

Learning Trajectory-Word Alignments for Video-Language Tasks

Xu Yang¹ Zhangzikang Li^{1,2} Haiyang Xu^{2*} Hanwang Zhang³ Qinghao Ye²
Chenliang Li² Ming Yan² Yu Zhang^{1*} Fei Huang² Songfang Huang²

¹School of Computer Science & Engineering, Key Lab of New Generation Artificial Intelligence Technology
& Its Interdisciplinary Applications (Ministry of Education), Southeast University

²DAMO Academy, Alibaba Group ³School of Computer Science & Engineering, Nanyang Technological University
{xuyang_palm, zhang_yu}@seu.edu.cn, lizhangzikang@gmail.com, {shuofeng.xhy, yeqinghao.yqh,
lc1193798, ym119608, f.huang, songfang.hsfc}@alibaba-inc.com, hanwangzhang@ntu.edu.sg


Abstract

In a video, an object usually appears as the **trajectory**, i.e., it spans over a few spatial but longer temporal patches, that contains abundant spatiotemporal contexts. However, modern Video-Language BERTs (VDL-BERTs) neglect this trajectory characteristic that they usually follow image-language BERTs (IL-BERTs) to deploy the patch-to-word (P2W) attention that may over-exploit trivial spatial contexts and neglect significant temporal contexts. To amend this, we propose a novel **TW-BERT** to learn Trajectory-Word alignment by a newly designed trajectory-to-word (T2W) attention for solving video-language tasks. Moreover, previous VDL-BERTs usually uniformly sample a few frames into the model while different trajectories have diverse graininess, i.e., some trajectories span longer frames and some span shorter, and using a few frames will lose certain useful temporal contexts. However, simply sampling more frames will also make pre-training infeasible due to the largely increased training burdens. To alleviate the problem, during the fine-tuning stage, we insert a novel Hierarchical Frame-Selector (HFS) module into the video encoder. HFS gradually selects the suitable frames conditioned on the text context for the later cross-modal encoder to learn better trajectory-word alignments. By the proposed T2W attention and HFS, our TW-BERT achieves SOTA performances on text-to-video retrieval tasks, and comparable performances on video question-answering tasks with some VDL-BERTs trained on much more data. The code will be available in the supplementary material.

1. Introduction

By witnessing the boom of BERT-like models in single domains, e.g., vision or language[52, 44, 21], researchers

*Corresponding authors.



A boy is playing **basketball**.

Figure 1. The comparisons of five different ways to build object-word alignments. In (b), by patch-to-word (P2W) attention, implicit object-word alignment can be built, e.g., the “ball region” is aligned with **basketball**. While in (c), P2W attention may concentrate the attention of an object on only one frame, e.g., **basketball** only attends over on the “ball region” of the first frame. (d) When pre-training TW-BERT, the trajectory of an object is constructed, e.g., the “ball trajectory” through 4 frames built. Note that we only use 4 frames during pre-training, we show 8 frames for comparing with (e) and the transparent ones denote that they are not input into the model. (e) When fine-tuning TW-BERT, we use HFS to select 4 frames from 8 and we can see that the selected frames are different from the uniformly sampled ones in (d).

begin to build vision-language BERTs[3, 61, 10] for learning robust cross-modal associations. Compared with the images, videos provide more details for building robust associations, e.g., the spatiotemporal contexts in videos can

better describe the actions. However, directly addressing dynamic videos may cost more storage and computation resources. To circumvent such huge cost, researchers first study image-language BERT (IL-BERT)[17, 25, 8, 27, 61] and then exploit the research fruits to build the challenging, while more pragmatic, video-language BERT (VDL-BERT)[47, 2, 26, 36, 13, 35, 46, 37, 54, 51].

Pioneering techniques are directly inherited from IL-BERT into VDL-BERT, *e.g.*, the image encoders are used to embed sampled video frames [23, 62]. However, since image encoders can hardly learn temporal contexts, the resultant VDL-BERTs will degrade to IL-BERT even if they are trained by video-text pairs. To ameliorate it, researchers apply the video encoders [32, 4, 12, 24] to embed spatiotemporal contexts. Though promising improvements are observed, these VDL-BERTs still neglect or ill-consider a significant factor in IL-BERT: the object-word alignment, which helps learn robust associations.

To learn object-word alignments, as Figure 1(a) shows, researchers use an RoI-based extractor to embed an image into a series of RoI features [34, 27, 8]. However, this RoI-based extractor is offline trained by object detection with limited label inventory, which will weaken the IL-BERT since the extractor will not be updated during large-scale pre-training. To enable the end-to-end training, researchers substitute the RoI-based extractor with visual transformers whose outputs are a series of grid embeddings, which will be used to build the cross-modal connections. Although a single grid usually does not construct an integral object, fortunately, the widely applied patch-to-word (P2W) attention of IL-BERT can softly seek the salient visual regions for a given query word. Then the object-word alignments can still be built between this softly detected object and the word, *e.g.*, as shown in Figure 1(b), the object “basketball” can be implicitly attended by the corresponding word query.

Although the P2W attention remedies the loss of RoI-level features for learning object-word alignments in IL-BERT, its effectiveness is weakened in the video case. This is because the objects usually act as the **Trajectories** which span a few spatial while multiple temporal grids in the videos. Thus, directly applying the P2W attention may over-exploit the trivial spatial contexts while neglecting the significant temporal contexts and then make the model attend to only one or two frames. Figure 1(c) shows this limitation that P2W attention only aligns the “ball” in the first frame to the word ball.

To address this limitation, we propose to learn **Trajectory-to-Word** alignments to solve video-language tasks and name this model as **TW-BERT**. Specifically, such alignment is learnt by a novel designed trajectory-to-word **T2W** attention, which first uses the word as the query to seek the salient parts of each frame and the sought parts are sequenced to form the trajectory. Then the query word at-

tends over the trajectories again for capturing cross-modal associations. In this way, the trivial spatial regions are weakened and the temporal contexts will be strengthened, *e.g.*, as shown in Figure 1(d), the attention weights of the word will be concentrated on the object trajectory instead of only one frame as in (c). In the implementation, we follow most VDL-BERTs to set up the network: two single-modal encoders for the video and text and one cross-modal encoder, which is sketched in Figure 2. For the cross-modal encoder, since our T2W attention does not have the same structure as the word-to-patch (W2P) attention, our cross-modal encoder is asymmetric.

Moreover, previous VDL-BERTs usually uniformly sample a few frames into the model, which contains two drawbacks. Firstly, using a few frames may lose temporal context and secondly, uniform sampling can hardly capture the varying graininess of the trajectories, *i.e.*, some trajectories span longer frames and some span shorter. However, simply sampling more frames will largely increase the pre-training burdens that are beyond the computation resources we own. To alleviate this problem, **in the fine-tuning stage**, we sample more frames into the video encoder while only keeping the most relevant frames according to the corresponding text by a novel designed **Hierarchical Frame-Selector (HFS)**. For example, as shown in Figure 1(e), 4 frames are selected by HFS from the 8 uniformly sampled ones for learning trajectory-word alignment. Specifically, HFS inserts a few lightweight layers into the video encoder and these layers can gradually filter frames conditioned on the language context. In this way, HFS learns the coarse-grained trajectory-word alignments, *i.e.*, frame-word alignments, to help the later T2W attention learn more fine-grained trajectory-word alignments.

To sum up, our contributions are:

- We propose a novel perspective to consider the videos that are composed of moving object trajectories, which may inspire the researchers to build more advanced VDL-BERTs.
- We propose a simple while effective **T2W** attention to learn **Trajectory-to-Word** alignments.
- We propose a novel hierarchical frame-selector (**HFS**) in TW-BERT during fine-tuning to capture the varying graininess of the trajectory while not largely increasing the training burdens.
- We achieve SOTA performances compared with other VDL-BERTs trained by the same amount of data.

2. Related Work

2.1. Image-Language BERT (IL-BERT)

Recently, various techniques have been proposed in IL-BERT to learn vision-language connections. Most of them aim at capturing robust object-word alignments since such

alignments construct the foundations of visual reasoning. In the beginning, a pre-trained Faster-RCNN [41] is used to extract a series of RoI poolings, where each one contains one or a few salient objects, to facilitate building object-word alignments [34, 48, 8, 27]. However, this Faster-RCNN is usually pre-trained by the object annotations from COCO [7] and VG [20], whose concept space is much narrower than the data used to train IL-BERT, which can be almost unlimitedly collected from the websites. Moreover, this Faster-RCNN is not updated during the end-to-end training of the IL-BERT, which means that the visual encoder may hardly learn novel knowledge from the web-collected data. Thus the performances of these IL-BERTs are limited.

To further release the potential of hugely web-collected data, the offline Faster RCNN is switched into vision Transformers [11, 32] and thus a homogeneous IL-BERT, *i.e.*, all the components are transformer-based, is built, which is more easily trained end-to-end [34]. Compared with the RoI-based encoder, the vision Transformer outputs a series of patch embeddings and thus may lose object-level knowledge. To remedy such loss, various strategies are proposed to improve the vision-language alignments. For example, the align-before-fuse strategy [25] aligns the paired image-text embeddings before cross-modal fusion for facilitating the subsequent fusion. And the fine-grained contrastive objective [58] amplifies the local details for learning the object-word alignments.

2.2. Video-Language BERT (VDL-BERT)

Since the spatiotemporal contexts in videos can hardly be learnt by image extractors, only inheriting the techniques which are successful in IL-BERT to VDL-BERT is not enough. Thus, based on the fruits of IL-BERT, most VDL-BERTs aim at exploiting more spatiotemporal contexts to build cross-modal associations. One straightforward way is to learn such spatiotemporal contexts through video Transformers [4, 28]. Besides this, some more advanced techniques are proposed, *e.g.*, VIOLET [12] tokenizes the dynamic video patches and predicts the labels of these tokens; BridgeFormer [14] erases the words (nouns or verbs) from the text and learn to match the visual embeddings queried by the erased words and the remained texts; or ALPRO [24] computes the similarities between the video embeddings with the generated entity prompts. TS2-Net [31] shifts token features and selects informative tokens in both temporal and spatial dimensions to produce the fine-grained spatial-temporal video representation. Although substantial improvements are observed, these methods use video patches in the cross-modal encoder, which neglects that an object usually acts as the trajectory in the videos, and thus they may over-exploit the trivial spatial contexts. To ameliorate this limitation, we propose TW-BERT to capture

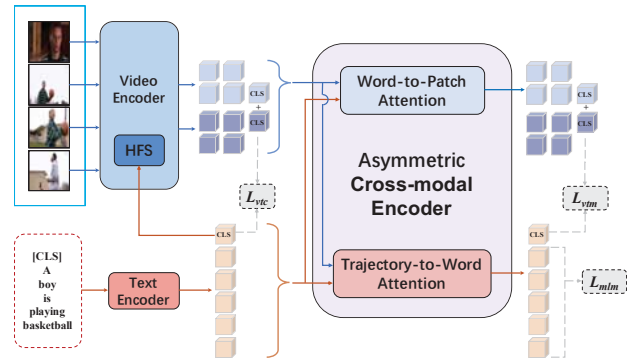


Figure 2. The architecture of TW-BERT, which contains two single-modal encoders and one asymmetric cross-modal encoder. Totally three losses in the grey blocks are used to train the whole model: L_{vtc} , L_{mlm} , and L_{vtm} . Note that HFS is only used in the fine-tuning stage.

trajectory-word alignments for more robust vision-language alignment.

3. TW-BERT

Figure 2 sketches TW-BERT, which has two single-modal encoders for embedding the video and text (cf. Sec. 3.1) and one cross-modal encoder for learning video-language associations (cf. Sec. 3.2). Different from the previous VDL-BERTs, our cross-modal encoder is an asymmetric one that contains a classic word-to-patch (W2P) attention and a novel proposed trajectory-to-word (T2W) attention (cf. Sec. 3.2) for learning trajectory-word alignments. In Sec. 3.3, we introduce how to use Hierarchical Frame-Selector (HFS) in the fine-tuning stage to gradually filter the frames for capturing coarse-grained frame-word alignments. Lastly, we introduce the losses used to train our TW-BERT in Sec. 3.4.

3.1. Single-Modal Encoders

Video Encoder. For a video, we sample a few 224×224 frames and input them into the 12-layer TimeSformer [4, 24] for embedding. TimeSformer first partitions each frame into 14×14 non-overlapping patches, which are flattened and fed to a linear projection layer to produce a sequence of patch tokens. Then TimeSformer applies self-attention along the temporal and spatial dimensions to calculate per-frame features. These features are further mean-pooled along the height and width dimensions. Learnable spatial positional embeddings are added to each video token in the same spatial location of different frames. The final output embedding set is $\mathbf{V} = \{v_{cls}, v_1, \dots, v_{N_V-1}\}$, where $v_n \in \mathbb{R}^d$ and v_{cls} is the global [CLS] embedding.

Text Encoder. For a text, we use a 6-layer transformer to embed it and the output is $\mathbf{X} = \{x_{cls}, x_1, \dots, x_{N_X-1}\}$, where $x_n \in \mathbb{R}^d$ and x_{cls} is the global [CLS] embedding. Similar to the video encoder, we also add positional embed-

dings to the text tokens.

3.2. Asymmetric Cross-Modal Encoder

After embedding videos and texts, a cross-modal encoder is used to fuse them by calculating bi-directional associations: vision-to-language and language-to-vision. No matter what the direction is, the motivation is to assign semantic knowledge from one domain to another. Since a single word contains integral semantic knowledge, we follow previous VDL-BERT [24, 2, 23, 47] to set a traditional word-to-patch (W2P) attention to assign the words to a patch. However, different from the word, only the object instead of one grid conveys integral semantic knowledge. In videos, an object usually spans both spatial and temporal axes and thus the previously used patch-to-word (P2W) attention may fail to transfer the semantic knowledge. To amend this, we design a novel Trajectory-to-Word (T2W) attention for transferring semantic knowledge from videos to texts. Since W2P and T2W attentions have diverse structures, our cross-encoder is asymmetric.

Specifically, both W2P and T2W attentions are built on the Multi-Head Attention (MHA) operation, here we first formalize MHA and then introduce how to use it to build W2P and T2W attentions. Formally, MHA is^{*}:

$$\begin{aligned}
 \text{Input: } & \mathcal{Q}, \mathcal{K}, \mathcal{V} \\
 \text{Att: } & \mathcal{A}_i = \text{Softmax}\left(\frac{\mathcal{Q}\mathcal{W}_i^{\mathcal{Q}}(\mathcal{K}\mathcal{W}_i^{\mathcal{K}})^T}{\sqrt{d}}\right) \\
 \text{Head: } & \mathcal{H}_i = \mathcal{A}_i\mathcal{V}\mathcal{W}_i^{\mathcal{V}}, \\
 \text{Multi-Head: } & \mathcal{H} = [\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_h]\mathcal{W}^{\mathcal{H}}, \\
 \text{Output: } & \mathcal{Z} = \text{LN}(\mathcal{H} + \mathcal{Q}),
 \end{aligned} \tag{1}$$

where $\mathcal{W}_i^{\mathcal{Q}}, \mathcal{W}_i^{\mathcal{K}}, \mathcal{W}_i^{\mathcal{V}}, \mathcal{W}_i^{\mathcal{H}}$ are all trainable matrices; h is the head number and $d_h = d/h$; \mathcal{A}_i is the i -th attention matrix corresponding to the i -th head matrix; $[\cdot]$ is the concatenation operation; and LN is the Layer Normalization.

Word-to-Patch (W2P) Attention. To calculate the W2P alignment, we apply the conventional W2P attention [24]:

$$\mathcal{Z}_{W2P} = \text{MHA}(\mathcal{Q} = \mathcal{V}, \mathcal{K} = \mathcal{V} = \mathcal{X}), \tag{2}$$

where $\mathcal{V} \in \mathbb{R}^{N_v \times D}$, $\mathcal{X} \in \mathbb{R}^{N_x \times D}$ are respectively video and word embedding sets got from two single-modal encoders. By setting the query \mathcal{Q} to the video patch embeddings, Eq. (2) learns to assign suitable words to each video patch and thus captures the W2P alignment.

Trajectory-to-Word (T2W) Attention. As shown in Figure. 3 (b), we propose the T2W attention that uses two steps to learn the T2W alignment: it first constructs a trajectory for a given word and then uses the word as the query to attend over the trajectory for capturing the associations. For




Figure 3. The comparisons between Patch-to-Word (P2W) and Trajectory-to-Word (T2W) attentions, where the green block denotes a query word and the blue blocks denote the video patches. In P2W attention, the word attends over all the video patches and may only concentrate on one frame, while in T2W attention, the salient red parts of each frame are found to construct a trajectory, which are connected by the red lines, *e.g.*, y_1, y_2, y_3 .

convenience, we introduce how the T2W attention calculates the fusion embedding z_{T2W} for a single word x and it is straightforward to extend it to a sequence of the words.

In the first step, T2W attention uses x as the query to find the salient parts for each frame and then sequence these parts to construct the trajectory. Assuming \mathcal{V}_t is the embedding set of the t -th frame, the salient part y_t is got as:

$$y_t = \text{MHA}(\mathcal{Q} = x, \mathcal{K} = \mathcal{V} = \mathcal{V}_t). \tag{3}$$

Then the salient parts at different time frames construct a continuous flow $\mathcal{Y} = \{y_1, \dots, y_T\}$, which is the trajectory of the given word.

In the second step, to get the trajectory-to-word fusion embedding z_{T2W} , we treat x as the query again while using the trajectory \mathcal{Y} as the key and value in MHA:

$$z_{T2W} = \text{MHA}(\mathcal{Q} = x, \mathcal{K} = \mathcal{V} = \mathcal{Y}). \tag{4}$$

By Eq. (3), T2W attention finds the salient parts for the given word at each frame, which enforces Eq. (4) to attend over the **continuous** frames instead of concentrating the attention only on one or some **episodic** frames as the previous P2W attentions. In this way, the whole T2W block exploits more temporal contexts to build vision-language associations, which facilitates the video reasoning tasks that usually require the correct recognition of the temporal patterns. Figure. 3 compares P2W and T2W attentions.

3.3. Hierarchical Frame-Selector

During fine-tuning, to improve the performance, previous studies tend to sample more frames than pre-training for introducing more temporal information. However, on the one hand, using too many frames will dramatically increase the computational resources, which causes infeasible training. On the other hand, in a video, different trajectories have varying graininess, *e.g.*, some trajectories may cover longer frames and some cover shorter. In this way, uniform sampling may introduce redundant temporal information for these longer trajectories and deficient information for the shorter ones. Interestingly, the time duration of the trajectory is also reflected in the corresponding text. In this way, we can use the text to choose suitable frames and then

^{*}To avoid symbol confusions, we use the calligraphic font ($\mathcal{}$ in LaTeX) to denote the built-in variables of the MHA module.




Figure 4. The architecture of the proposed Hierarchical Frame-Selector. At beginning, four frames are input into the video encoder and the first frame-selector behind the 6-th layer filters out the 2-nd frame. At last, only the tokens of the 1-st and 4-th frames are preserved.

only input them into the cross-modal attention for learning trajectory-word alignments.

Motivated by this, we propose a novel **Hierarchical Frame-Selector (HFS)** that can gradually select the most relevant video frames to reduce the computational burdens while providing more temporal knowledge. Figure 4 illustrates the architecture of our proposed frame-selector module. HFS contains a few frame selector layers which have the same structure with different parameters. For each frame-selector layer, given a sequence of video frame [CLS] tokens $\mathbf{V}_{cls} = \{v_{cls}^1, v_{cls}^2, \dots, v_{cls}^T\} \in \mathbb{R}^{T \times D}$, we concatenate the text [CLS] embedding \mathbf{x}_{cls} with each v_{cls}^t : $\mathbf{f}_t = [v_{cls}^t, \mathbf{x}_{cls}] \in \mathbb{R}^{2D}$ ($1 \leq t \leq T$) and feed them into a scorer network including a FC layer and a softmax layer:

$$\mathbf{s} = \text{Softmax}(FC(\mathbf{f})) \in \mathbb{R}^{T \times 1}. \quad (5)$$

Here, T denotes the number of frames input into the video encoder, and $\mathbf{s} \in \mathbb{R}^T$ is the score vector of T frames. Then for this frame selector, we keep the tokens of k frames with the top- k s_t for the further embedding in the subsequent video Transformer layers. During training, we employ the perturbed maximum method [5] to construct a differentiable Top-K operator. In the video encoder, we totally insert 2 selector layers and insert them in the 6-th and 12-th layers as in Figure 4.

3.4. Training Objectives

To train TW-BERT, as the grey blocks shown in Figure 2, we totally use three losses which are masked language modeling (MLM), video-text matching (VTM), and video-text contrastive loss (VTC).

Masked language Modeling (MLM) [26, 24, 12, 50]. MLM aims to predict the masked word tokens given both the video and the text contexts. To get it, we first randomly replace the input text tokens with the [MASK] token with a probability of 15% and then use the [MASK] embedding output from the cross-modal encoder to predict the masked word by calculating a cross-entropy loss:

$$L_{MLM} = \mathbb{E}_{(V, \hat{X}) \sim D} H(y^{msk}, p^{msk}(V, \hat{X})) \quad (6)$$

where \hat{X} is the masked text, H is the cross-entropy loss, y^{msk}/p^{msk} are the ground-truth/predicted masked tokens.

Video-text Matching (VTM) [24, 12, 50]. VTM calculates whether the given video and text are matched or not. To get it, for a given video-text pair, we first randomly replace the

text with the ones from a different video in the same batch. Then we concatenate the video and text [CLS] embeddings output from the cross-modal encoder and input the concatenated embedding into a binary classifier to judge whether the given video-text pair is matched or not:

$$L_{VTM} = \mathbb{E}_{(V, X) \sim D} H(y^{vtm}, p^{vtm}(V, X)) \quad (7)$$

where $y^{vtm}/p^{vtm}(V, X)$ are the ground-truth/predicted values indexing whether V and X are matched or not.

Video-text Contrastive (VTC) [24, 50, 49, 14, 15, 2, 29]. As detailed in Section 3, VTC contrasts the outputs of two single-modal encoders to pull close their embedding space to help the subsequent cross-modal encoder build more robust vision-language associations. Suppose s_{ij} denotes the similarity score of the i -th video and the j -th text, for each video and text, we calculate the softmax-normalized video-to-text and video-to-image similarity as:

$$p_{ij}^{v2t}(V) = \frac{\exp(s_{ij}/\tau)}{\sum_j \exp(s_{ij}/\tau)}, p_{ij}^{t2v}(X) = \frac{\exp(s_{ij}/\tau)}{\sum_i \exp(s_{ij}/\tau)} \quad (8)$$

where τ is a learnable temperature parameter. Let $y^{v2t}(V)$ and $y^{t2v}(X)$ denote the ground-truth one-hot similarity. This loss contains two symmetric parts, where the left term forces the i -th text embedding to be close to the i -th video embedding compared with the other texts and the right term has a similar effect:

$$L_{VTC} = \frac{1}{2} \mathbb{E}_{(V, X) \sim D} [H(y^{v2t}(V), p^{v2t}(V)) + H(y^{t2v}(X), p^{t2v}(X))] \quad (9)$$

In the implementation, we follow [25] to use the momentum queue as a continuously-evolving teacher to provide more negative samples.

4. Experiments

4.1. Pre-training Dataset

Following recent work [24, 14, 15, 49, 2], we pre-train TW-BERT on Google Conceptual Captions (CC3M) [45] containing 3.3M image-text pairs and WebVid-2M [2] containing 2.5M video-text pairs. For CC3M, the image is treated as a one-frame video data during pre-training. Note that due to the limited storage and computation resources, we do not use some much larger datasets like HowTo100M [38] containing 136M video-text pairs as [26,

55]. Also, we do not distill knowledge from CLIP [39], which is pre-trained on 400M image-text pairs, as [36].

4.2. Downstream Tasks

Text-to-Video Retrieval. (i) **MSRVTT** contains 10K YouTube videos with 200K descriptions. We follow [56] to use 9K train+val videos for training and report results on the 1K test split. (ii) **DiDeMo** [16] contains 10K Flickr videos annotated with 40K sentences. (iii) **LSMDC** consists of 118,081 video clips sourced from 202 movies, where the validation set and the test set contain 7,408 and 1,000 videos. (iv) **ActivityNet Caption** contains 20K YouTube videos annotated with 100K sentences. The training set contains 10K videos, and we use val1 set with 4.9K videos to report results. For MSRVTT and LSMDC, we perform standard text-to-video retrieval. For DiDeMo and ActivityNet Caption, we concatenate all the text captions in the same video as a single query and evaluate paragraph-to-video retrieval. For two tasks, we measure the performances by average recall at K(R@K) and Median Rank on zero-shot and fine-tune setups.

Video Question Answering. (i) **MSRVTT-QA** [53] is built upon videos and captions from MSRVTT [56], which contains 10K videos with 243K open-ended questions and 1.5K answer candidates. (ii) **MSVD** [6] contains 50K question-answer pairs with 2423 answer candidates. We use standard train/val/test splits for the two tasks, and report accuracy.

4.3. Implementation Details

We initialize our video encoder by ViT-B/16 [42] and the text encoder by the first six BERT-Base layers [9]. For the cross-modal encoder, the self-attentions in all 3 cross-modal attentions (W2P contains 1 and T2W contains 2) are initialized by the last 6 BERT-Base layers [9]. The model is trained end-to-end during both pre-training and fine-tuning.

In pre-training, the feature dimension is set to 256 when calculating the contrastive losses and the temperature is set to 0.05. For the momentum queue, the momentum value is 0.995 and the size of the queue is 65,536. The above implementation details follow the recent work [25, 24] for a fair comparison. We pre-train the model on CC3M and WebVid-2M for 10 epochs on 8 NVIDIA A100 GPUs where the batch size is 128. We use AdamW [18] optimizer with a weight decay of 0.001 and betas (0.9, 0.98). The learning rate is first warmed-up to $1e-4$ and then decays following a linear decay schedule.

During fine-tuning text-to-video retrieval, We sample 20 frames and use **HFS** to preserve 8 frames to feed into the cross-modal encoder. The model is trained with both VTC and VTM losses, and we obtain similarity scores from the output of the VTM head during inference. For video question answering, we sample 32 frames per video and preserve 16 frames. Since MSRVTT-QA and MSVD-QA [53]

are open-ended VQA, in which the answers are in free-form natural language, it is common to convert the task to a classification task by predicting the answer’s label. We input the concatenation of the video and question [CLS] tokens into a two-layer MLP [9] for calculating the cross-entropy loss. All the fine-tuning experiments are conducted on 8 NVIDIA V100 GPUs.

4.4. Ablation Studies

We conduct comprehensive ablation studies to evaluate the effectiveness of the proposed trajectory-to-word (T2W) attention and Hierarchical Frame-Selector (HFS) module.

Comparing Methods. Base: We use a symmetric cross-modal encoder that contains patch-to-word (P2W) and word-to-patch (W2P) attentions. **T2W:** We replace P2W attention in Base by our T2W attention. **MeanP:** Another simple way to consider temporal knowledge is to mean pool the embeddings along the temporal axis. Specifically, we simply average $\{y_1, \dots, y_T\}$. These three ablations are implemented using the zero-shot setting.

[L1,L2]: We insert frame-selector layer into the L1-th and L2-th layers. **F@T:** We sample T frames into the model. **F@T1-T2-T3:** We use Hierarchical Frame Selector where “F@T1-T2-T3” denotes that T1 frames are originally input into the model and then T2 and T3 frames are respectively remained by the first and second frame-selector layers. All these ablations are implemented using the fine-tuning setting.

Quantitative Results. Table 1 compares the performances of diverse baselines. From this table, we can see that T2W outperforms Base, e.g., T2W achieves 3.6% R@5 improvements on MSRVTT for zero-shot evaluation, which suggests that our T2W attention can exploit more temporal contexts to better solve video-language tasks.

Also, T2W achieves higher scores than MeanP on different tasks, e.g., T2W achieves 28.4% of R@1 score in DiDeMo zero-shot text-to-video retrieval task while MeanP only has 27.4%. Applying mean-pooling strategy indeed fuse all the temporal knowledge, however, simply averaging all y_t also means that the contexts *over the whole time axis* are used, while lots of objects may only *span a few frames* and using all the temporal contexts may introduce trivial or even harmful noises. This proves that our T2W attention can capture more important cross-modal associations in this asymmetric cross-modal encoder case.

We further analyze the impact of the layer settings and frame number on the Hierarchical Frame-Selector module. As shown in Table 2, frame-selector in deep layers (e.g. 6-th and 12-th layers) brings better performances on both retrieval and QA tasks. Moving frame-selector module into shallower layers deteriorates the scores. For example, the accuracy drops 0.9% when the frame-selector module is placed before the third layer. We hypothesize that discard-

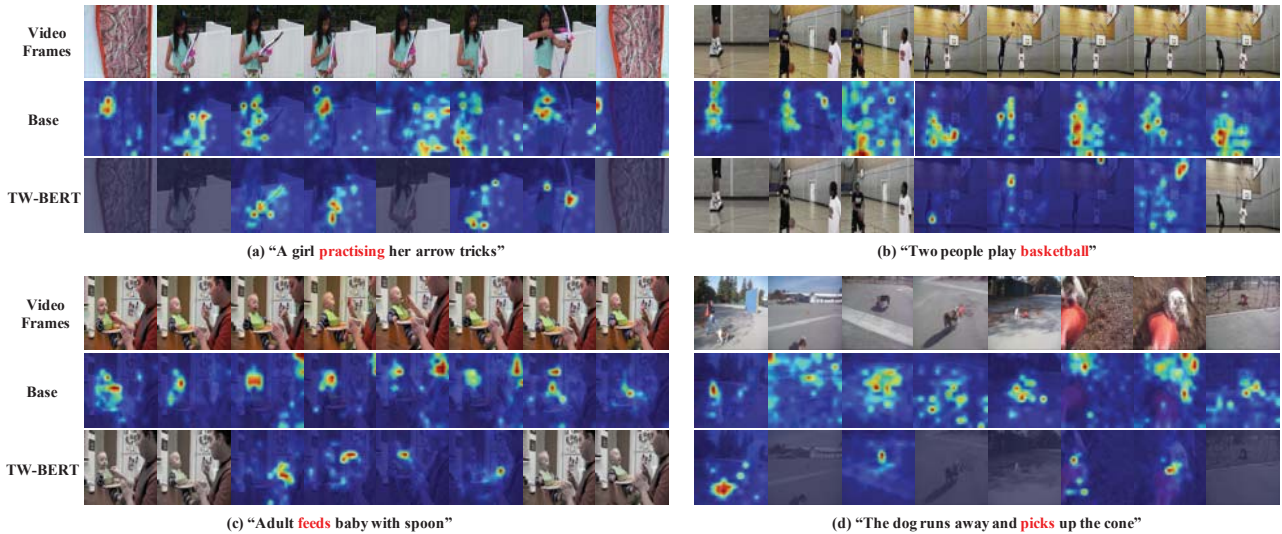


Figure 5. Visualizations of the attention maps from cross-modal encoder. Sample (a) and (b) are from MSRVT [56], (c) and (d) are from DiDeMo [16] retrieval dataset. **TW-BERT** attends to the patches related to given query word by Trajectory-to-Word attention.

Table 1. Performances of various baselines. R@K and MedR respectively denote recall (%) with K retrieval efforts and median ranking for retrieved videos where higher R@K and lower MedR indicate better performance.

Method	MSRVTT-ZS				DiDeMo-ZS				MSVD-QA
	R@1↑	R@5↑	R@10↑	MedR↓	R@1↑	R@5↑	R@10↑	MedR↓	Acc.
Base	25.1	46.4	57.3	7.0	26.6	52.8	62.7	5.0	47.4
T2W	26.8	50.0	59.9	5.0	28.4	52.9	64.5	4.0	48.2
MeanP	25.8	47.9	58.0	6.0	27.4	53.1	64.0	5.0	48.0

Table 2. Performances comparison with different layer settings of the **HFS** in fine-tuning.

layers	MSRVTT-FT				MSVD-QA	
	R@1↑	R@5↑	R@10↑	MedR↓	Acc.	
[2,4]	36.7	63.2	74.9	3.0	46.2	
[3,6]	37.4	63.7	75.4	3.0	46.9	
[4,8]	38.1	64.6	76.2	3.0	47.8	
[5,10]	38.2	64.9	76.3	3.0	48.3	
[6,12]	38.4	65.1	76.6	3.0	48.5	

Table 3. Performances comparison with different frame settings of the **HFS** in fine-tuning.

Frame	MSRVTT-FT				Frame	MSVD-QA	
	R@1↑	R@5↑	R@10↑	MedR↓		Acc.	
F@1	35.8	63.3	74.2	3.0	F@8	47.2	
F@4-2-1	35.9	63.2	73.9	3.0	F@20-14-8	47.2	
F@4	37.1	63.9	75.3	3.0	F@12	47.7	
F@8-6-4	37.5	64.2	75.5	3.0	F@24-18-12	47.9	
F@8	38.1	64.9	76.0	3.0	F@16	48.2	
F@20-14-8	38.4	65.1	76.6	3.0	F@32-24-16	48.5	

ing frames too early would result in a lack of temporal knowledge in subsequent ViT layers. We thus choose to insert the frame-selector module in the 6-th and 12-th layers in the following experiments. In terms of frame number, as shown in Table 3, on the one hand, as the number of frame we sample increases, both retrieval and QA performance obtain improvements. For example, in MSRVT

fine-tuning text-to-video retrieval task, F@20-14-8 outperforms F@8-6-4 by 0.9% on R@1 or on MSVD video question answering task, F@32-24-16 outperforms F@24-18-12 by 0.6%, which suggests that more temporal knowledge benefit our model whether using uniform sampling or our proposed **HFS**. On the other hand, using **HFS** gets higher scores than uniform sampling, *e.g.*, F@20-14-8 achieves 38.4% of R@1 score in MSRVT retrieval task while F@8 only has 38.1%, or on MSVD video question answering task, F@32-24-16 outperforms F@16 by 0.3%, which validates the power of **HFS** module.

Qualitative Results. We show the heat maps of the attentions of **Base** and **TW-BERT** in Figure 5. We see that **TW-BERT** can select the most relevant frames according to the whole text, then it implicitly form a trajectory among these frames for a given query word to avoid over-exploiting the trivial spatial contexts as in **Base**. For example, in (a), **TW-BERT** first discards the 1-st and 8-th frames which is completely irrelevant to the text, then tracks the hand of the girl in each frame that it keeps according to the query word “practising” while **Base** cannot avoid impact of the irrelevant frames, and in those relevant frames, it only attends to the larger while trivial regions about the whole body of the girl. Moreover, in (d), **TW-BERT** keeps the 6-th and 7-th frames for these two frames show the dog’s action more

Table 4. Experiments of text-to-video retrieval on MSRVT, DiDeMo and LSMDC datasets. “#PT Pairs” lists the number of video-text pairs for pre-training. We show results with zero-shot evaluation (top) and fine-tuning evaluation (bottom). R@A is the average of R@1, R@5 and R@10.

Method	#PT Pairs	DiDeMo					LSMDC					MSRVTT				
		R@1↑	R@5↑	R@10↑	R@A↑	MedR↓	R@1↑	R@5↑	R@10↑	R@A↑	MedR↓	R@1↑	R@5↑	R@10↑	R@A↑	MedR↓
Frozen [2]	5.5M	21.1	46.0	56.2	41.1	7.0	9.3	22.0	30.1	20.5	51.0	18.7	39.5	51.6	36.6	10.0
VIOLET [12]	186M	23.5	49.8	59.8	44.4	-	-	-	-	-	-	25.9	49.5	59.7	45.0	-
OA-Trans [49]	5.5M	23.5	50.4	59.8	44.6	6.0	-	-	-	-	-	23.4	47.5	55.6	42.2	8.0
ALPRO [24]	5.5M	23.8	47.3	57.9	43.0	6.0	-	-	-	-	-	24.1	44.7	55.4	41.4	8.0
BridgeFormer [14]	5.5M	25.6	50.6	61.1	45.8	5.0	12.2	25.9	32.2	23.4	42.0	26.0	46.4	56.4	42.9	7.0
MILES [15]	5.5M	27.2	50.3	63.6	47.0	5.0	11.1	24.7	30.6	22.1	50.7	26.1	47.2	56.9	43.4	7.0
TW-BERT	5.5M	28.4	52.9	64.5	48.6	4.0	14.2	30.4	36.0	26.9	28.0	26.8	50.0	59.9	45.6	5.0
Frozen [2]	5.5M	31.0	59.8	72.4	54.4	3.0	15.0	30.8	39.8	28.5	20.0	31.0	59.5	70.5	53.7	3.0
VIOLET [12]	186M	32.6	62.8	74.7	56.7	-	16.1	36.6	41.2	31.3	-	34.5	63.0	73.4	57.0	-
ALPRO [24]	5.5M	35.9	67.5	78.8	60.7	3.0	-	-	-	-	-	33.9	60.7	73.2	55.9	3.0
OA-Trans [49]	5.5M	34.8	64.4	75.1	58.1	3.0	18.2	34.3	43.7	32.1	18.5	35.8	63.4	76.5	58.6	3.0
BridgeFormer [14]	5.5M	37.0	62.2	73.9	57.7	3.0	17.9	35.4	44.5	32.6	15.0	37.6	64.8	75.1	59.2	3.0
MILES [15]	5.5M	36.6	63.9	74.0	58.2	3.0	17.8	35.6	44.1	32.5	15.5	37.7	63.6	73.8	58.4	3.0
TW-BERT	5.5M	41.8	71.1	81.2	64.7	2.0	21.0	38.8	49.2	36.3	11.0	38.4	65.1	76.6	60.0	3.0

Table 5. ActivityNet Caption with fine-tuning setting.

Method	#PT Pairs	R@1↑	R@5↑	R@10↑	MedR↓
Dense [19]	-	14.0	32.0	-	34.0
FSE [60]	-	18.2	44.8	-	7.0
CE [30]	-	18.2	47.7	-	6.0
HSE [60]	-	20.5	49.3	-	-
Clipbert [23]	5.6M	21.3	49.0	63.5	6.0
TW-BERT	5.5M	31.7	62.3	74.9	3.0

Table 6. Experiments of video question answering on MSRVT and MSVD datasets in top-1 accuracy (%).

Method	#PT Pairs	MSRVTT	MSVD
Clipbert [23]	5.6M	37.4	-
ALPRO [24]	5.5M	42.1	45.9
SINGULARITY [22]	5.5M	42.7	45.9
LGDN [33]	15.2M	43.1	-
SSML [1]	100M	35.1	35.1
JustAsk [57]	69M	41.5	46.3
MERLOT [59]	180M	43.1	-
VIOLET [12]	186M	43.9	47.9
TW-BERT	5.5M	43.6	48.5

clearly. Then TW-BERT attends to the part where the dog is in contact with the cone according to the query word “picks” while Base only focuses on the body of the dog.

4.5. Comparisons with SOTA

We compare **TW-BERT** with previous methods on two frequently applied tasks: video-text retrieval (VDTR) and video question answering (VDQA). Table 4 and 5 report the performances of VDTR on MSRVT [56], DiDeMo [16], LSMDC [43], and ActivityNet Caption [19], respectively, where the former three datasets contain both zero-shot and fine-tuning setups and the last one only has fine-tuning setup. Table 6 reports the VDQA on MSRVT [56] and MSVD [6]. Among the compared models, MILES [15], BridgeFormer [14], OA-Trans [49], Clipbert [23] and VIOLET [12] are SOTA models proposed in recently 1-2 years. Note that VIOLET [12] and ALPRO [24] distill knowledge from additional large-scale BERTs while we do not. Also, we show the number of pre-training video-text pairs in these

tables for more clear comparisons.

From these tables, we can find that when the pre-training data is in the same scale, TW-BERT achieves the best performance compared with all the other models on both VDTR and VDQA. For example, on DiDeMo VDTR, TW-BERT outperforms BridgeFormer by 4.8% on R@1, or on MSVD VDQA, TW-BERT outperforms ALPRO by 2.6%. Moreover, compared with the models trained on much more data, TW-BERT can still achieve the best performances on various tasks, *e.g.*, on LSMDC VDTR, TW-BERT outperforms VIOLET by 8.0% on R@10 or on MSRVT VDQA, TW-BERT outperforms MERLOT [59] by 0.5%.

Note that the videos in LSMDC are longer than MSRVT and DiDeMo, which means that the videos in this dataset contain more temporal contexts than the other datasets. Then as shown in Table 4, the improvements of TW-BERT over other SOTAs are larger than the improvements on the other datasets. For example, compared with BridgeFormer, TW-BERT has an average 3.5% improvement in the zero-shot setting, while on DiDeMo dataset, the corresponding average improvement over MILES is only 1.6%. Such comparisons further validate the effectiveness of TW-BERT in exploiting temporal contexts.

Among these SOTAs, only when compared with VIOLET, which uses an additional large-scale model DALL-E [40] and 32 more times pre-training data than ours (180M VS. 5.5M), TW-BERT cannot comprehensively surpass VIOLET on all the tasks. For example, on MSRVT VDQA, VIOLET achieves 0.3% higher than TW-BERT, while such marginal improvements are got at the cost of much more training resources. Furthermore, TW-BERT still outperforms VIOLET on the other tasks, *e.g.*, on DiDeMo VDTR, TW-BERT outperforms VIOLET by 9.2% on R@1 or on MSVD VDQA, TW-BERT outperforms VIOLET by 0.6%. These comparisons confirm the effectiveness of the proposed TW-BERT.

5. Conclusion

We propose a novel Trajectory-Word BERT (TW-BERT) that builds Trajectory-to-Word alignments for solving video-language tasks. In particular, we introduce an asymmetric cross-modal encoder that contains word-to-patch (W2P) and Trajectory-to-Word (T2W) to capture cross-modal associations. Moreover, a novel Hierarchical Frame-Selector (HFS) module is proposed in the fine-tuning stage to filter out irrelevant frames to help the later T2W attention learn better trajectory-word alignments. Extensive experiments across diverse tasks confirm the effectiveness of the proposed TW-BERT.

6. Acknowledgements

This work is supported by National Science Foundation of China (62206048), National Key R&D Program of China (2018AAA0100104, 2018AAA0100100), Natural Science Foundation of Jiangsu Province (BK20220819, BK20211164), and Young Elite Scientists Sponsorship Program of Jiangsu Association for Science and Technology Tj-2022-027.

References

- [1] Elad Amrani, Rami Ben-Ari, Daniel Rotman, and Alex Bronstein. Noise estimation using density estimation for self-supervised multimodal learning. In *AAAI*, pages 6644–6652, 2021. 8
- [2] 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*, page 1728–1738, 2021. 2, 4, 5, 8
- [3] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 1
- [4] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding. *arXiv preprint arXiv:2102.05095*, 2021. 2, 3
- [5] Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. Learning with differentiable perturbed optimizers. *Advances in neural information processing systems*, 33:9508–9519, 2020. 5
- [6] David Chen and William B Dolan. Collecting highly parallel data for paraphrase evaluation. *meeting of the association for computational linguistics*, 2011. 6, 8
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, , and C. Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 3
- [8] 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 *ECCV*, 2020. 2, 3
- [9] 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. 6
- [10] Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Fang Wen, and Nenghai Yu. Peco: Perceptual codebook for bert pre-training of vision transformers. *arXiv preprint arXiv:2111.12710*, 2021. 1
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 3
- [12] 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, 3, 5, 8
- [13] Valentin Gabeur, Chen Sun, Kartek Alahari, and Cordelia Schmid. Multi-modal transformer for video retrieval. In *ECCV*, pages 214–229, 2020. 2
- [14] Yuying Ge, Yixiao Ge, Xihui Liu, Dian Li, Ying Shan, Xiaohu Qie, and Ping Luo. Bridgeformer: Bridging video-text retrieval with multiple choice questions. In *CVPR*, pages 16167–16176, 2022. 3, 5, 8
- [15] Yuying Ge, Yixiao Ge, Xihui Liu, Alex Jinpeng Wang, Jianping Wu, Ying Shan, Xiaohu Qie, and Ping Luo. Miles: Visual bert pre-training with injected language semantics for video-text retrieval. *arXiv preprint arXiv:2204.12408*, 2022. 5, 8
- [16] Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *ICCV*, page 5804–5813, 2017. 6, 7, 8
- [17] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, pages 4904–4916, 2021. 2
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [19] Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *ICCV*, page 706–715, 2017. 8
- [20] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, and David A Shamma. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 2017. 3
- [21] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019. 1
- [22] Jie Lei, Tamara L. Berg, and Mohit Bansal. Revealing single frame bias for video-and-language learning. *arXiv preprint arXiv:2206.03428*, 2022. 8

- [23] 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 *CVPR*, page 7331–7341, 2021. 2, 4, 8
- [24] Dongxu Li, Junnan Li, Hongdong Li, Juan Carlos Niebles, and Steven C.H. Hoi. Align and prompt: Video-and-language pre-training with entity prompts. In *CVPR*, page 4953–4963, 2022. 2, 3, 4, 5, 6, 8
- [25] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven C. H. Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In *Advances in neural information processing systems*, volume 34, pages 9694–9705, 2021. 2, 3, 5, 6
- [26] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. Hero: Hierarchical spatio-temporal reasoning with contrastive action correspondence for end-to-end video object grounding. *arXiv preprint arXiv:2005.00200*, 2020. 2, 5, 6
- [27] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020. 2, 3
- [28] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mvitv2: Improved multiscale vision transformers for classification and detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4804–4814, 2022. 3
- [29] Chengzhi Lin, Ancong Wu, Junwei Liang, Jun Zhang, Wenhong Ge, Wei-Shi Zheng, and Chunhua Shen. Text-adaptive multiple visual prototype matching for video-text retrieval. *Advances in neural information processing systems*, 35:38655–38666, 2022. 5
- [30] Yang Liu, Samuel Albanie, Arsha Nagrani, and Andrew Zisserman. Use what you have: Video retrieval using representations from collaborative experts. *arXiv preprint arXiv:1907.13487*, 2019. 8
- [31] Yuqi Liu, Pengfei Xiong, Luhui Xu, Shengming Cao, and Qin Jin. Ts2-net: Token shift and selection transformer for text-video retrieval. In *ECCV*, pages 319–335, 2022. 3
- [32] 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 *ICCV*, pages 10012–10022, 2021. 2, 3
- [33] Haoyu Lu, Mingyu Ding, Nanyi Fei, Yuqi Huo, and Zhiwu Lu. Lgdn: Language-guided denoising network for video-language modeling. *arXiv preprint arXiv:2209.11388*, 2022. 8
- [34] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in neural information processing systems*, 2019. 2, 3
- [35] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Xilin Chen, and Ming Zhou. Univlm: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2002.06353*, 2020. 2
- [36] 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. *arXiv preprint arXiv:2104.08860*, 2021. 2, 6
- [37] 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 *CVPR*, pages 9879–9889, 2020. 2
- [38] 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 *CVPR*, page 9879–9889, 2020. 5
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, and Jack Clark. Learning transferable visual models from natural language supervisions. In *ICML*, pages 8748–8763, 2021. 6
- [40] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092*, 2021. 8
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2015. 3
- [42] Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihni Zelnik-Manor. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*, 2021. 6
- [43] Anna Rohrbach, Marcus Rohrbach, Niket Tandon, and Bernt Schiele. A dataset for movie description. In *CVPR*, page 3202–3212, 2015. 8
- [44] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019. 1
- [45] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL*, page 2556–2565, 2018. 5
- [46] Chen Sun, Fabien Baradel, Kevin Murphy, and Cordelia Schmid. Learning video representations using contrastive bidirectional transformer. *arXiv preprint arXiv:1906.05743*, 2019. 2
- [47] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *ICCV*, page 7464–7473, 2019. 2, 4
- [48] Hao Tan and Mohit Bansal. Lxmert: Learning cross-modality encoder representations from transformers. *arXiv preprint arXiv:1908.07490*, 2019. 3
- [49] Alexander Wang, Yixiao Ge, Guanyu Cai, Rui Yan, Xudong Lin, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Object-aware video-language pre-training for retrieval. In *CVPR*, pages 3313–3322, 2022. 5, 8

- [50] Alex Jinpeng Wang, Yixiao Ge, Rui Yan, Yuying Ge, Xudong Lin, Guanyu Cai, Jianping Wu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. All in one: Exploring unified video-language pre-training. *arXiv preprint arXiv:2203.07303*, 2022. 5
- [51] Wenzhe Wang, Mengdan Zhang, Runnan Chen, Guanyu Cai, Penghao Zhou, Pai Peng, Xiaowei Guo, Jian Wu, and Xing Sun. Dig into multi-modal cues for video retrieval with hierarchical alignment. In *IJCAI*, pages 1113–1121, 2021. 2
- [52] Yau-Shian Wang, Hung yi Lee, and Yun-Nung Chen. Tree transformer: Integrating tree structures into self-attention. *arXiv preprint arXiv:1909.06639*, 2019. 1
- [53] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *ACM International Conference on Multimedia*, pages 1645–1653, 2017. 6
- [54] Hu Xu, Gargi Ghosh, Po-Yao Huang, Prahal Arora, Masoumeh Aminzadeh, Christoph Feichtenhofer, Florian Metze, and Luke Zettlemoyer. Vlm: Task-agnostic video-language model pre-training for video understanding. *arXiv preprint arXiv:2105.09996*, 2021. 2
- [55] Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *arXiv preprint arXiv:2109.14084*, 2021. 6
- [56] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msrvt: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296, 2016. 6, 7, 8
- [57] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *ICCV*, pages 1686–1697, 2021. 8
- [58] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. *arXiv preprint arXiv:2111.07783*, 2021. 3
- [59] Rowan Zellers, Ximing Lu, Jack Hessel, Youngjae Yu, Jae Sung Park, Jize Cao, Ali Farhadi, and Yejin Choi. Merlot: Multimodal neural script knowledge models. In *Advances in neural information processing systems*, volume 34, page 23634–23651, 2021. 8
- [60] Bowen Zhang, Hexiang Hu, and Fei Sha. Cross-modal and hierarchical modeling of video and text. In *ECCV*, pages 374–390, 2018. 8
- [61] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan L. Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. 1, 2
- [62] Linchao Zhu and Yi Yang. Actbert: Learning global-local video-text representations. In *CVPR*, page 8746–8755, 2020. 2