

# Expanding Language-Image Pretrained Models for General Video Recognition

Bolin Ni<sup>1,2,\*</sup>, Houwen Peng<sup>4,†</sup>, Minghao Chen<sup>5,\*</sup>, Songyang Zhang<sup>6</sup>,  
Gaofeng Meng<sup>1,2,3,†</sup>, Jianlong Fu<sup>4</sup>, Shiming Xiang<sup>1,2</sup>, Haibin Ling<sup>5</sup>

<sup>1</sup> NLPR, Institute of Automation, Chinese Academy of Sciences

<sup>2</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences

<sup>3</sup> CAIR, HK Institute of Science and Innovation, Chinese Academy of Sciences

<sup>4</sup> Microsoft Research    <sup>5</sup> Stony Brook University    <sup>6</sup> University of Rochester

**Abstract.** Contrastive language-image pretraining has shown great success in learning visual-textual joint representation from web-scale data, demonstrating remarkable “zero-shot” generalization ability for various image tasks. However, how to effectively expand such new language-image pretraining methods to video domains is still an open problem. In this work, we present a simple yet effective approach that adapts the pretrained language-image models to video recognition directly, instead of pretraining a new model from scratch. More concretely, to capture the long-range dependencies of frames along the temporal dimension, we propose a cross-frame attention mechanism that explicitly exchanges information across frames. Such module is lightweight and can be plugged into pretrained language-image models seamlessly. Moreover, we propose a video-specific prompting scheme, which leverages video content information for generating discriminative textual prompts. Extensive experiments demonstrate that our approach is effective and can be generalized to different video recognition scenarios. In particular, under fully-supervised settings, our approach achieves a top-1 accuracy of 87.1% on Kinetics-400, while using 12× fewer FLOPs compared with Swin-L and ViViT-H. In zero-shot experiments, our approach surpasses the current state-of-the-art methods by +7.6% and +14.9% in terms of top-1 accuracy under two popular protocols. In few-shot scenarios, our approach outperforms previous best methods by +32.1% and +23.1% when the labeled data is extremely limited. Code and models are available at [here](#).

**Keywords:** Video Recognition, Contrastive Language-Image Pretraining

## 1 Introduction

Video recognition is one of the most fundamental yet challenging tasks in video understanding. It plays a vital role in numerous vision applications, such as micro-video recommendation [62], sports video analysis [40], autonomous driving [18],

---

\* Work done during internship at Microsoft Research.

† Corresponding authors: houwen.peng@microsoft.com, gfmeng@nlpr.ia.ac.cn.

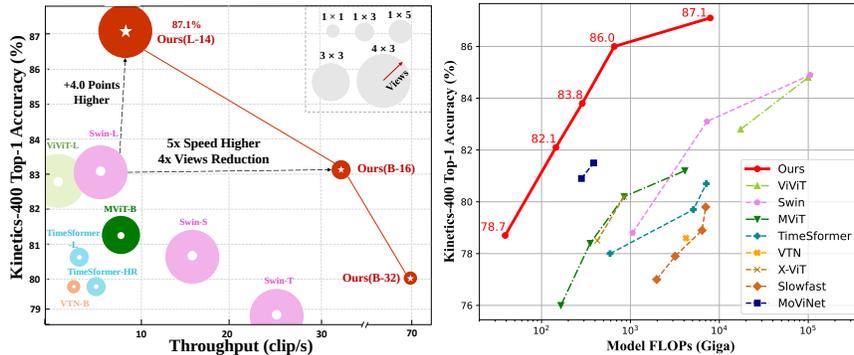


Fig. 1: Comparison with state-of-the-art methods on Kinetics-400 [22] in terms of throughput, the number of views, and FLOPs. Best viewed in color.

and so on. Over the past few years, based upon convolutional neural networks and now transformers, video recognition has achieved remarkable progress [62,21]. Most existing works follow a closed-set learning setting, where all the categories are pre-defined. Such method is unrealistic for many real-world applications, such as automatic tagging of web videos, where information regarding new video categories is not available during training. It is thus very challenging for closed-set methods to train a classifier for recognizing unseen or unfamiliar categories.

Fortunately, recent work in large-scale contrastive language-image pretraining, such as CLIP [36], ALIGN [19], and Florence [54], has shown great potentials in addressing this challenge. The core idea is to learn visual or visual-language representation with natural language supervision using web-scale image-text data. After pretraining, natural language is used to reference learned visual concepts (or describe new ones), thus enabling zero/few-shot transfer of the models to downstream tasks. Inspired by these works [36,19,54], we consider to use text as the supervision signals to learn a new video representation for general recognition scenarios, including zero-shot, few-shot, and fully-supervised.

However, directly training a language-video model is unaffordable for many of us, because it requires large-scale video-text pretraining data as well as a massive number of GPU resources (*e.g.*, thousands of GPU days). A feasible solution is to adapt the pretrained language-image models to video domain. Very recently, there are several studies exploring how to transfer the knowledge from the pretrained language-image models to other downstream tasks, *e.g.*, point cloud understanding [58] and dense prediction [37,59]. However, the transfer and adaptation to video recognition is not well explored. When adapting the pretrained cross-modality models from image to video domain, there are two key issues to be solved: 1) how to leverage the temporal information contained in videos, and 2) how to acquire discriminative text representation for a video.

For the first question, we present a new architecture for video temporal modeling. It consists of two key components: a cross-frame communication transformer and a multi-frame integration transformer. Specifically, the cross-

frame communication transformer takes raw frames as input and provides a frame-level representation using a pretrained language-image model, while allowing information exchange between frames with message tokens. Each message token not only depicts the semantics of the current frame, but also communicates with other frames to model their dependencies. The multi-frame integration transformer then simply transfer the frame-level representations to video-level.

For the second question, we employ the text encoder pretrained in the language-image models and expand it with a video-specific prompting scheme. The key idea is to leverage video content information to enhance text prompting. The intuition behind is that appropriate contextual information can help the recognition. For example, if there is extra video content information about “**in the water**”, the actions “**swimming**” and “**running**” will be much easier to be distinguished. In contrast to prior work manually designing a fixed set of text prompts, this work proposes a learnable prompting mechanism, which integrates both semantic labels and representation of videos for automatic prompt generation.

With the above two issues addressed, we can smoothly adapt the existing image-level cross-modality pretrained models to video domains. Without loss of generality, here we choose the available CLIP [36] and Florence [54] models and eXpand them for general video recognition, forming new model families called X-CLIP and X-Florence, respectively. Comprehensive experiments demonstrate our expanded models are generally effective. In particular, under the fully-supervised setting, X-CLIP-L/14 achieves competitive performance on Kinetics-400/600 with a top-1 accuracy of 87.1%/88.3%, surpassing ViViT-H [3] by 2.3%/2.5% while using  $12\times$  fewer FLOPs, as shown in Fig. 1. In zero-shot experiments, X-Florence surpasses the state-of-the-art ActionCLIP [48] by +7.6% and +14.9% under two popular protocols. In few-shot experiments, X-CLIP outperforms other prevailing methods by +32.1% and +23.1% when the data is extremely limited.

In summary, our contributions are three-fold:

- We propose a new cross-frame communication attention for video temporal modeling. This module is light and efficient, and can be seamlessly plugged into existing language-image pretrained models, without undermining their original parameters and performance.
- We design a video-specific prompting technique to yield instance-level textual representation automatically. It leverages video content information to enhance the textual prompt generation.
- Our work might pave a new way of expanding existing large-scale language-image pretrained models for general video recognition and other potential video tasks. Extensive experiments demonstrate the superiority and good generalization ability of our method under various learning configurations.

## 2 Related Work

**Visual-language Pretraining.** Visual-language pretraining has achieved remarkable progress over the past few years [43,42,31,61]. In particular, contrastive language-image pretraining demonstrates very impressive “zero-shot” transfer

and generalization capacities [36,19,54]. One of the most representative works is the recent CLIP [36]. A large amount of follow-up works have been proposed to leverage the pretrained models for downstream tasks. For example, CoOp [60], CLIP-Adapter [15] and Tip-Adapter [57] use the pretrained CLIP for improving the few-shot transfer, while PointCLIP [58] and DenseCLIP [37,59] transfer the knowledge to point cloud understanding and dense prediction, respectively. VideoCLIP [51] extends the image-level pretraining to video by substituting the image-text data with video-text pairs [31]. However, such video-text pretraining is computationally expensive and requires a large amount of curated video-text data which is not easy to acquire. In contrast, our method directly adapts the existing pretrained model to video recognition, largely saving the training cost.

There are two concurrent works mostly related to ours. One is ActionCLIP [48], while the other is [20]. Both of them introduce visual-language pretrained models to video understanding. ActionCLIP proposes a “pretrain, prompt and finetune” framework for action recognition, while [20] proposes to optimize a few random vectors for adapting CLIP to various video understanding tasks. In contrast, our method is more general. It supports adapting various language-image models, such as CLIP and Florence [54], from image to video. Moreover, we propose a lightweight and efficient cross-frame attention module for video temporal modeling, while presenting a new video-specific text prompting scheme.

**Video Recognition.** One key factor to build a robust video recognition model is to exploit the temporal information. Among many methods, 3D convolution is widely used [44,45,35,50], while it suffers from high computational cost. For efficiency purposes, some studies [45,35,50] factorize convolutions across spatial and temporal dimensions, while others insert the specific temporal modules into 2D CNNs [27,25,30]. Nevertheless, the limited receptive field of CNNs gives the rise of transformer-based methods [3,5,29,11,53], which achieve very promising performance recently. However, these transformer-based methods are either computationally intensive or insufficient in exploiting the temporal information. For example, ViViT [3] disregards the temporal information in the early stage. Video Swin [29] utilizes 3D attention while having high computational cost.

The temporal modeling scheme in our method shares a similar spirit with the recent proposed video transformers, *i.e.*, VTN [32], ViViT [3], and AVT [17]. They all use a frame-level encoder followed by a temporal encoder, but our method has two fundamental differences. 1) In [32,3,17], each frame is encoded separately, resulting in no temporal interaction before final aggregation. This late fusion strategy does not fully make use of the temporal cues. By contrast, our method replaces the spatial attention with the proposed cross-frame attention, which allows global spatio-temporal modeling for all frames. 2) Similar to previous works [29,11,12,5], both ViViT [3] and VTN [32] adopt a dense temporal sampling strategy and ensemble the predictions of multiple views at inference, which is time-consuming. On the contrary, we empirically analyze different sampling methods for late fusion, and demonstrate that a sparse sampling is good enough, achieving better performance with fewer FLOPs than the dense strategy, as verified in Sec. 4.5 (Analysis).

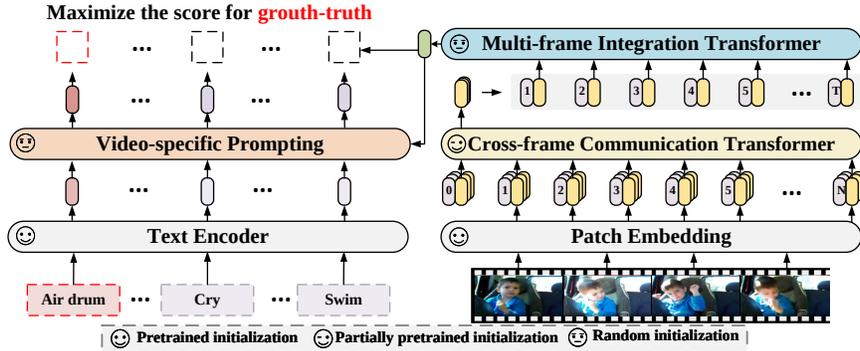


Fig. 2: An overview of our framework. The details are elaborated in Sec. 3.1.

### 3 Approach

In this section, we present our proposed framework in detail. First, we briefly overview our video-text framework in Sec. 3.1. Then, we depict the architecture of the video encoder, especially for the proposed cross-frame attention in Sec. 3.2. Finally, we introduce a video-specific prompting scheme in Sec. 3.3.

#### 3.1 Overview

Most prior works in video recognition learn discriminative feature embeddings supervised by a one-hot label [3,5,12,47]. While in this work, inspired by the recent contrastive language-image pretraining [36,19,54], we propose to use text as the supervision, since the text provides more semantic information. As shown in Fig. 2, our method learns to align the video representation and its corresponding text representation by jointly training a video encoder and a text encoder. Rather than pretraining a new video-text model from scratch, our method is built upon prior language-image models and expands them with video temporal modeling and video-adaptive textual prompts. Such a strategy allows us to fully take advantage of existing large-scale pretrained models while transferring their powerful generalizability from image to video in a seamless fashion.

Formally, given a video clip  $V \in \mathcal{V}$  and a text description  $C \in \mathcal{C}$ , where  $\mathcal{V}$  is a set of videos and  $\mathcal{C}$  is a collection of category names, we feed the video  $V$  into the video encoder  $f_{\theta_v}$  and the text  $C$  into the text encoder  $f_{\theta_c}$  to obtain a video representation  $\mathbf{v}$  and a text representation  $\mathbf{c}$  respectively, where

$$\mathbf{v} = f_{\theta_v}(V), \quad \mathbf{c} = f_{\theta_c}(C). \quad (1)$$

Then, a video-specific prompt generator  $f_{\theta_p}$  is employed to yield instance-level textual representation for each video. It takes the video representation  $\mathbf{v}$  and text representation  $\mathbf{c}$  as inputs, formulated as

$$\hat{\mathbf{c}} = f_{\theta_p}(\mathbf{c}, \mathbf{v}). \quad (2)$$

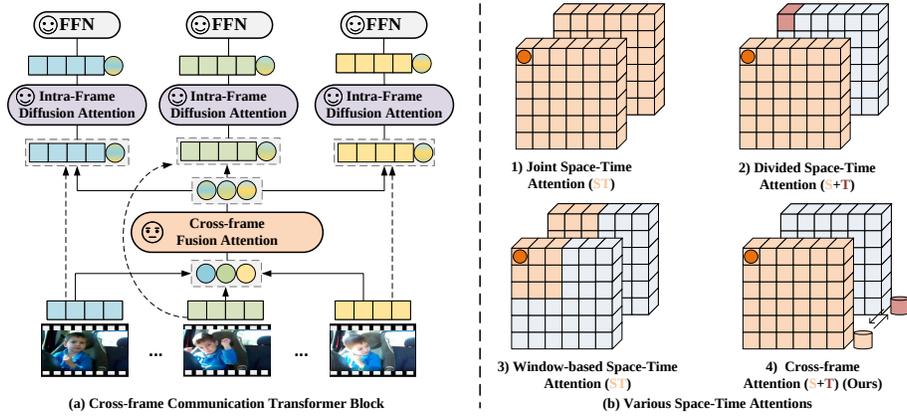


Fig. 3: (a) Cross-frame Attention. (b) compares different space-time attention mechanisms used in existing video transformer backbones [3,5,29].

Finally, a cosine similarity function  $\text{sim}(\mathbf{v}, \hat{\mathbf{c}})$  is utilized to compute the similarity between the visual and textual representations:

$$\text{sim}(\mathbf{v}, \hat{\mathbf{c}}) = \langle \mathbf{v}, \hat{\mathbf{c}} \rangle / (\|\mathbf{v}\| \|\hat{\mathbf{c}}\|). \quad (3)$$

The goal of our method is to maximize the  $\text{sim}(\mathbf{v}, \hat{\mathbf{c}})$  if  $V$  and  $C$  are matched and otherwise minimize it.

### 3.2 Video Encoder

Our proposed video encoder is composed of two cascaded vision transformers: a cross-frame communication transformer and a multi-frame integration transformer. The cross-frame transformer takes raw frames as input and provides a frame-level representation using a pretrained language-image model, while allowing information exchange between frames. The multi-frame integration transformer then simply integrates the frame-level representations and outputs video features.

Specifically, given a video clip  $V \in \mathbb{R}^{T \times H \times W \times 3}$  of  $T$  sampled frames with  $H$  and  $W$  denote the spatial resolution, following ViT [10], the  $t$ -th frame is divided into  $N$  non-overlapping patches  $\{\mathbf{x}_{t,i}\}_{i=1}^N \in \mathbb{R}^{P^2 \times 3}$  with each of size  $P \times P$  pixels, where  $t \in \{1, \dots, T\}$  denotes the temporal index, and  $N=HW/P^2$ . The patches  $\{\mathbf{x}_{t,i}\}_{i=1}^N$  are then embedded into patch embeddings using a linear projection  $\mathbf{E} \in \mathbb{R}^{3P^2 \times D}$ . After that, we prepend a learnable embedding  $\mathbf{x}_{class}$  to the sequence of embedded patches, called [class] token. Its state at the output of the encoder serves as the frame representation. The input of the cross-frame communication transformer at the frame  $t$  is denoted as:

$$\mathbf{z}_t^{(0)} = [\mathbf{x}_{class}, \mathbf{E}\mathbf{x}_{t,1}, \mathbf{E}\mathbf{x}_{t,2}, \dots, \mathbf{E}\mathbf{x}_{t,N}] + \mathbf{e}^{spa}, \quad (4)$$

where  $\mathbf{e}^{spa}$  represents the spatial position encoding.

Then we feed the patch embeddings into an  $L_c$ -layer Cross-frame Communication Transformer (CCT) to obtain the frame-level representation  $\mathbf{h}_t$ :

$$\begin{aligned} \mathbf{z}_t^{(l)} &= \text{CCT}^{(l)}(\mathbf{z}_t^{(l-1)}), \quad l = 1, \dots, L_c \\ \mathbf{h}_t &= \mathbf{z}_{t,0}^{(L_c)}, \end{aligned} \quad (5)$$

where  $l$  denotes the block index in CCT,  $\mathbf{z}_{t,0}^{(L_c)}$  represents the final output of the [class] token. CCT is built-up with the proposed cross-frame attention, as will be elaborated later.

At last, the  $L_m$ -layer Multi-frame Integration Transformer (MIT) takes all frame representation  $\mathbf{H} = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_T]$  as input and outputs the video-level representation  $\mathbf{v}$  as following:

$$\mathbf{v} = \text{AvgPool}(\text{MIT}(\mathbf{H} + \mathbf{e}^{temp})), \quad (6)$$

where AvgPool and  $\mathbf{e}^{temp}$  denote the average pooling and temporal position encoding, respectively. We use standard learnable absolute position embeddings [46] for  $\mathbf{e}^{spa}$  and  $\mathbf{e}^{temp}$ . The multi-frame integration transformer is constructed by the standard multi-head self-attention and feed-forward networks [46].

**Cross-frame Attention.** To enable a cross-frame information exchange, we propose a new attention module. It consists of two types of attentions, *i.e.*, cross-frame fusion attention (CFA) and intra-frame diffusion attention (IFA), with a feed-forward network (FFN). We introduce a *message token* mechanism for each frame to abstract, send and receive information, thus enabling visual information to exchange across frames, as shown in Fig. 3(a). In detail, the message token  $\mathbf{m}_t^{(l)}$  for the  $t$ -th frame at the  $l$ -th layer is obtained by employing a linear transformation on the [class] token  $\mathbf{z}_{t,0}^{(l-1)}$ . This allows message tokens to abstract the visual information of the current frame.

Then, the cross-frame fusion attention (CFA) involves all message tokens to learn the global spatio-temporal dependencies of the input video. Mathematically, this process at  $l$ -th block can be expressed as:

$$\hat{\mathbf{M}}^{(l)} = \mathbf{M}^{(l)} + \text{CFA}(\text{LN}(\mathbf{M}^{(l)})), \quad (7)$$

where  $\hat{\mathbf{M}}^{(l)} = [\hat{\mathbf{m}}_1^{(l)}, \hat{\mathbf{m}}_2^{(l)}, \dots, \hat{\mathbf{m}}_T^{(l)}]$  and LN indicates layer normalization [4].

Next, the intra-frame diffusion (IFA) takes the frame tokens with the associated message token to learn visual representation, while the involved message token could also diffuse global spatio-temporal dependencies for learning. Mathematically, this process at  $l$ -th block can be formulated as:

$$[\hat{\mathbf{z}}_t^{(l)}, \bar{\mathbf{m}}_t^{(l)}] = [\mathbf{z}_t^{(l-1)}, \hat{\mathbf{m}}_t^{(l)}] + \text{IFA}(\text{LN}([\mathbf{z}_t^{(l-1)}, \hat{\mathbf{m}}_t^{(l)}])), \quad (8)$$

where  $[\cdot, \cdot]$  concatenates the features of frame tokens and message tokens.

Finally, the feed-forward network(FFN) performs on the frame tokens as:

$$\mathbf{z}_t^{(l)} = \hat{\mathbf{z}}_t^{(l)} + \text{FFN}(\text{LN}(\hat{\mathbf{z}}_t^{(l)})). \quad (9)$$

Note that the message token is dropped before the FFN layer and does not pass through the next block, since it is generated online and used for frames communication within each block. Alternating the fusion and diffusion attentions through  $L_c$  blocks, the cross-frame communication transformer (CCT) can encode the global spatial and temporal information of video frames. Compared to other space-time attention mechanisms [3,5,29], as presented in Fig. 3(b), our proposed cross-frame attention models the global spatio-temporal information while greatly reducing the computational cost.

*Initialization.* When adapting the pretrained image encoder to the video encoder, there are two key modifications. 1) The intra-frame diffusion attention (IFA) inherits the weights directly from the pretrained models, while the cross-frame fusion attention (CFA) is randomly initialized. 2) The multi-frame integration transformer is appended to the pretrained models with random initialization.

### 3.3 Text Encoder

We employ the pretrained text encoder and expand it with a video-specific prompting scheme. The key idea is to use video content to enhance the text representation. Given a description  $C$  about a video, the text representation  $\mathbf{c}$  is obtained by the text encoder, where  $\mathbf{c} = f_{\theta_c}(C)$ . For video recognition, how to generate a good text description  $C$  for each video is a challenging problem. Previous work, such as CLIP [36], usually defines textual prompts manually, such as “A photo of a {label}”. However, in this work, we empirically show that such manually-designed prompts do not improve the performance for video recognition (as presented in Tab. 9). In contrast, we just use the “{label}” as the text description  $C$  and then propose a learnable text prompting scheme.

**Video-specific prompting.** When understanding an image or a video, human can instinctively seek helps from discriminative visual cues. For example, the extra video semantic information of “in the water” will make it easier to distinguish “swimming” from “running”. However, it is difficult to acquire such visual semantics in video recognition tasks, because 1) the datasets only provide the category names, such as “swimming” and “running”, which are pre-defined and fixed; and 2) the videos in the same class share the identical category name, but their visual context and content are different. To address these issues, we propose a learnable prompting scheme to generate textual representation automatically. Concretely, we design a video-specific prompting module, which takes the video content representation  $\bar{\mathbf{z}}$  and text representation  $\mathbf{c}$  as inputs. Each block in the video-specific prompting module is consisting of a multi-head self-attention (MHSA) [46] followed by a feed-forward network to learn the prompts,

$$\bar{\mathbf{c}} = \mathbf{c} + \text{MHSA}(\mathbf{c}, \bar{\mathbf{z}}) \text{ and } \tilde{\mathbf{c}} = \bar{\mathbf{c}} + \text{FFN}(\bar{\mathbf{c}}), \quad (10)$$

where  $\mathbf{c}$  is the text embedding,  $\bar{\mathbf{z}} \in \mathbb{R}^{N \times d}$  is the average of  $\{\mathbf{z}_t^{(L_c)}\}_{t=1}^T$  along the temporal dimension, and  $\tilde{\mathbf{c}}$  is the video-specific prompts. We use text representation  $\mathbf{c}$  as query and the video content representation  $\bar{\mathbf{z}}$  as key and value. This

Table 1: Comparison with state-of-the-art on Kinetics-400. We report the FLOPs and throughput per view. \* indicates video-text pretraining.

Method	Pretrain	Frames	Top-1	Top-5	Views	FLOPs(G)	Throughput
<i>Methods with random initialization</i>							
MViTv1-B, 64×3 [11]	-	64	81.2	95.1	3 × 3	455	7
<i>Methods with ImageNet pretraining</i>							
Uniformer-B [24]	IN-1k	32	83.0	95.4	4 × 3	259	-
TimeSformer-L [5]	IN-21k	96	80.7	94.7	1 × 3	2380	3
Mformer-HR [33]	IN-21k	16	81.1	95.2	10 × 3	959	-
Swin-L [29]	IN-21k	32	83.1	95.9	4 × 3	604	6
Swin-L (384↑) [29]	IN-21k	32	84.9	96.7	10 × 5	2107	-
MViTv2-L (312↑) [26]	IN-21k	40	86.1	97.0	5 × 3	2828	-
<i>Methods with web-scale image pretraining</i>							
ViViT-H/16x2 [3]	JFT-300M	32	84.8	95.8	4 × 3	8316	-
TokenLearner-L/10 [39]	JFT-300M	-	85.4	96.3	4 × 3	4076	-
CoVeR [55]	JFT-3B	-	<b>87.2</b>	-	1 × 3	-	-
<i>Methods with web-scale language-image pretraining</i>							
ActionCLIP-B/16 [48]	CLIP-400M	32	83.8	96.2	10 × 3	563	-
A6 [20]	CLIP-400M	16	76.9	93.5	-	-	-
MTV-H [53]	WTS*	32	<b>89.1</b>	98.2	4 × 3	3705	-
X-Florence (384↑)	FLD-900M	8	86.2	96.6	4 × 3	2114	6
X-Florence	FLD-900M	32	86.5	96.9	4 × 3	2822	2
X-CLIP-B/16	IN-21k	8	81.1	94.7	4 × 3	145	33
X-CLIP-B/32		8	80.4	95.0	4 × 3	39	136
X-CLIP-B/32		16	81.1	95.5	4 × 3	75	69
X-CLIP-B/16	CLIP-400M	8	83.8	96.7	4 × 3	145	33
X-CLIP-B/16		16	84.7	96.8	4 × 3	287	17
X-CLIP-L/14		8	87.1	97.6	4 × 3	658	8
X-CLIP-L/14 (336↑)		16	<b>87.7</b>	97.4	4 × 3	3086	2

implementation allow the text representation to extract the related visual context from videos. We then enhance the text embedding  $\mathbf{c}$  with the video-specific prompts  $\hat{\mathbf{c}}$  as follows,  $\hat{\mathbf{c}} = \mathbf{c} + \alpha\tilde{\mathbf{c}}$ , where  $\alpha$  is a learnable parameter with an initial value of 0.1. The  $\hat{\mathbf{c}}$  is finally used for classification in Eq. (3).

## 4 Experiments

In this section, we conduct experiments on different settings, *i.e.*, fully-supervised, zero-shot and few-shot, followed by the ablation studies of the proposed method.

### 4.1 Experimental Setup

**Architectures and Datasets.** We expand CLIP and Florence to derive four variants: X-CLIP-B/32, X-CLIP-B/16, X-CLIP-L/14 and X-Florence, respectively. X-CLIP-B/32 adopts ViT-B/32 as parts of the cross-frame communication transformer, X-CLIP-B/16 uses ViT-B/16, while X-CLIP-L/14 employs ViT-L/14. For all X-CLIP variants, we use a simple 1-layer multi-frame integration transformer, and the number of the video-specific prompting blocks is 2. We

Table 2: Comparison with state-of-the-art on Kinetics-600.

Method	Pretrain	Frames	Top-1	Top-5	Views	FLOPs	Throughput
<i>Methods with random initialization</i>							
MViT-B-24, 32×3 [11]	-	32	83.8	96.3	5 × 1	236	-
<i>Methods with ImageNet pretraining</i>							
Swin-L (384↑) [29]	IN-21k	32	86.1	97.3	10 × 5	2107	-
<i>Methods with web-scale pretraining</i>							
ViViT-L/16x2 320 [3]	JFT-300M	32	83.0	95.7	4 × 3	3992	-
ViViT-H/16x2 [3]	JFT-300M	32	85.8	96.5	4 × 3	8316	-
TokenLearner-L/10 [39]	JFT-300M	-	86.3	97.0	4 × 3	4076	-
Florence (384↑) [54]	FLD-900M	-	87.8	97.8	4 × 3	-	-
CoVeR [55]	JFT-3B	-	<b>87.9</b>	-	1 × 3	-	-
MTV-H [53]	WTS*	32	<b>89.6</b>	98.3	4 × 3	3705	-
X-CLIP-B/16		8	85.3	97.1	4 × 3	145	74
X-CLIP-B/16	CLIP-400M	16	85.8	97.3	4 × 3	287	40
X-CLIP-L/14		8	<b>88.3</b>	97.7	4 × 3	658	20

evaluate the efficacy of our method on four benchmarks: *Kinetics-400&600* [22,7], *UCF-101* [41] and *HMDB-51* [23]. More details about architectures and datasets are provided in the *supplementary materials*.

## 4.2 Fully-supervised Experiments

**Training and Inference.** We sample 8 or 16 frames in fully-supervised experiments. The detailed hyperparameters are showed in the *supplementary materials*.

**Results.** In Tab. 1, we report the results on Kinetics-400 and compare with other SoTA methods under different pretraining, including random initialization, IN-1k/21k [9] pretraining, web-scale image and language-image pretraining.

Compared to the methods pretrained on IN-21k [9], our X-CLIP-B/16<sub>8f</sub> (8 frames) surpasses Swin-L [28] by +0.7% with 4× fewer FLOPs and running 5× faster (as presented in Fig. 1). The underlying reason is that the shift-window attention in Swin is inefficient. Also, our X-CLIP-L/14<sub>8f</sub> outperforms MViTv2-L [26] by +1.0% with 5× fewer FLOPs. In addition, when using IN-21k pretraining, our method surpasses TimeSformer-L [5] with fewer FLOPs.

When compared to the methods using web-scale image pretraining, our X-CLIP is also competitive. For example, X-CLIP-L/14<sub>8f</sub> achieves +2.3% higher accuracy than ViViT-H [3] with 12× fewer FLOPs. MTV-H [53] achieves better results than ours, but it uses much more pretraining data. Specifically, MTV-H uses a 70M video-text dataset including about 17B images, which are much larger than the 400M image-text data used in CLIP pretraining.

Moreover, compared to ActionCLIP [48], which also adopt CLIP as the pretrained model, our X-CLIP-L/14<sub>8f</sub> is still superior, getting +3.3% higher accuracy with fewer FLOPs. There are two factors leading to the smaller FLOPs of our method. One is that X-CLIP does not use 3D attention like [29] and has fewer layers. The other factor is that X-CLIP samples fewer frames for each video clip, such as 8 or 16 frames, while ActionCLIP [48] using 32 frames.

Table 3: Zero-shot performances on HMDB51 [23] and UCF101 [41].

Method	HMDB-51	UCF-101
MTE [52]	19.7 ± 1.6	15.8 ± 1.3
ASR [49]	21.8 ± 0.9	24.4 ± 1.0
ZSECOC [34]	22.6 ± 1.2	15.1 ± 1.7
UR [63]	24.4 ± 1.6	17.5 ± 1.6
TS-GCN [14]	23.2 ± 3.0	34.2 ± 3.1
E2E [6]	32.7	48
ER-ZSAR [8]	35.3 ± 4.6	51.8 ± 2.9
ActionCLIP [48]	40.8 ± 5.4	58.3 ± 3.4
X-CLIP-B/16	<b>44.6 ± 5.2</b> (+3.8)	<b>72.0 ± 2.3</b> (+13.7)
X-Florence	<b>48.4 ± 4.9</b> (+7.6)	<b>73.2 ± 4.2</b> (+14.9)

Table 4: Zero-shot performance on Kinetics-600 [7].

Method	Top-1 Acc.	Top-5 Acc.
DEWISE [13]	23.8 ± 0.3	51.0 ± 0.6
ALE [1]	23.4 ± 0.8	50.3 ± 1.4
SJE [2]	22.3 ± 0.6	48.2 ± 0.4
ESZSL [38]	22.9 ± 1.2	48.3 ± 0.8
DEM [56]	23.6 ± 0.7	49.5 ± 0.4
GCN [16]	22.3 ± 0.6	49.7 ± 0.6
ER-ZSAR [8]	42.1 ± 1.4	73.1 ± 0.3
X-CLIP-B/16	<b>65.2 ± 0.4</b> (+23.1)	<b>86.1 ± 0.8</b> (+13.0)
X-Florence	<b>68.8 ± 0.9</b> (+26.7)	<b>88.4 ± 0.6</b> (+15.3)

In addition, we report the results on Kinetics-600 in Tab. 2. Using only 8 frames, our X-CLIP-B/16<sub>8f</sub> surpasses ViViT-L, while using 27× fewer FLOPs. More importantly, our X-CLIP-L/14<sub>8f</sub> achieves 88.3% top-1 accuracy while using 5× fewer FLOPs compared to the current state-of-the-art method MTV-H [53].

From the above fully-supervised experiments, we can observe that, our X-CLIP method achieves very competitive performance compared to prevailing video transformer models [55,54,53,48,20]. This mainly attributes to two factors. 1) The proposed cross-frame attention can effectively model temporal dependencies of video frames. 2) The joint language-image representation is successfully transferred to videos, unveiling its powerful generalization ability for recognition.

### 4.3 Zero-shot Experiments

**Training and Inference.** We pretrain X-CLIP-B/16 on Kinetics-400. More details about the evaluation protocols are provided in the *supplementary materials*.

**Results.** Zero-shot video recognition is very challenging, because the categories in the test set are unseen to the model during training. We report the results in Tab. 3 and Tab. 4. On HMDB-51 [23] and UCF-101 [41] benchmarks, our X-CLIP outperforms the previous best results by +3.8% and +13.7% in terms of top-1 accuracy respectively, as reported in Tab. 3. On Kinetics-600 [7] as presented in Tab. 4, our X-CLIP outperforms the state-of-the-art ER-ZSAR [8] by +23.1%. Such remarkable improvements can be attributed to the proposed video-text learning framework, which leverages the large-scale visual-text pretraining and seamlessly integrates temporal cues and textual prompts.

### 4.4 Few-shot Experiments

**Training and Inference.** A general  $K$ -shot setting is considered, *i.e.*,  $K$  examples are sampled from each category randomly for training. We compare with some representative methods. More details about the comparison methods and evaluation protocols are provided in the *supplementary materials*.

Table 5: Few-shot results. Top-1 accuracy is reported with 32 frames.

Method	HMDB-51				UCF-101			
	$K=2$	$K=4$	$K=8$	$K=16$	$K=2$	$K=4$	$K=8$	$K=16$
TSM [27]	17.5	20.9	18.4	31.0	25.3	47.0	64.4	61.0
TimeSformer [5]	19.6	40.6	49.4	55.4	48.5	75.6	83.7	89.4
Swin-B [29]	20.9	41.3	47.9	56.1	53.3	74.1	85.8	88.7
X-CLIP-B/16	<b>53.0</b>	<b>57.3</b>	<b>62.8</b>	<b>64.0</b>	<b>76.4</b>	<b>83.4</b>	<b>88.3</b>	<b>91.4</b>
	(+32.1)	(+16.0)	(+13.4)	(+7.9)	(+23.1)	(+7.8)	(+2.5)	(+2.0)
X-Florence	<b>51.6</b>	<b>57.8</b>	<b>64.1</b>	<b>64.2</b>	<b>84.0</b>	<b>88.5</b>	<b>92.5</b>	<b>94.8</b>
	(+30.7)	(+16.5)	(+14.7)	(+8.1)	(+30.7)	(+12.9)	(+6.7)	(+5.4)

**Results.** Tab. 5 presents the results of  $K$ -shot learning. For the extreme case where  $K=2$ , we observe that for those single-modality methods, the performance drops significantly, demonstrating that over-fitting occurs due to the serious lack of data. In contrast, X-CLIP shows robustness by surpassing them with large margins. For example, X-CLIP-B/16 outperforms Swin-B by +32.1% and +23.1% in terms of top-1 accuracy on HMDB-51 and UCF-101 with  $K=2$ , respectively. Such large improvements are mainly due to the exploitation of the semantics in text representation. It further verifies the efficacy of transferring the knowledge of the pretrained language-image models to the few-shot models. We also observe that the performance gap between our method and others decreases as the sample size increases. It demonstrates increasing data can mitigate the over-fitting for other methods. Besides, it is noteworthy that the comparison of methods with CLIP pretraining and ImageNet pretraining is not fair enough. Hence, in Sec. 4.5, we provide an additional ablation analysis and verify the performance gains mainly comes from the use of textual information, rather than the CLIP pretraining.

#### 4.5 Ablation and Analysis

Unless stated otherwise, the fully-supervised experiments are performed on Kinetics-400, while the few-shot experiments are conducted on HMDB-51 with  $K=2$ . The zero-shot evaluation is on the first split of the validation set of UCF-101. We use X-CLIP-B/16<sub>8f</sub> with single-view inference in all experiments.

**Ablation.** *The effects of the proposed components.* Tab. 6 shows the performance evolution from the pretrained image CLIP to our expanded video X-CLIP. First, we design a simple baseline that averages the CLIP features of all video frames for classification, called CLIP-Mean. It uses the text supervision but does not utilize prompting technique. We can observe that equipping the original transformer in CLIP with our proposed cross-frame communication mechanism, *i.e.* Eq. (7-9), can improve the accuracy by +1.2%. Then, appending 1-layer multi-frame integration transformer (MIT) can further improve the accuracy by +0.5%. This illustrates that our X-CLIP framework can effectively leverage temporal cues in video clips. With the proposed video-specific prompting, X-CLIP can surpass the CLIP-Mean baseline by +2.3%. It demonstrates that the video-specific prompting scheme can generate more discriminative textual representation. Meanwhile, additionally

Table 6: Component-wise analysis of our X-CLIP and other techniques.

Components	Top-1.(%)
Baseline(CLIP-Mean)	80.0
+ Cross-frame Communication	81.2(+1.2)
+ Multi-frame Integration	81.7(+1.7)
+ Video-specific Prompt	82.3(+2.3)
Techniques	
+ 4×3-views Inference	83.8(+3.8)

Table 7: Ablation study on which part to finetune. ✓ means finetuning. The CUDA memory is calculated on 2 video inputs, each containing 8 frames.

Visual	Text	Zero.	Few.	Fully.	Mem.(G)
✓	✓	<b>72.9</b>	<b>54.6</b>	<b>82.4</b>	22
✓	✗	<b>70.0</b>	50.8	<b>82.3</b>	6
✗	✓	66.8	<b>53.4</b>	79.3	20
✗	✗	64.2	47.3	79.1	4

Table 8: Ablation study on the effect of the text information.

Method	Zero-shot	Few-shot	Fully.
w/o text	/	32.0	81.6
w/ text	<b>70.0</b>	<b>50.8(+18.8)</b>	<b>82.3(+0.7)</b>

Table 9: Comparison with different prompting methods.

Method	Fully.	Few.	Zero.
w/o prompt	81.7	49.6	63.2
Ensemble. [36]	81.7	49.6	63.9
Vectors. [60]	82.0	49.9	63.2
Ours	<b>82.3(+0.3)</b>	<b>50.8(+0.9)</b>	<b>70.0(+6.1)</b>

using multi-view inference can boost the performance by +1.5%. Overall, with our proposed methods and all the techniques mentioned above, X-CLIP can boost the top-1 accuracy of the CLIP-Mean baseline from 80.0% to 83.8%.

*Which branch to finetune?* In order to demonstrate which branch should be finetuned when transferred to different downstream tasks, we separately freeze the parameters of the pretrained image and text encoder. Note that the randomly initialized parameters are always finetuned. From Tab. 7, we summarize the following observations. 1) For fully-supervised setting, finetuning the image encoder brings +3.0% improvements, while freezing the text encoder reduces the CUDA memory from 22G to 6G with minor performance loss. 2) For few-shot setting, we find the top-2 results are achieved by finetuning the text encoder. We conjecture the reason is that with few samples, the text encoder suffers less from the over-fitting than the over-parameterized image model. 3) For zero-shot setting, finetuning both the image and the text encoder achieves the best results.

*The effects of text.* To evaluate the impact of text, we replace the text encoder with a randomly initialized fully-connected layer as the classification head. From Tab. 8, we can observe that, without the text branch, the model cannot adapt to zero-shot setting, because there is no data to initialize the head. For the few-shot and fully-supervised experiments, text information can bring +18.8% and +0.7% gains, respectively. This indicates the semantic information involved in text representation is beneficial to classification, especially for low-shot learning.

*The effects of pretraining.* In Tab. 10, we investigate the effects of pretraining. We use ViT-B/16 pretrained on IN-1k/21k as the video encoder in our framework. Though the pretrained image encoder and text encoder are not in a joint embed-

Table 10: Ablation study on the different pretraining.

Pretrain	Top-1. (%)	Top-5. (%)
ImageNet-1k	75.9	90.2
ImageNet-21k	79.8	94.0

Table 11: Comparison of two sampling methods.

#F	Train \ Test	multi-view $\rightarrow$ single-view	
		Dense	Sparse
8	Dense	81.9 $\rightarrow$ 77.8(-4.1)	82.4 $\rightarrow$ 81.1(-1.3)
	Sparse	82.2 $\rightarrow$ 77.3(-4.9)	<b>83.4</b> $\rightarrow$ <b>82.3</b> (-1.1)
32	Dense	82.8 $\rightarrow$ 78.8(-4.0)	83.2 $\rightarrow$ 83.0(-0.2)
	Sparse	83.0 $\rightarrow$ 77.9(-5.1)	<b>84.4</b> $\rightarrow$ <b>84.2</b> (-0.2)

ding space, the model with IN-21k and IN-1k pretraining still achieve 79.8% and 75.9% top-1 accuracy on Kinetics-400, yet much inferior to the original CLIP large-scale pretraining (82.3%).

**Analysis.** *Comparison with other prompting methods.* We compare with two existing methods in Tab. 9: prompt ensembling [36] with 16 handcraft templates and learnable vectors [60] with length 16. It can be seen that our video-specific prompts outperforms others, especially in zero-shot setting (+6.1%). This demonstrates the efficacy of our method, which generates more adaptive prompts and better textual representation for unseen videos.

*Dense v.s. sparse sampling.* We further explore what is the best sampling strategy for our method in Tab. 11. We find that the dense sampling does not perform well as in previous works [29,12,3]. In contrast, the sparse sampling best matches our method. Regardless of the number of frames and views, using sparse sampling both in training and inference achieves the best performance.

*Single-view v.s. multi-view inference.* Although it can improve performance, multi-view inference takes relatively high computational cost, because the cost grows linearly with the number of views. In Tab. 11, we show that our multi-modality models with sparse sampling is robust to the number of views, *i.e.*, single-view can achieve comparable performance to 10 temporal views. The underlying reason is the language-image models provide robust representation.

## 5 Conclusion

In this work, we present a simple approach that adapts the pretrained language-image models to video recognition. To capture the temporal information, we propose a cross-frame attention mechanism that explicitly exchanges information across frames. A video-specific prompting technique is designed to yield instance-level discriminative textual representation. Extensive experiments under three different learning scenarios demonstrate the effectiveness of our method. In future work, we plan to extend our method to different video tasks beyond classification.

**Acknowledgements.** This research was supported in part by the National Key Research and Development Program of China under Grant No. 2018AAA0100400, and the National Natural Science Foundation of China under Grants 61976208, 62071466 and 62076242, and the InnoHK project. HL was not supported by any fund for this research.

## References

1. Akata, Z., Perronnin, F., Harchaoui, Z., Schmid, C.: Label-embedding for image classification. *IEEE T-PAMI* pp. 1425–1438 (2015)
2. Akata, Z., Reed, S., Walter, D., Lee, H., Schiele, B.: Evaluation of output embeddings for fine-grained image classification. In: *CVPR*. pp. 2927–2936 (2015)
3. Arnab, A., Dehghani, M., Heigold, G., Sun, C., Lučić, M., Schmid, C.: Vivit: A video vision transformer. In: *ICCV*. pp. 6836–6846 (2021)
4. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. *arXiv preprint arXiv:1607.06450* (2016)
5. Bertasius, G., Wang, H., Torresani, L.: Is space-time attention all you need for video understanding? In: *ICML*. pp. 813–824 (2021)
6. Brattoli, B., Tighe, J., Zhdanov, F., Perona, P., Chalupka, K.: Rethinking zero-shot video classification: End-to-end training for realistic applications. In: *CVPR*. pp. 4613–4623 (2020)
7. Carreira, J., Noland, E., Banki-Horvath, A., Hillier, C., Zisserman, A.: A short note about kinetics-600. *arXiv preprint arXiv:1808.01340* (2018)
8. Chen, S., Huang, D.: Elaborative rehearsal for zero-shot action recognition. In: *ICCV*. pp. 13638–13647 (2021)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *CVPR*. pp. 248–255. *Ieee* (2009)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *ICLR* (2021)
11. Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., Feichtenhofer, C.: Multiscale vision transformers. In: *ICCV*. pp. 6824–6835 (2021)
12. Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: *ICCV*. pp. 6202–6211 (2019)
13. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: Devise: A deep visual-semantic embedding model. In: *NeurIPS*. pp. 2121–2129 (2013)
14. Gao, J., Zhang, T., Xu, C.: I know the relationships: Zero-shot action recognition via two-stream graph convolutional networks and knowledge graphs. In: *AAAI*. vol. 33, pp. 8303–8311 (2019)
15. Gao, P., Geng, S., Zhang, R., Ma, T., Fang, R., Zhang, Y., Li, H., Qiao, Y.: Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544* (2021)
16. Ghosh, P., Saini, N., Davis, L.S., Shrivastava, A.: All about knowledge graphs for actions. *arXiv preprint arXiv:2008.12432* (2020)
17. Girdhar, R., Grauman, K.: Anticipative video transformer. In: *ICCV* (2021)
18. Herath, S., Harandi, M., Porikli, F.: Going deeper into action recognition: A survey. *Image and vision computing* (2017)
19. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: *ICML*. pp. 4904–4916 (2021)
20. Ju, C., Han, T., Zheng, K., Zhang, Y., Xie, W.: Prompting visual-language models for efficient video understanding. In: *CVPR* (2022)
21. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: *CVPR*. pp. 1725–1732 (2014)

22. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
23. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: Hmdb: a large video database for human motion recognition. In: ICCV. pp. 2556–2563 (2011)
24. Li, K., Wang, Y., Zhang, J., Gao, P., Song, G., Liu, Y., Li, H., Qiao, Y.: Uniformer: Unifying convolution and self-attention for visual recognition. In: ICLR (2022)
25. Li, Y., Ji, B., Shi, X., Zhang, J., Kang, B., Wang, L.: Tea: Temporal excitation and aggregation for action recognition. In: CVPR. pp. 909–918 (2020)
26. Li, Y., Wu, C.Y., Fan, H., Mangalam, K., Xiong, B., Malik, J., Feichtenhofer, C.: Improved multiscale vision transformers for classification and detection. In: CVPR (2022)
27. Lin, J., Gan, C., Han, S.: Tsm: Temporal shift module for efficient video understanding. In: ICCV. pp. 7083–7093 (2019)
28. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: ICCV (2021)
29. Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: CVPR (2022)
30. Liu, Z., Wang, L., Wu, W., Qian, C., Lu, T.: Tam: Temporal adaptive module for video recognition. In: ICCV. pp. 13708–13718 (2021)
31. Miech, A., Zhukov, D., Alayrac, J.B., Tapaswi, M., Laptev, I., Sivic, J.: Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In: ICCV. pp. 2630–2640 (2019)
32. Neimark, D., Bar, O., Zohar, M., Asselmann, D.: Video transformer network. arXiv preprint arXiv:2102.00719 (2021)
33. Patrick, M., Campbell, D., Asano, Y., Misra, I., Metze, F., Feichtenhofer, C., Vedaldi, A., Henriques, J.F.: Keeping your eye on the ball: Trajectory attention in video transformers. In: NeurIPS (2021)
34. Qin, J., Liu, L., Shao, L., Shen, F., Ni, B., Chen, J., Wang, Y.: Zero-shot action recognition with error-correcting output codes. In: CVPR. pp. 2833–2842 (2017)
35. Qiu, Z., Yao, T., Mei, T.: Learning spatio-temporal representation with pseudo-3d residual networks. In: ICCV. pp. 5533–5541 (2017)
36. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: ICML (2021)
37. Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: Denseclip: Language-guided dense prediction with context-aware prompting. In: CVPR (2022)
38. Romera-Paredes, B., Torr, P.: An embarrassingly simple approach to zero-shot learning. In: ICML. pp. 2152–2161 (2015)
39. Ryoo, M., Piergiovanni, A., Arnab, A., Deghani, M., Angelova, A.: Tokenlearner: Adaptive space-time tokenization for videos. In: NeurIPS (2021)
40. Selva, J., Johansen, A.S., Escalera, S., Nasrollahi, K., Moeslund, T.B., Clapés, A.: Video transformers: A survey. arXiv preprint arXiv:2201.05991 (2022)
41. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. arXiv preprint arXiv:1212.0402 (2012)
42. Sun, C., Baradel, F., Murphy, K., Schmid, C.: Learning video representations using contrastive bidirectional transformer. In: ECCV (2020)
43. Sun, C., Myers, A., Vondrick, C., Murphy, K., Schmid, C.: Videobert: A joint model for video and language representation learning. In: ICCV. pp. 7464–7473 (2019)

44. Tran, D., Bourdev, L., Fergus, R., Torresani, L., Paluri, M.: Learning spatiotemporal features with 3d convolutional networks. In: ICCV. pp. 4489–4497 (2015)
45. Tran, D., Wang, H., Torresani, L., Ray, J., LeCun, Y., Paluri, M.: A closer look at spatiotemporal convolutions for action recognition. In: CVPR (2018)
46. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017)
47. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L.: Temporal segment networks: Towards good practices for deep action recognition. In: ECCV. pp. 20–36 (2016)
48. Wang, M., Xing, J., Liu, Y.: Actionclip: A new paradigm for video action recognition. arXiv preprint arXiv:2109.08472 (2021)
49. Wang, Q., Chen, K.: Alternative semantic representations for zero-shot human action recognition. In: ECML PKDD. pp. 87–102 (2017)
50. Xie, S., Sun, C., Huang, J., Tu, Z., Murphy, K.: Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: ECCV. pp. 305–321 (2018)
51. Xu, H., Ghosh, G., Huang, P.Y., Okhonko, D., Aghajanyan, A., Metze, F., Zettlemoyer, L., Feichtenhofer, C.: Videoclip: Contrastive pre-training for zero-shot video-text understanding. In: EMNLP (2021)
52. Xu, X., Hospedales, T.M., Gong, S.: Multi-task zero-shot action recognition with prioritised data augmentation. In: ECCV. pp. 343–359 (2016)
53. Yan, S., Xiong, X., Arnab, A., Lu, Z., Zhang, M., Sun, C., Schmid, C.: Multiview transformers for video recognition. In: CVPR (2022)
54. Yuan, L., Chen, D., Chen, Y.L., Codella, N., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., et al.: Florence: A new foundation model for computer vision. arXiv preprint arXiv:2111.11432 (2021)
55. Zhang, B., Yu, J., Fifty, C., Han, W., Dai, A.M., Pang, R., Sha, F.: Co-training transformer with videos and images improves action recognition. arXiv preprint arXiv:2112.07175 (2021)
56. Zhang, L., Xiang, T., Gong, S.: Learning a deep embedding model for zero-shot learning. In: CVPR. pp. 2021–2030 (2017)
57. Zhang, R., Fang, R., Gao, P., Zhang, W., Li, K., Dai, J., Qiao, Y., Li, H.: Tip-adapter: Training-free clip-adapter for better vision-language modeling. arXiv preprint arXiv:2111.03930 (2021)
58. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point cloud understanding by clip. In: CVPR (2021)
59. Zhou, C., Loy, C.C., Dai, B.: Denseclip: Extract free dense labels from clip. arXiv preprint arXiv:2112.01071 (2021)
60. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. arXiv preprint arXiv:2109.01134 (2021)
61. Zhu, L., Yang, Y.: Actbert: Learning global-local video-text representations. In: CVPR. pp. 8746–8755 (2020)
62. Zhu, Y., Li, X., Liu, C., Zolfaghari, M., Xiong, Y., Wu, C., Zhang, Z., Tighe, J., Manmatha, R., Li, M.: A comprehensive study of deep video action recognition. arXiv preprint arXiv:2012.06567 (2020)
63. Zhu, Y., Long, Y., Guan, Y., Newsam, S., Shao, L.: Towards universal representation for unseen action recognition. In: CVPR. pp. 9436–9445 (2018)