on MSR-VTT, LSMDC, and DiDeMo datasets, *: selecting results without post-processing, e.g. dual softmax (Cheng et al. 2021). †: using additional training corpus. ‡: pre-training with the same Web-Vid2M dataset, plus an image-text dataset CC3M(Sharma et al. 2018). We neglect CC3M because it’s hard to model object trajectories and actions with only one image. ◊: additional results via initializing the model from CLIP.

Clip4Caption (Tang et al. 2021), UniPerceiver (Zhu et al. 2022) and MV-GPT (Seo et al. 2022).

We conduct experiments on the commonly used MSR-VTT and MSVD datasets. Table 1 shows the results of the following four metrics: BLEU-4 (B-4) (Papineni et al. 2002), CIDEr (C) (Vedantam, Zitnick, and Parikh 2015), METEOR (M) (Banerjee and Lavie 2005) and ROUGE-L (R-L) (Lin 2004). Our proposed STOA-VLP surpasses other methods on most of the metrics, even when pre-trained with less video-text data.

	PT data	MSR-VTT	MSVD	MC
		Acc	Acc	Acc
ActBERT	136M	-	-	48.6
NoiseE	136M	35.1	35.1	-
HCRN	-	35.6	36.1	-
ClipBERT	5.6M	37.4	-	88.2
VLM	136M	-	-	91.6
VideoClip	136M	-	-	92.1
DeCEMBERT	136M	37.4	-	-
CoMVT	35M	39.5	42.6	-
HDVILA	103M	40.0	-	97.1
JustAsk	69M	41.5	46.3	-
MV-GPT	53M	41.7	-	-
ALPRO‡	5.5M	42.1	45.9	-
Merlot	180M	43.1	-	90.9
VIOLET	185.5M	43.9	47.9	91.9
STOA-VLP	2.5M	43.2	50.8	98.5

Table 3: Video question answering results on MSR-VTT-QA, MSVD-QA, and MSR-VTT MC datasets. ‡: pre-training with the same Web-Vid2M dataset, plus an image-text dataset CC3M.

Although MV-GPT utilizes subtitles as an additional input, our model still significantly outperforms it on Rouge-L(+4.4) and CIDEr(+0.2) without any additional text input during fine-tuning. The steady improvement demonstrates that STOA-VLP can better utilize both the coarse representation and the fine-grained trajectory information from the video to generate consistent and fluent captions.

Text-Video Retrieval We list Recall at Rank K ($R@K$) results of Frozen (Bain et al. 2021), ALPRO (Li et al. 2022), VIOLET (Fu et al. 2021), HDVILA (Xue et al. 2022), OA-Trans (Wang et al. 2022), BridgeFormer (Ge et al. 2022) and those initialize weights from CLIP (Radford et al. 2021): CLIP4CLIP (Luo et al. 2021), CAMOE (Cheng et al. 2021), CLIP2Video(Fang et al. 2021), CLIP2TV (Gao et al. 2021), MDMMT-2 (Kunitsyn et al. 2022), and STOA-VLP on MSR-VTT-1k-A (Xu et al. 2016), LSMDC (Hendricks et al. 2017) and DiDeMo (Rohrbach et al. 2015).

Table 2 shows the results. Though not explicitly enhancing the overall alignment, our proposed method surpasses most previous methods, except for MDMMT-2 which uses additional training corpora of 10 downstream text-video datasets. We observe that initializing weights from CLIP (Radford et al. 2021) boosts the performance of retrieval. Our STOA-VLP surpasses other pre-training methods initialized from CLIP (Radford et al. 2021), especially the OA-Trans and the BridgeFormer, which combine the Web-Vid2M we used and the CC3M during pre-training. This further demonstrates that dynamically modeling object trajectories and action features can also build a fine-grained connection that benefits the overall video-text alignment.

Video Question Answering We evaluate our model on MSR-VTT-QA (Xu et al. 2017), MSVD-QA (Xu et al. 2017) and MSR-VTT-MC (Yu, Kim, and Kim 2018) datasets with

	Obj.	Act.	OTA	ASP	MSR-VTT Captioning				MSVD Captioning				MSVD QA
					C	M	R	B-4	C	M	R	B-4	Acc
Base					58.3	30.6	68.2	44.8	127.3	42.1	82.7	62.3	49.6
	✓				58.7	30.8	68.2	45.2	128.7	42.3	83.0	62.7	50.1
	✓		✓		59.6	30.7	68.2	45.1	131.7	42.9	83.4	63.0	50.7
	✓	✓	✓		58.9	30.9	68.3	44.9	128.0	43.3	83.5	63.6	50.3
STOA-VLP	✓	✓	✓	✓	60.2	31.0	68.4	45.8	131.8	43.4	83.9	64.4	50.8

Table 4: The ablation of our proposed modules, We add our proposed modules and pre-training tasks in sequence to evaluate their effectiveness on video captioning and video question answering. Note that Obj. means the object trajectory tokens, Act. means the spatial-temporal action tokens.

the prediction accuracy metric. We compare our results with task-specific method HCRN (Le et al. 2020) and pre-training methods: ActBERT (Zhu and Yang 2020), NoisE (Amrani et al. 2021), ClipBERT (Lei et al. 2021), VLM (Xu et al. 2021a), VideoClip (Xu et al. 2021b), DeCeMBERT (Tang, Lei, and Bansal 2021), CoMVT (Seo, Nagrani, and Schmid 2021), HDVILA (Xue et al. 2022), JustAsk (Yang et al. 2021), MV-GPT (Seo et al. 2022), ALPRO (Li et al. 2022), Merlot (Zellers et al. 2021) and VIOLET (Fu et al. 2021).

Table 3 shows the results. With only 2.5M pre-training data, our model surpasses all other methods on MSVD-QA, except for VIOLET which pre-trains on a 185.5M dataset and incorporates DALL-E (Ramesh et al. 2021) pre-trained on 250 million text-image pairs to generate discrete visual features for learning. We observe a performance gain of 2.9% on MSVD-QA, 1.4% on MSR-VTT-MC compared to the previous best models. We conjecture that by explicitly modeling object trajectories and actions, a better alignment is built between the question and the visual feature, and fine-grained information from video is observed and utilized to better answer text questions.




Figure 3: (a): the number of object trajectory tokens and the corresponding retrieval results. (b): the number of action tokens and the corresponding retrieval results. Note that the action token number is set to 4 for (a), and the object trajectory token number is set to 20 for (b).

Ablation Studies

We conduct ablation studies on the MSR-VTT text-video retrieval task and the MSR-VTT, MSVD video captioning task to evaluate the effects of different components.

Effect of Numbers of Spatial-temporal Tokens We study the effect of the number of temporal-aware object

trajectory tokens and spatial-temporal action tokens on the MSR-VTT retrieval task. As shown in Figure 3 (a), we find that the incorporation of the action token does help to align between video and text. Generating 4 action tokens yields the best retrieval performance. Continue to enlarge the action token number, we observe a performance drop. We conjecture this is because some of the action tokens are redundant and failed to be matched with the verb tokens in the text. We found a similar pattern for varying the number of object trajectory tokens in Figure 3 (b).

Effect of Proposed Components We further analyze the effect of our incorporated features and proposed tasks to model them on two different downstream tasks including video captioning and video question answering, which naturally require more fine-grained information and a spatial-temporal understanding of the video scene. The results are shown in Table 4. Note that the base model removes all spatial-temporal modules and auxiliary modeling tasks. Compared to the base model, the incorporation of temporal-aware object trajectories alone can bring improvements to all tasks. Our OTA task further builds a fine-grained alignment between object trajectories and noun tokens that benefits the downstream tasks. We also find that adding a spatial-temporal action modeling module without the ASP task makes some of the captioning results worse. Video captioning and video question answering tasks require a fine-grained semantic understanding of the visual part. We conjecture that without the guidance of our ASP task, the meaning of our action tokens is ambiguous, which might cause a performance drop. Together with our proposed spatial-temporal modules and tasks, our best result is achieved.

Conclusion

In this paper, we explicitly model fine-grained information in a spatial-temporal way to build a better cross-modal alignment. Our proposed STOA-VLP introduces two novel modules to model object trajectories and action features across space and time dimensions. Two auxiliary tasks are designed to build a coarse-to-fine cross-modal alignment. We observe advanced improvements in downstream tasks with moderate-size pre-training data. In addition, the proposed framework can also be applied to more tasks, we regard it as a future working direction.

Acknowledgments

Xiaocheng Feng is the corresponding author of this work. We thank the anonymous reviewers for their insightful comments. This work was supported by the National Key R&D Program of China via grant 2020AAA0106502, National Natural Science Foundation of China (NSFC) via grant 62276078 and the Major Key Project of PCL, PCL2021A06.

References

- Amrani, E.; Ben-Ari, R.; Rotman, D.; and Bronstein, A. 2021. Noise estimation using density estimation for self-supervised multimodal learning. In *Proc. of AAAI*, 6644–6652.
- 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 *Proc. of ICCV*, 1728–1738.
- Banerjee, S.; and Lavie, A. 2005. METEOR: An Automatic Metric for MT Evaluation with Improved Correlation with Human Judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, 65–72.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Proc. of ECCV*, 213–229.
- Chen, H.; Li, J.; and Hu, X. 2020. Delving deeper into the decoder for video captioning. arXiv:2001.05614.
- Chen, H.; Lin, K.; Maye, A.; Li, J.; and Hu, X. 2020. A semantics-assisted video captioning model trained with scheduled sampling. *Frontiers in Robotics and AI*, 475767.
- 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:2109.04290.
- Cordonnier, J.; Loukas, A.; and Jaggi, M. 2020. On the Relationship between Self-Attention and Convolutional Layers. In *Proc. of ICLR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. arXiv:2010.11929.
- Fang, H.; Xiong, P.; Xu, L.; and Chen, Y. 2021. CLIP2Video: Mastering Video-Text Retrieval via Image CLIP. arXiv:2106.11097.
- Fu, T.-J.; Li, L.; Gan, Z.; Lin, K.; Wang, W. Y.; Wang, L.; and Liu, Z. 2021. VIOLET : End-to-End Video-Language Transformers with Masked Visual-token Modeling. arXiv:2111.12681.
- Gao, Z.; Liu, J.; Sun, W.; Chen, S.; Chang, D.; and Zhao, L. 2021. CLIP2TV: Align, Match and Distill for Video-Text Retrieval. arXiv:2111.05610.
- Ge, Y.; Ge, Y.; Liu, X.; Li, D.; Shan, Y.; Qie, X.; and Luo, P. 2022. Bridging Video-Text Retrieval With Multiple Choice Questions. In *Proc. of CVPR*, 16167–16176.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. B. 2017. Mask R-CNN. In *Proc. of ICCV*, 2980–2988.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proc. of CVPR*, 770–778.
- Hendricks, L. A.; Wang, O.; Shechtman, E.; Sivic, J.; Darrell, T.; and Russell, B. C. 2017. Localizing Moments in Video with Natural Language. In *Proc. of ICCV*, 5804–5813.
- Herzig, R.; Ben-Avraham, E.; Mangalam, K.; Bar, A.; Chechik, G.; Rohrbach, A.; Darrell, T.; and Globerson, A. 2022. Object-region video transformers. In *Proc. of CVPR*, 3148–3159.
- Kuhn, H. W. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly*, 83–97.
- Kunitsyn, A.; Kalashnikov, M.; Dzabrayev, M.; and Ivaniuta, A. 2022. MDMMT-2: Multidomain Multimodal Transformer for Video Retrieval, One More Step Towards Generalization. arXiv:2203.07086.
- Le, T. M.; Le, V.; Venkatesh, S.; and Tran, T. 2020. Hierarchical Conditional Relation Networks for Video Question Answering. In *Proc. of CVPR*, 9969–9978.
- 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 *Proc. of CVPR*, 7331–7341.
- Li, D.; Li, J.; Li, H.; Niebles, J. C.; and Hoi, S. C. 2022. Align and Prompt: Video-and-Language Pre-training with Entity Prompts. In *Proc. of CVPR*, 4953–4963.
- Li, J.; Selvaraju, R.; Gotmare, A.; Joty, S.; Xiong, C.; and Hoi, S. C. H. 2021. Align before Fuse: Vision and Language Representation Learning with Momentum Distillation. In *Proc. of NeurIPS*, 9694–9705.
- Li, L.; Chen, Y.-C.; Cheng, Y.; Gan, Z.; Yu, L.; and Liu, J. 2020. HERO: Hierarchical Encoder for Video+Language Omni-representation Pre-training. In *Proc. of EMNLP*, 2046–2065.
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*, 74–81.
- Lin, K.; Gan, Z.; and Wang, L. 2021. Augmented partial mutual learning with frame masking for video captioning. In *Proc. of AAAI*, 2047–2055.
- Lin, K.; Li, L.; Lin, C.-C.; Ahmed, F.; Gan, Z.; Liu, Z.; Lu, Y.; and Wang, L. 2022. SwinBERT: End-to-end transformers with sparse attention for video captioning. In *Proc. of CVPR*, 17949–17958.
- Luo, H.; Ji, L.; Shi, B.; Huang, H.; Duan, N.; Li, T.; Li, J.; Bharti, T.; and Zhou, M. 2020. UniVL: A Unified Video and Language Pre-Training Model for Multimodal Understanding and Generation. arXiv:2002.06353.
- 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:2104.08860.
- Miech, A.; Zhukov, D.; Alayrac, J.; Tapaswi, M.; Laptev, I.; and Sivic, J. 2019. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *Proc. of ICCV*, 2630–2640.

- Pan, B.; Cai, H.; Huang, D.; Lee, K.; Gaidon, A.; Adeli, E.; and Niebles, J. C. 2020. Spatio-Temporal Graph for Video Captioning With Knowledge Distillation. In *Proc. of CVPR*, 10867–10876.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a Method for Automatic Evaluation of Machine Translation. In *Proc. of ACL*, 311–318.
- Perez-Martin, J.; Bustos, B.; and Pérez, J. 2021. Improving video captioning with temporal composition of a visual-syntactic embedding. In *Proc. of WACV*, 3039–3049.
- 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 *Proc. of ICML*, 8748–8763.
- Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *Proc. of ICML*, 8821–8831.
- Ren, S.; He, K.; Girshick, R. B.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Proc. of NeurIPS*, 91–99.
- Rohrbach, A.; Rohrbach, M.; Schiele, B.; et al. 2015. The long-short story of movie description. In *German conference on pattern recognition*, 209–221.
- Seo, P. H.; Nagrani, A.; Arnab, A.; and Schmid, C. 2022. End-to-end generative pretraining for multimodal video captioning. In *Proc. of CVPR*, 17959–17968.
- Seo, P. H.; Nagrani, A.; and Schmid, C. 2021. Look before you speak: Visually contextualized utterances. In *Proc. of CVPR*, 16877–16887.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual Captions: A Cleaned, Hypernymed, Image Alt-text Dataset For Automatic Image Captioning. In *Proc. of ACL*, 2556–2565.
- Sun, C.; Baradel, F.; Murphy, K.; and Schmid, C. 2019a. Learning Video Representations using Contrastive Bidirectional Transformer. arXiv:1906.05743.
- Sun, C.; Myers, A.; Vondrick, C.; Murphy, K.; and Schmid, C. 2019b. VideoBERT: A Joint Model for Video and Language Representation Learning. In *Proc. of ICCV*, 7463–7472.
- Tang, M.; Wang, Z.; Liu, Z.; Rao, F.; Li, D.; and Li, X. 2021. Clip4caption: Clip for video caption. In *Proc. of ACM MM*, 4858–4862.
- Tang, Z.; Lei, J.; and Bansal, M. 2021. DeCEMBERT: Learning from Noisy Instructional Videos via Dense Captions and Entropy Minimization. In *Proc. of NAACL*, 2415–2426.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *Proc. of NeurIPS*, 5998–6008.
- Vedantam, R.; Zitnick, C. L.; and Parikh, D. 2015. CIDEr: Consensus-based image description evaluation. In *Proc. of CVPR*, 4566–4575.
- Wang, J.; Ge, Y.; Cai, G.; Yan, R.; Lin, X.; Shan, Y.; Qie, X.; and Shou, M. Z. 2022. Object-aware Video-language Pre-training for Retrieval. In *Proc. of CVPR*, 3313–3322.
- Wang, Z.; Yu, J.; Yu, A. W.; Dai, Z.; Tsvetkov, Y.; and Cao, Y. 2021. SimVLM: Simple Visual Language Model Pre-training with Weak Supervision. arXiv:2108.10904.
- Xie, S.; Sun, C.; Huang, J.; Tu, Z.; and Murphy, K. 2018. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proc. of ECCV*, 305–321.
- Xu, D.; Zhao, Z.; Xiao, J.; Wu, F.; Zhang, H.; He, X.; and Zhuang, Y. 2017. Video Question Answering via Gradually Refined Attention over Appearance and Motion. In *Proc. of ACM MM*, 1645–1653.
- Xu, H.; Ghosh, G.; Huang, P.-Y.; Arora, P.; Aminzadeh, M.; Feichtenhofer, C.; Metze, F.; and Zettlemoyer, L. 2021a. VLM: Task-agnostic Video-Language Model Pre-training for Video Understanding. In *Proc. of ACL Findings*, 4227–4239.
- Xu, H.; Ghosh, G.; Huang, P.-Y.; Okhonko, D.; Aghajanyan, A.; Metze, F.; Zettlemoyer, L.; and Feichtenhofer, C. 2021b. VideoCLIP: Contrastive Pre-training for Zero-shot Video-Text Understanding. arXiv:2109.14084.
- Xu, J.; Mei, T.; Yao, T.; and Rui, Y. 2016. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *Proc. of CVPR*, 5288–5296.
- Xue, H.; Hang, T.; Zeng, Y.; Sun, Y.; Liu, B.; Yang, H.; Fu, J.; and Guo, B. 2022. Advancing High-Resolution Video-Language Representation with Large-Scale Video Transcriptions. In *Proc. of CVPR*, 5036–5045.
- Yang, A.; Miech, A.; Sivic, J.; Laptev, I.; and Schmid, C. 2021. Just ask: Learning to answer questions from millions of narrated videos. In *Proc. of ICCV*, 1686–1697.
- Yu, Y.; Kim, J.; and Kim, G. 2018. A joint sequence fusion model for video question answering and retrieval. In *Proc. of ECCV*, 471–487.
- Zellers, R.; Lu, X.; Hessel, J.; Yu, Y.; Park, J. S.; Cao, J.; Farhadi, A.; and Choi, Y. 2021. MERLOT: Multimodal Neural Script Knowledge Models. In *Proc. of NeurIPS*, 23634–23651.
- Zhang, P.; Li, X.; Hu, X.; Yang, J.; Zhang, L.; Wang, L.; Choi, Y.; and Gao, J. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proc. of CVPR*, 5579–5588.
- Zhang, Z.; Shi, Y.; Yuan, C.; Li, B.; Wang, P.; Hu, W.; and Zha, Z. 2020. Object Relational Graph With Teacher-Recommended Learning for Video Captioning. In *Proc. of CVPR*, 13275–13285.
- Zhao, B.; Li, X.; and Lu, X. 2018. Video Captioning with Tube Features. In *Proc. of IJCAI*, 1177–1183.
- Zhu, L.; and Yang, Y. 2020. ActBERT: Learning Global-Local Video-Text Representations. In *Proc. of CVPR*, 8743–8752.
- Zhu, X.; Zhu, J.; Li, H.; Wu, X.; Li, H.; Wang, X.; and Dai, J. 2022. Uni-perceiver: Pre-training unified architecture for generic perception for zero-shot and few-shot tasks. In *Proc. of CVPR*, 16804–16815.