

# ViCTR: Video-conditioned Text Representations for Activity Recognition

Kumara Kahatapitiya<sup>1\*</sup> Anurag Arnab<sup>2</sup> Arsha Nagrani<sup>2</sup> Michael S. Ryoo<sup>1,2</sup>

<sup>1</sup>Stony Brook University <sup>2</sup>Google Research

kkahatapitiy@cs.stonybrook.edu

## Abstract

Vision-Language models (VLMs) have excelled in the image-domain—especially in zero-shot settings—thanks to the availability of vast pretraining data (i.e., paired image-text samples). However for videos, such paired data is not as abundant. Therefore, video-VLMs are usually designed by adapting pretrained image-VLMs to the video-domain, instead of training from scratch. All such recipes rely on augmenting visual embeddings with temporal information (i.e., image  $\rightarrow$  video), often keeping text embeddings unchanged or even being discarded. In this paper, we argue the contrary, that better video-VLMs can be designed by focusing more on augmenting text, rather than visual information. More specifically, we introduce Video-conditioned Text Representations ( $V_iCTR$ ): a form of text embeddings optimized w.r.t. visual embeddings, creating a more-flexible contrastive latent space. Our model can further make use of freely-available semantic information, in the form of visually-grounded auxiliary text (e.g. object or scene information). We evaluate our model on few-shot, zero-shot (HMDB-51, UCF-101), short-form (Kinetics-400) and long-form (Charades) activity recognition benchmarks, showing strong performance among video-VLMs.

## 1. Introduction

Video understanding poses significant challenges, often adding to the complications in image domain such as model complexity and annotation costs. The additional temporal dimension and different modalities of data introduce useful cues, but also can be redundant, raising interesting questions about trade-offs. Activity Recognition (i.e., classification) in particular—as the prominent task in video understanding—has long been explored by the community in these research directions. Whether it is efficient architecture variants ranging from CNNs [13, 31, 57] to Transformers [2, 5, 12], training schemes from fully-supervised [7, 14] to self-supervised [15, 47, 53] or data regimes from unimodal [63, 73] to multimodal [20, 39], the progress has been steady and exciting. More recently, with the availability of internet-scale paired image-text data, the direction of vision-language models

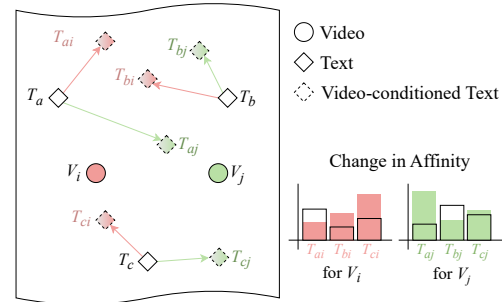


Figure 1. **Video-conditioned Text Representations:** Pretrained image-VLMs can generate reasonable visual embeddings for videos (e.g. by temporally-pooling frame embeddings), together with paired text embeddings. However, usually, these text embeddings are not dependent on visual information—meaning, they are common for every video. Such representations lack the flexibility to align properly in a shared vision-language latent space, when optimized based on a contrastive similarity (i.e., *Affinity*) w.r.t. all videos. However, with *Video-conditioned Text* representations that specialize uniquely for each video, we grant more freedom for text embeddings to move in the latent space, and adapt to different scenarios (e.g. more-challenging recognition tasks).

(VLMs) [23, 50] have emerged dominant, achieving strong generalization across numerous benchmarks. However, the progress of VLMs in the video domain is yet to be caught-up to its full potential.

Following the seminal VLMs such as CLIP [50] and ALIGN [23], there have been significant strides in tasks such as image classification [80, 85, 87], open-vocabulary object detection [18, 38], text-to-image retrieval [61, 82] and robot manipulation [24, 86]. Such models are usually pretrained on paired image-text data based on a contrastive learning framework. The idea is to have two separate backbones—an Image Encoder and a Text Encoder, that generate embeddings in a joint latent space. To optimize this space, the corresponding pairs of embeddings are drawn closer, by increasing their similarity (i.e., *Affinity*). The key advantage of such models is that, at inference, any semantic concept (given as a text input) can be embedded in the same space, giving intriguing zero-shot or few-shot transfer capabilities [1, 86]. For instance, CLIP [50] excels at classifying unseen attribute categories (e.g. objects, scenes), or even counting

\*Work done as a student researcher at Google.

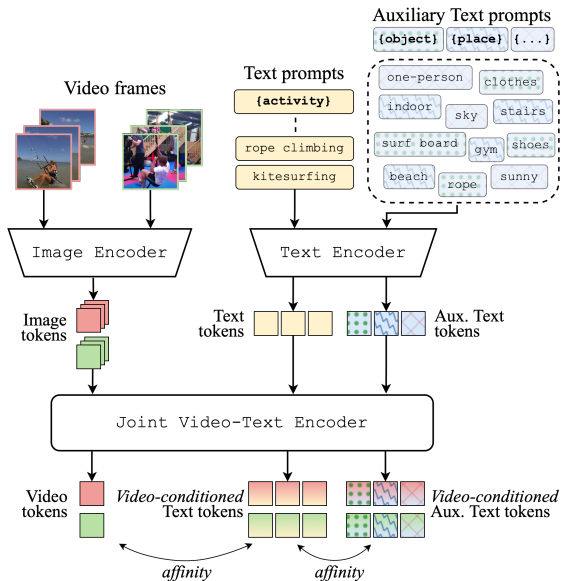


Figure 2. **Overview of VicTR:** First, we extract image (*i.e.*, frame) and text tokens using a pretrained image-VLM. Next, such tokens go through a joint video-text encoder, generating video tokens and *video-conditioned* text tokens, based on which, we compute affinity-based logits for classification. Optionally, any semantic concept (given as auxiliary text) can also be processed similarly, to help guide the classifier. This is motivated based on the co-occurrence of semantics (*e.g.* rope, gym, one-person) and categories-of-interest, *i.e.*, activity classes in our setting (*e.g.* rope climbing). Here, the color change of text tokens represents the idea of video-conditioning.

such occurrences [86]. However, these VLMs do not perform well in tasks that require specialized knowledge, such as localizing (*e.g.* detection/segmentation) or temporal reasoning (*e.g.* activity recognition), at least not out-of-the-box, as their training objective has not seen any location or temporal cues. Yet, with task-specific finetuning, such models can readily be adapted to specialized domains [18, 40].

In the video domain, training VLMs from scratch may show a limited success [74]—while also being expensive—due to the lack of paired data at scale. As a compromise, the common practice is to adapt pretrained image-VLMs to video, by introducing temporal information. Such methods either insert temporal modules within the image backbone itself to have cross-frame interactions [40], or use a post-processing video head on-top of the image backbone [4, 33, 36, 65]. In both cases, image embeddings are enhanced as video embeddings. However, the use of text embeddings varies among different approaches. Text may either be discarded [33], kept frozen [36, 65], used as conditioning [4] (to further enhance video embeddings), or fully-updated jointly with video [40]. More often than not, the main focus is on visual embeddings (*i.e.*, converting image  $\rightarrow$  video), and the impact of updating text has been limited.

Nevertheless, video models benefit from semantic in-

formation [22, 68, 86]. In fact, certain attributes (*e.g.* objects, scene or human subjects) are directly tied with specific activities, and can simplify their recognition. For instance, the presence of attributes such as [rope, gym, one-person] can narrow down the potential activity to battling ropes or rope climbing. VLMs are especially suited to take advantage of such semantics. Any concept represented as text can be visually-grounded based on paired embeddings (in zero-shot), to extract relevant attributes for a given input that benefit recognition tasks. Such visually-grounded semantics are cheap in-terms of both annotation and compute costs, yet highly-useful.

Motivated by the above we propose VicTR, focusing on adapting text information to the video domain. More specifically, we generate *Video-conditioned Text* embeddings (see Fig. 1), while jointly-training both textual and visual features generated by an image-VLM. By finetuning text embeddings, we observe significant gains in our framework, compared to just finetuning visual embeddings (similar to the observations in [87]). We can also make use of freely-available auxiliary semantic information, represented in the form of visually-grounded text embeddings. Fig. 2 shows an overview of the proposed architecture. Our video-conditioned text embeddings are unique to each video, allowing more-flexibility to move in the latent space and generalize to complex downstream tasks. Optionally, our video-conditioned auxiliary text can further help optimize this latent space. We evaluate VicTR on few-shot, zero-shot, short-form and long-form activity recognition, validating its strong generalization capabilities among video-VLMs.

## 2. Related Work

**Video understanding** is about reasoning based on spatio-temporal inputs. Compared to image inputs, videos bring additional useful cues such as motion or multiple modalities (*e.g.* audio) into play, but also any associated complications such as increased compute requirements and redundancy in data. Convolutional networks (CNNs) [7, 63, 67, 73] and Recurrent models [11, 83] have been the state-of-the-art in video modeling, prior to the rise of Transformers [2, 5, 34, 55]. Multi-stream models [7, 14] that make use of different spatio-temporal views [14, 53] or modalities (*e.g.* optical-flow [7, 20], audio [21, 39, 54]) have emerged, tackling benchmark tasks such as activity recognition [28, 29], localization [17, 60, 83] or text-to-video retrieval [75]. To handle longer video inputs, models have focused on efficient temporal modeling [26, 44, 45], or memory mechanisms [58, 70, 71]. While Neural Architecture Search (NAS) has enabled efficient model designs [13, 56, 57], self-supervised methods [15, 20, 47, 53] have alleviated the high demand for annotated data. More recently, language-supervision has been of interest for video understanding due to the strong generalization capabilities shown in the image domain.

**Vision-Language Models (VLMs)** are usually trained on internet-scale paired visual-language (*e.g.* image-text) data. Seminal work such as CLIP [50] and ALIGN [23] have shed the light on the capabilities of such models, especially for zero-shot transfer. Since then, VLM literature has flourished, with applications in open-vocabulary object detection [18, 38], open-set classification [48], retrieval [3, 61, 82], captioning [81], segmentation [51, 76], robot manipulation [24, 27, 86] and many other domains. Although VLMs are generally trained on image-text data, there are intuitive variants which are trained either only on images [64] or only on text [41]. The commonly-used similarity-based objective of VLMs has also been repurposed to specialized domains, through prompt learning [90] or engineering [18, 42]. The text encoder of VLMs can be a powerful mapping from semantic concepts to latent embeddings [37]. Many foundation models [1, 84, 85] follow similar design principles as VLMs, thriving in zero-shot [19] or few-shot [90] settings. Recent work combining Large Language Models (LLMs) with VLMs show how language can act as a communication-medium between models [68, 86, 89]. In [37], authors use an LLM to represent object classes as a set of its semantic attributes, to learn a better classifier.

As for video-VLMs, they are either trained from scratch on video-text data [74, 81], or more-often than not, finetuned initializing from a pretrained image-VLM [9, 32, 80]. Some are even trained on both image and video data paired with text [3]. The success of VLMs in the image domain has fueled similar research directions in the video domain.

**Adapting image-text models to video** is a common practice when designing video-VLMs. A general and effective recipe for such adaptation is proposed in [9]. It consists of temporal modeling, multi-modal fusion, auxiliary training objectives, and both image/video data at scale. All others usually make use of a subset of these concepts. CLIP-ViP [78] is trained with different sources of data and multiple cross-modal training objectives. VideoCoCa [80] extends CoCa [84] with attention-pooled frame embeddings, which are used to decode text captions in a generative framework. MOV [48] is trained with additional audio/flow encoders through cross-modal attention, keeping image-text encoders frozen. Video-specific prompts can also be learned with such frozen encoders [25]. Vi-Fi [52] shows that simply finetuning CLIP image-text encoders without any specialized modules can generate video representations efficiently.

Apart from the above, there exists a body of prior work that closely-relates to ViCTR. ActionCLIP [65] upgrades its CLIP image-encoder with (1) parameter-free temporal layers (TSM [31]) within the backbone, and (2) a temporal transformer head, while keeping the text-encoder fixed. Similarly, CLIP4clip [36] just uses a temporal transformer head to update visual embeddings. CLIPHitchhiker’s [4] generates *text-conditioned video* embeddings by temporally-pooling

frame embeddings, conditioned on each text query. In this case, a given video generates multiple different visual embeddings, one per each text embedding. EVL [33] completely discards text. It acts as an initialization for a visual-only backbone, consisting of CLIP image encoder and a temporal, class-conditioned decoder. X-CLIP [40] introduces trainable temporal layers within its backbone image encoder, and generates video-specific text prompts. Meaning, it finetunes both encoders similar to ours. However, it does not allow interaction among text embeddings, nor with fine-grained visual information (but only, with temporally-aggregated information). Hence, it shows limited gains from adapting text to video domain. In contrast, our *video-conditioned text* embeddings that are unique for each video, interacts with both fine-grained visual embeddings and other text embeddings, to enable a better contrastive framework, and in-turn, a more-flexible alignment in the latent space.

### 3. Background: image-VLMs to video

In this section, we introduce the generic framework for adapting image-VLMs to video, and discuss how prior work fit into it. We consider CLIP [50] as the image-VLM, which is widely-adapted thanks to its convincing performance and open-source models. It consists of two encoders: Image and Text, optimized together on internet-scale paired image-text data. Image Encoder ( $\text{ENC}_{\text{img}}$ ) is a ViT [10]. Given an input image  $I \in \mathbb{R}^{H \times W \times 3}$ , it is broken down to patch embeddings (*i.e.*, tokens) and processed through multiple transformer layers. The class token [cls] is sampled as the visual embedding  $e_{\text{img}}$ . Text Encoder ( $\text{ENC}_{\text{txt}}$ ) is a causal transformer, operating on tokenized text. Each class-label (or, any semantic concept) given as text  $T$ , is first converted into a prompt based on a template such as “a photo of {class}.”, and tokenized with Byte Pair Encoding (BPE) [59] at the input of Text Encoder. Following multiple causal transformer layers, the [EOS] (*i.e.*, end-of-sequence) token is extracted as the text embedding  $e_{\text{txt}}$ .

$$\begin{aligned} e_{\text{img}} &= \text{ENC}_{\text{img}}(I), \\ e_{\text{txt}} &= \text{ENC}_{\text{txt}}(T). \end{aligned}$$

The two encoders are jointly-optimized with Cross-Entropy loss, where logits are computed based on the similarities (*i.e.*, *affinities*) between visual and text embeddings. The corresponding pairs of embeddings (*i.e.*, positives) are drawn together ( $\uparrow$  affinity) in a joint embedding space, whereas the others (*i.e.*, negatives) are pushed apart ( $\downarrow$  affinity).

$$\text{Affinity}(e_{\text{img}}, e_{\text{txt}}) = \frac{\langle e_{\text{img}}, e_{\text{txt}} \rangle}{\|e_{\text{img}}\|_2 \|e_{\text{txt}}\|_2}.$$

When adapting this framework to the video domain, the above Image encoder, Text encoder and the learning objective usually stays the same. But now, video frames  $V \in \mathbb{R}^{T \times H \times W \times 3} = [I^1, I^2, \dots, I^T]$  become inputs to

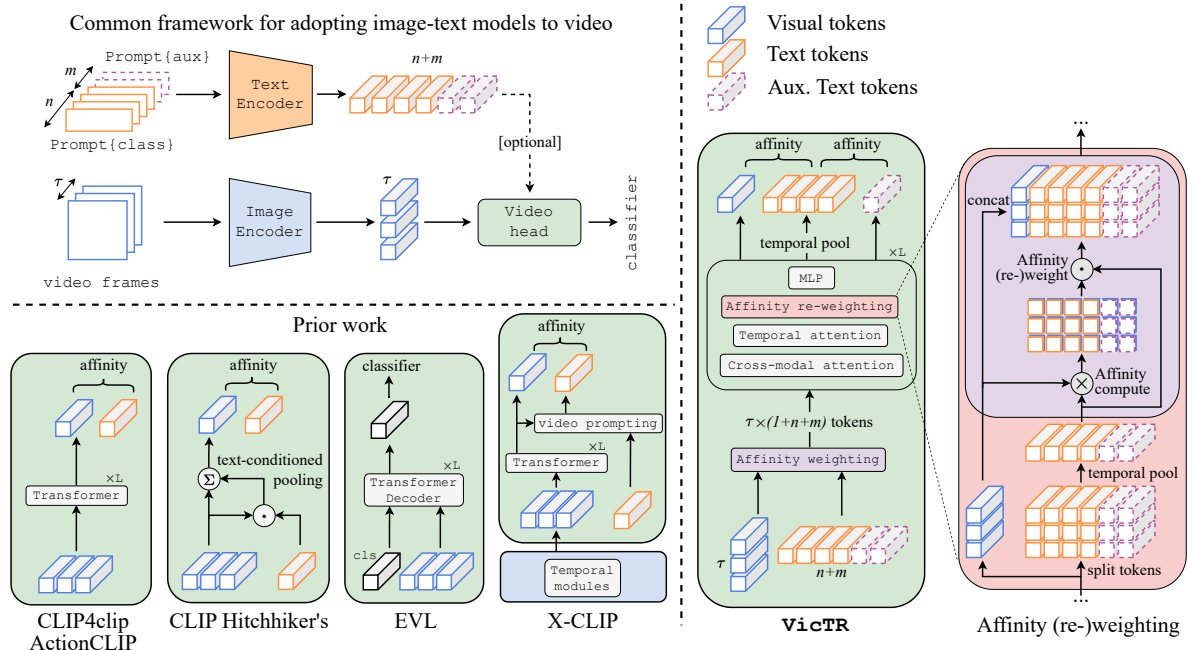


Figure 3. **Detailed view of ViCTR compared to prior art:** There exist multiple closely-related work on adapting pretrained image-VLMs to video, such as CLIP4Clip [36], ActionCLIP [65], CLIP Hitchhiker’s [4], EVL [33] and X-CLIP [40]. All these follow a common framework (top-left). Text prompts and video frames are first encoded using two-separate encoders, and then fed into a video head to enable temporal reasoning. It is optional to use text tokens within the video head. Often, text information is kept unchanged [36, 65], or even discarded [33] (bottom-left). CLIP Hitchhiker’s [4] however, use text as conditioning to generate *text-conditioned video* embeddings. X-CLIP [40]—which is the closest to our method, jointly-optimizes visual and text tokens. But, it provides limited information for text to contrast-against: only temporally-aggregated visual embeddings, showing marginal gains from updating text. In contrast, ViCTR allows text to contrast against both fine-grained visual and other text information, while also jointly-optimizing both modalities. We generate *video-conditioned text* representations, *i.e.*, text uniquely-specialized for each video (refer to Fig. 1). Our video head consists of three key operations: (1) *Token-boosting*, (2) *Cross-modal attention*, and (3) *Affinity (re-)weighting* (right). Token-boosting creates dedicated text tokens per video and per timestep, weighted by per-frame affinities of a given video. These enable us to model variations of semantics (represented as text) over time. Affinity (re-)weighting highlights or down-plays each text class, grounded on visual information. Such affinity weights are similar to the ones in CLIP [50] training objective, making the optimization more-consistent. Cross-modal attention enables message passing between both visual-textual and textual-textual modes, creating a better contrastive representation. Also, optionally, ViCTR can make use of auxiliary semantics (*e.g.* object, scene, human-subjects) given as visually-grounded text (refer to Fig. 2). Such auxiliary semantics help align our video-conditioned text embeddings in the latent space.

the Image encoder (while each being processed separately), and further go through a Video Head  $\text{Head}_{\text{vid}}$  to induce temporal reasoning capabilities. Optionally, text embedding  $e_{\text{txt}}$  may also be updated or used as a conditioning within the Video Head.

$$e_{\text{vid}}, [e_{\text{txt}}] = \text{Head}_{\text{vid}}(e_{\text{img}}^1, \dots, e_{\text{img}}^{\mathcal{T}}, [e_{\text{txt}}]).$$

Here,  $[\cdot]$  denotes optional embeddings. This Video Head may just be a temporal pooling layer or a temporal transformer as in [36, 65], or may even consist of more-specialized modules. Text embeddings could either be discarded as in [33], used as a conditioning as in [4], or jointly-updated with video embeddings as in [40]. Finally, logits are computed based on video-text affinities if text is not discarded, or as a linear mapping of video embeddings if text is discarded. This generic framework is shown in Fig. 3 (top-left), along with variations of prior work in Fig. 3 (bottom-left).

#### 4. Video-conditioned Text Representations

In ViCTR, we adapt a pretrained image-VLM (*e.g.* CLIP [50]) to video, focusing more on text representations. Refer to Fig. 3 (right) for a detailed view. The image-VLM has not seen any temporal information during training. While it obviously affects the temporal reasoning capabilities of the visual embeddings— which most prior work focus on addressing, it also affects the text embeddings as well. The learnt latent space (and, the affinity-based objective) depends on both these embeddings. Thus, we consider text equally as important, if not more, in contrast to prior work

ViCTR consists of a joint video-text model as  $\text{Head}_{\text{vid}}$ , which consumes both visual and text embeddings from the image-VLM. It outputs text embeddings uniquely-specified for each video, *i.e.*, *Video-conditioned Text* embeddings. It relies on three main components: (1) *Token-boosting*, (2)

Cross-modal attention, and (3) Affinity (re-)weighting. Optionally, it can also benefit from any semantic concept available as auxiliary text, to optimize its latent space. Following subsections look at each of these in detail.

Let us first introduce a few additional notations. Consider a fixed vocabulary of  $n$  activity-classes given by  $[T^1, T^2, \dots, T^n]$ , and optional  $m$  auxiliary semantic categories given by  $[A^1, A^2, \dots, A^m]$ . The corresponding text embeddings can be denoted as  $\{e_{\text{txt}}^x \mid x = 1, 2, \dots, n\}$  and  $\{e_{\text{aux}}^y \mid y = 1, 2, \dots, m\}$ . Also, given an input video  $V^i$  of  $\mathcal{T}$  frames, the corresponding image embeddings can be denoted as  $\{e_{\text{img}}^{i,t} \mid t = 1, 2, \dots, \mathcal{T}\}$ . The inputs to our Video Head are  $e_{\text{img}}^{i,t}$ ,  $e_{\text{txt}}^x$  and  $e_{\text{aux}}^y$  tokens. As visual embeddings are extracted per-frame and the text embeddings per prompt, there is no interaction among frame tokens, among text tokens or, across frame-text tokens up to this point.

### 4.1. Token-boosting

To introduce *video-conditioned text* embeddings, we first create a dedicated set of text tokens per video, by replicating the outputs of the backbone text encoder. Going further, we also create text tokens per each frame. This is done by weighting text tokens with the corresponding frame-text affinities. Formally, given  $(n + m)$  text tokens, we end up with  $\mathcal{T} \times (n + m)$  dedicated text tokens per video, at the input of our video head. Refer to Fig. 3 (right).

$$\begin{aligned} e_{\text{txt}}^{i,t,x} &= e_{\text{txt}}^x \cdot \text{SigAffinity}(e_{\text{img}}^{i,t}, e_{\text{txt}}^x), \\ e_{\text{aux}}^{i,t,y} &= e_{\text{aux}}^y \cdot \text{SigAffinity}(e_{\text{img}}^{i,t}, e_{\text{aux}}^y). \end{aligned}$$

Here,  $\text{SigAffinity}(\cdot)$  corresponds to affinity-weights normalized in  $[0, 1]$  range. We convert the values given by  $\text{Affinity}(\cdot)$  that lie in  $[-1, 1]$ , to be affinity-weights, by scaling with a learnable weight ( $w$ ) and feeding through a sigmoid activation.

$$\text{SigAffinity}(\cdot) = \text{Sigmoid}(w \cdot \text{Affinity}(\cdot)).$$

Although such affinity-weights based on the original image-VLM embeddings are not ideal for temporal reasoning, it initializes a noisy-version of our *video-conditioned text* embeddings that gets updated iteratively, later in the network. Such a token-boosting brings multiple other benefits. (1) More tokens means higher the model capacity. It can help learn better representations, but also adds a compute overhead (which we handle through other measures, as discussed later). (2) It also highlights relevant text tokens by grounding text on visual embeddings, while diminishing irrelevant ones. Subsequent attention mechanisms attend less to such diminished tokens, simplifying the gradient flow during learning. In other words, it acts as a soft-selection of relevant semantics, specific to each video. (3) Finally, it enables our model to capture variations of semantic categories over time. How

certain attributes appear (or, disappear) over time is an important motion cue for activity recognition.

Next, we concatenate such boosted text tokens with visual tokens (corresponding to  $\mathcal{T}$  frames), and feed  $\mathcal{T} \times (1+n+m)$  tokens to the subsequent layers.

$$z^{i,t} = \text{Concat}(e_{\text{img}}^{i,t}, e_{\text{txt}}^{i,t,x} \big|_{x=\{1,\dots,n\}}, e_{\text{aux}}^{i,t,y} \big|_{y=\{1,\dots,m\}}).$$

Such  $Z_0^i = [z^{i,1}, \dots, z^{i,\mathcal{T}}]$  tokens go through  $L$  transformer layers in our Video Head. Each layer ( $l$ ) consists of cross-modal attention, temporal attention, affinity (re-)weighting and linear (MLP) layers.

### 4.2. Cross-modal and Temporal attention

We consider our token representation to be two-dimensional (*i.e.*, cross-modal and temporal), and apply divided self-attention (MSA) on each axis as in [2, 5]. First, we have a Cross-modal attention layer. Here, each visual token could attend to all text tokens at the same timestep, and each text token could attend to both the visual token and other text tokens at the same timestep. Since text tokens are already affinity-weighted, attention weights do not draw information from irrelevant semantic classes. Next, we have a Temporal attention layer. Here, both visual and text tokens go through a shared set of parameters, learning temporal cues in visual modality (*i.e.*,  $e_{\text{img}} \rightarrow e_{\text{vid}}$ ), and modeling variations of semantics across time in textual modality.

$$\begin{aligned} \hat{Z}_l^i &= Z_l^i + \text{MSA}_{\text{cross}}(\text{LN}(Z_l^i)), \\ \bar{Z}_l^i &= \hat{Z}_l^i + \text{MSA}_{\text{temporal}}(\text{LN}(\hat{Z}_l^i)). \end{aligned}$$

Here,  $\text{LN}(\cdot)$  stands for LayerNorm operation. Having a divided attention across two-axes instead of a joint-attention eases the compute requirement of our video head.

### 4.3. Affinity (re-)weighting

As previously discussed, the original affinities based on the image-VLM embeddings can be noisy, in the context of temporal reasoning. Now, as we have updated both our visual (*i.e.*, video) and text tokens with cross-modal and temporal information, they are in a better state to re-compute affinities. Hence, we compute new affinity values and re-weight the text tokens accordingly. Refer to Fig. 3 (rightmost). First, we split video and text tokens as in,

$$[\bar{e}_{\text{vid},l}^{i,t}, \bar{e}_{\text{txt},l}^{i,t,x} \big|_{x=\{1,\dots,n\}}, \bar{e}_{\text{aux},l}^{i,t,y} \big|_{y=\{1,\dots,m\}}] = \bar{z}_l^{i,t}.$$

Next, we temporally-pool the text tokens to come up with a compressed representation, on which we perform affinity re-weighting. This is similar to token-boosting, but done with updated video-text embeddings that are already video-conditioned. Without loss of generality, the same operations apply for auxiliary text tokens.

$$\begin{aligned} \bar{e}_{\text{txt},l}^{i,x} &= \text{Pool}(\bar{e}_{\text{txt},l}^{i,t,x}), \\ \bar{e}_{\text{txt},l}^{i,t,x} &= \bar{e}_{\text{txt},l}^{i,x} \cdot \text{SigAffinity}(\bar{e}_{\text{vid},l}^{i,t}, \bar{e}_{\text{txt},l}^{i,x}). \end{aligned}$$

Finally, such affinity (re-)weighted text tokens are concatenated with visual tokens, as  $\bar{Z}_l^i$ , and go through an MLP.

$$Z_{l+1}^i = \bar{Z}_l^i + \text{MLP}(\bar{Z}_l^i).$$

#### 4.4. Classifier

Following  $L$  transformer layers in our Video Head, we temporally-pool all tokens. We end up with a single video embedding,  $n$  activity-text embeddings and  $m$  aux-text embeddings. We further aggregate auxiliary embeddings, leaving a single embedding per each of the  $k$  semantic categories (e.g. object, scene, human-subjects). Finally, we compute logits based on affinity, similar to the CLIP [50] objective, and use Cross-Entropy loss for optimization.

$$\begin{aligned} \text{logit}^{i,x} &= \text{Affinity}(e_{\text{vid},L}^i, e_{\text{txt},L}^{i,x}), \\ \text{logit}_{\text{aux}}^{i,x,y} &= \text{Affinity}(e_{\text{txt},L}^{i,x}, e_{\text{aux},L}^{i,y}) \Big|_{y=\{1,\dots,k\}}. \end{aligned}$$

#### 4.5. Discussion on design decisions

**Auxiliary semantic information:** We rely on optional semantics (or, attributes) in the form of visually-grounded auxiliary text, to improve our *video-conditioned text* embeddings. This is guided by the loss on  $\text{logit}_{\text{aux}}$ . The vocabulary of such auxiliary texts is fixed (i.e., common for all videos) per dataset. On Charades, we consider 97 auxiliary text classes, and on Kinetics-400, we use 88 classes (refer the appendix for more details). To highlight only the relevant semantics for a given video, we visually-ground them via (1) cross-modal attention with visual embeddings, and (2) affinity weighting. Finally, to compute  $\text{logit}_{\text{aux}}$ , we create one *representative embedding* per each of the  $k$  semantic categories, by average pooling aux embeddings within a category ( $k = 4$  for Charades and  $k = 3$  for Kinetics-400).

**Alternative weighting schemes:** Our text (re-)weighting method is similar to a contrastive training objective (as in CLIP [50]), which is based on visual-text affinities. We find this complementary nature beneficial. It highlights relevant text (and diminish irrelevant ones) within each intermediate layer of our Video Head. This iterative process fixes the initial noisy affinities resulting from the original image-VLM embeddings, when fused with better temporal cues in subsequent layers. We also explored other weighting schemes such as learnable weights or attention-based weights, which are not directly-connected to the training objective. They do not provide any improvements.

**Visual-only or Text-only classifiers:** We also explored different classifiers (i.e., how we compute logits), considering (1) a visual-only classifier as in [33], (2) a text-only classifier, or (3) an affinity-based classifier as in [40, 50]. The last performs the best. Even though we primarily focus on updating text embeddings, it still makes sense to rely on video-text affinities to be the training objective (or, classifier), as it is complementary to the components within our Video Head.

## 5. Experiments

To validate the merits of ViCTR, we experiment on few-shot and zero-shot activity recognition (on HMDB-51 [29] and UCF-101 [62]), as well as short-form (on Kinetics-400 [28]) and long-form recognition (on Charades [60]). Following sub-sections will detail our implementation, evaluation settings, datasets and the results.

**Implementation details:** We use a pretrained CLIP [50] as our image-VLM backbone. Our Video Head is randomly-initialized having 4 transformer blocks similar to [65], which is applied on-top of CLIP backbones. We consider an embedding dimension of 512/768 (w/ heads 8/12) corresponding to CLIP B/16 and L/14 backbone variants. Our output video-text embeddings are further mapped into 256-dimensional embeddings prior to computing affinity-based logits. We use an AdamW [35] optimizer with a cosine schedule for training. On Kinetics-400 [28], we finetune our model for 30 epochs with a batch size of 256 using 8e-6/8e-5 learning rates for backbone/newly-initialized parameters, similar to [40]. On Charades [60], we finetune for 50k iterations with a batch size of 64 using 5e-7/5e-4 learning rates for backbone/newly-initialized parameters, similar to [4]. We use augmentations and input sampling strategies similar to [40] for Kinetics-400 and similar to [33] for Charades.

**Evaluation settings:** In our experiments, we compare against prior art VLMs on each dataset. Since the direction of adapting image-VLMs to video is relatively-recent, their absolute performance may not be the state-of-the-art in some cases (e.g. long-form recognition), but we report numbers in comparable settings. For each experiment, we report pretraining settings, #frames-per-view, #views-at-inference and compute-per-view (GFLOPs) as supplementary metrics. We evaluate single-label activity recognition performance with Top-1 (%) accuracy, and multi-label recognition with Average Precision (mAP%). When reporting FLOPs, we consider the cost of computing a single affinity-based logit (i.e., the cost for one video-text pair) similar to [40].

### 5.1. Few-shot and Zero-shot Transfer

**Data:** We consider the downstream datasets HMDB-51 [29] and UCF-101 [62] to evaluate few-shot and zero-shot performance of our model. UCF-101 is a classification dataset collected from YouTube. It contains  $\sim 13$ k clips annotated with 101 action classes. HMDB-51 is relatively small and contains  $\sim 7$ k clips with 51 annotated classes. Both datasets have three splits of training/test data. In few-shot evaluation, we randomly sample 2, 4, 8, or 16 clips per class to create our training sets, same as in [40]. We use a model pretrained on Kinetics-400 [28] for 10 epochs and finetune on few-shot examples for 50 epochs, using 32-frames per view as in [40].

**Few-shot results:** In Table 1, we report top-1 accuracy on the first test split among three, in each dataset, using a single

Model	HMDB-51				UCF-101				
	$k$ :	2	4	8	16	2	4	8	16
<i>Methods w/o image-text pretraining</i>									
TSM [31]	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	
Video-Swin-B [34]	20.9	41.3	47.9	56.1	53.3	74.1	85.8	88.7	
<i>Methods w/ image-text pretraining</i>									
X-CLIP [40]	53.0	57.3	62.8	64.0	76.4	83.4	88.3	91.4	
X-Florence [40]	51.6	57.8	64.1	64.2	84.0	88.5	92.5	94.8	
VicTR (B/16)	60.0	63.2	66.6	70.7	87.7	92.3	93.6	95.8	

Table 1. **Zero-shot Transfer:** On HMDB-51 [29] and UCF-101 [62], we compare our method against prior art, reporting top-1 accuracy (on the first split among three test splits as in [40]). We use models pretrained on Kinetics-400 [28] for 10 epochs, and finetune on few-shot samples for 50 epochs. We randomly-sample  $k = \{2, 4, 8, 16\}$  clips per class as few-shot training samples in each setting. VicTR shows a significant boost over X-CLIP [40]. Non-VLMs are de-emphasized.

Model	#Frames	HMDB-51	UCF-101
<i>Methods w/o image-text pretraining</i>			
MTE [77]	-	19.7 $\pm$ 1.6	15.8 $\pm$ 1.3
ASR [66]	16	21.8 $\pm$ 0.9	24.4 $\pm$ 1.0
ZSECOG [49]	-	22.6 $\pm$ 1.2	15.1 $\pm$ 1.7
UR [91]	1	24.4 $\pm$ 1.6	17.5 $\pm$ 1.6
TS-GCN [16]	16	23.2 $\pm$ 3.0	34.2 $\pm$ 3.1
E2E [6]	16	32.7	48.0
ER-ZSAR [8]	-	35.3 $\pm$ 4.6	51.8 $\pm$ 2.9
<i>Methods w/ image-text pretraining</i>			
ActionCLIP [65]	32	40.8 $\pm$ 5.4	58.3 $\pm$ 3.4
X-CLIP [40]	32	44.6 $\pm$ 5.2	72.0 $\pm$ 2.3
VicTR (B/16)	32	51.0 $\pm$ 1.3	72.4 $\pm$ 0.3

Table 2. **Zero-shot Transfer:** On HMDB-51 [29] and UCF-101 [62], we compare our method against prior art, reporting input format (#Frames) and top-1 accuracy (%) as mean $\pm$ std across the three splits of test set as in [40]. We use models pretrained on Kinetics-400 for 10 epochs. VicTR outperforms similar video-VLM adaptations. Non-VLMs are de-emphasized.

view at inference. VicTR significantly outperforms prior art, either w/o image-text pretraining (TSM [31], TimeSformer [5], Video-Swin [34]) or w/ such pretraining (X-CLIP [40], X-Florence [40]). Although our method uses similar backbones as X-CLIP, it even outperforms X-Florence (an extension of a more-generic foundation model) on both datasets consistently. This shows the effectiveness of our *video-conditioned text* embeddings when generalizing to downstream with few training samples.

**Zero-shot results:** We report zero-shot transfer performance in Table 2. We use a model pretrained for 10 epochs on Kinetics-400 [28] with 32-frames per view, similar to [40], and transfer to the downstream. We report mean and standard deviation on three-splits. VicTR-B/16 outperforms X-CLIP [40] by 6.4% on HMDB-51 and by 0.4% on UCF-101. Also, the performance of our model is more stable across splits. This validates that the learned *video-conditioned text*

Model	Pretrain	#Frames	#Views	GFLOPs	Top-1
<i>Methods w/o image-text pretraining</i>					
Video-Swin-L (384 $\uparrow$ ) [34]	IN-21K	32	10 $\times$ 5	2107	84.9
TimeSformer-L [5]	IN-21K	96	1 $\times$ 3	2380	80.7
MTV-L [79]	JFT-300M	32	4 $\times$ 3	1504	84.3
Video-SwinV2-G (384 $\uparrow$ ) [34]	IN-21K+	8	4 $\times$ 5	-	86.8
MViTv2-L [30] (312 $\uparrow$ )	-	40	5 $\times$ 3	2828	86.1
ViViT-L FE [2]	JFT-300M	32	1 $\times$ 3	3980	83.5
TokenLearner [55]	JFT-300M	64	4 $\times$ 3	4076	85.4
CoVeR-L [88]	JFT-3B	-	1 $\times$ 3	-	87.2
<i>Methods w/ image-text pretraining</i>					
ST-Adapter [43]	CLIP	32	1 $\times$ 3	2749	87.2
Text4Vis [72]	CLIP	32	1 $\times$ 3	1662	87.1
EVL [33]	CLIP	8	1 $\times$ 3	674	86.3
X-CLIP [40]	CLIP	8	4 $\times$ 3	658	87.1
VicTR (L/14)	CLIP	8	4 $\times$ 3	656	87.0

Table 3. **Short-form Activity Recognition:** On Kinetics-400 [28], we compare our method against prior art, reporting pretraining settings, input format, compute cost (GFLOPs) and top-1 accuracy (%). Here, #Frames represents the number of frames per view, while #Views represents the number of temporal $\times$ spatial crops during inference. The compute cost is reported per view. VicTR shows a competitive performance among video-VLMs with a similar cost. Non-VLMs are de-emphasized.

embeddings can be generalized, even w/o seeing the same categories as in the downstream, during pretraining.

## 5.2. Short-form Activity Recognition

**Data:** Kinetics-400 [28] is a large-scale activity recognition dataset, with 240k training and 20k validation videos. Each clip contains video-level annotations for a single activity out of 400 categories, having short  $\sim$ 10s duration.

**Results:** We report the performance of VicTR on Kinetics-400 short-form activity recognition in Table 3. We consider L/14 with 8-frames per view, while using  $4 \times 3$  such views at inference similar to [40]. Our method shows a competitive performance at a similar footprint to closely-related video-VLMs [33, 40]. It is also competitive with CoVeR-L [88] which is trained with  $10\times$  more data. VicTR outperforms MTV [79] by +2.7%, ViViT [2] by +3.5% and TokenLearner [55] by +1.6%, all trained on a similar scale of data, while being more-efficient.

## 5.3. Long-form Activity Recognition

**Data:** Charades [60] is a small-yet-challenging activity recognition dataset with  $\sim$ 9.8k long-form videos. It comes with frame-level annotations of 157 daily household activities. Yet, the benchmark setting requires making video-level predictions. The data is split as  $\sim$ 7.9k for training and  $\sim$ 1.8k for validation. Each video contains multiple overlapping activities, having an average duration of  $\sim$ 30s.

**Results:** We report the performance of VicTR on Charades long-form activity recognition in Table 4. Here, we consider both B/16 and L/14 model variants with 32-frames per view, while having  $4 \times 1$  such views at inference. Our method

Model	Pretrain	#Frames	#Views	GFLOPs	mAP
<i>Methods w/o image-text pretraining</i>					
I3D + NL [67]	K400	128	10×3	544	37.5
EvaNet [46]	K400	64	-	-	38.1
LFB-101 [69]	K400	32	10×3	529	42.5
SlowFast-50 [14]	K400	8+32	10×3	66	38.0
SlowFast-101 + NL [14]	K400	16+64	10×3	234	42.5
X3D-XL (312↑) [13]	K400	16	10×3	48	43.4
MViT [12]	K400	32	10×3	237	47.7
AssembleNet-101 [57]	-	128	5×1	1200	58.6
<i>Methods w/ image-text pretraining</i>					
ActionCLIP [65]	CLIP	32	10×3	563	44.6
CLIP4clip [36]	CLIP	32	1×1	-	32.0
CLIP Hitchhiker’s [4]	CLIP	32	1×1	-	44.9
VicTR (B/16)	CLIP	32	4×1	567	50.1
VicTR (L/14)	CLIP	32	4×1	2602	57.6

Table 4. **Long-form Activity Recognition:** On Charades [60], we compare our method against prior art, reporting pretraining settings, input format, compute cost (GFLOPs) and mean Average Precision (mAP%). The compute cost reported is per view. Here, #Views represents the number of temporal×spatial crops, each having #Frames per view). VicTR achieves strong a performance among the methods pretrained w/ image-text data by a considerable margin. Non-VLMs are de-emphasized.

Model	Kinetics-400	Charades
VicTR	84.4	50.1
VicTR (No Aux. Text)	84.2	49.8
VicTR (w/ CLIP Visual emb.)	84.0	49.7
VicTR (w/ CLIP Text emb.)	83.3	41.7

Table 5. **Ablating main hypotheses:** On Kinetics-400 [28] and Charades [60], we measure the importance of auxiliary text prompts. We also show that updating text is most critical in our framework, rather than updating visual embeddings (i.e., temporally-pooled CLIP image embeddings is as good as our video embeddings).

outperforms prior video-VLMs by a considerable margin. In fact, VicTR-B/16 shows +5.2% mAP boost over CLIP Hitchhiker’s [4], and +5.5% mAP boost over ActionCLIP [65] with a similar footprint. This is a significant improvement considering the challenging Charades settings. Our method is also competitive with non-VLMs, whereas other video-VLMs lag behind. It highlights the limitations of current VLMs in long-context temporal modeling.

#### 5.4. Ablation Study

In Table 5, we provide evidence to validate our main hypotheses. Namely, we evaluate the impact of auxiliary semantics and the effectiveness of updating text embeddings.

**Auxiliary semantics do help.** We rely on extra semantic information to guide our latent embedding space. We see that such auxiliary text is giving +0.2% gain on Kinetics-400 and +0.3% mAP gain on Charades. This conveys the potential of semantics, but also the limitations of not having ground-truth annotations corresponding to them.

Model	mAP
VicTR	50.1
VicTR (No Affinity weighting)	48.8
VicTR (w/ joint-attention)	44.8
VicTR (Text Classifier)	41.2
VicTR (Visual Classifier)	43.1

Table 6. **Ablating design decisions:** On Charades [60], we evaluate different design decisions of VicTR. First, we show the effectiveness of Affinity Weighting and divided attention in our framework. We also replace our visual-text affinity-based logits with simpler visual-only or text-only logits to show the benefit of ours.

**Updating text embeddings is more effective.** To evaluate which of our embeddings (video or *video-conditioned text*) are critical, we replace them with the corresponding original CLIP [50] embeddings (i.e., temporally-pooled frame, or text). We see that the proposed *video-conditioned text* are significantly-more effective, and when replaced, the performance drops −1.1% on Kinetics-400 and −8.4% mAP on Charades. In contrast, when our video embeddings are replaced, the performance drops only −0.4% and −0.4% mAP, respectively. Meaning, the CLIP frame embeddings are on-par with our video embeddings, but our *video-conditioned text* embeddings are significantly improved.

In Table 6, we ablate and justify our design decisions. Namely, we evaluate our affinity weighting mechanism, divided attention, and affinity-based classifier.

**Affinity-weighting and divided attention do help.** We see a +1.3% mAP performance gain by having our affinity (re-)weighting mechanism. While joint-attention may be more expressive compared to divided attention, it can incur training difficulties. As a result, we see the divided attention enjoying a significant +5.3% mAP boost.

**Affinity-based classifier is required.** As we previously discussed, our affinity weighting mechanism makes more-sense in the context of the same affinity-based loss formulation. To verify this, we replace such affinity-based logits with text-only or visual-only logits, which are just linear mappings of the corresponding embeddings. These significantly underperforms, with −8.9% mAP and −7.0% mAP, respectively.

## 6. Conclusion

In this paper, we introduced VicTR, a framework for adapting image-VLMs to video, with a focus on *video-conditioned text* embeddings. It can also benefit from freely-available auxiliary semantic information in the form of visually-grounded text, to guide the learned latent space. Our evaluations verified the importance of updating text embeddings, across multiple activity recognition benchmarks, under few-shot, zero-shot, short-form and long-form settings. We believe that this work reveals the importance of using language embeddings for temporal reasoning.



## References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a Visual Language Model for Few-Shot Learning. *NeurIPS*, 2022. 1, 3
- [2] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A Video Vision Transformer. In *ICCV*, pages 6836–6846, 2021. 1, 2, 5, 7
- [3] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in Time: A Joint Video and Image Encoder for End-to-End Retrieval. In *ICCV*, pages 1728–1738, 2021. 3
- [4] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. A CLIP-Hitchhiker’s Guide to Long Video Retrieval. *arXiv preprint arXiv:2205.08508*, 2022. 2, 3, 4, 6, 8
- [5] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is Space-Time Attention All You Need for Video Understanding? In *ICML*, page 4, 2021. 1, 2, 5, 7
- [6] Biagio Brattoli, Joseph Tighe, Fedor Zhdanov, Pietro Perona, and Krzysztof Chalupka. Rethinking Zero-shot Video Classification: End-to-end Training for Realistic Applications. In *CVPR*, pages 4613–4623, 2020. 7
- [7] Joao Carreira and Andrew Zisserman. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *CVPR*, pages 6299–6308, 2017. 1, 2
- [8] Shizhe Chen and Dong Huang. Elaborative Rehearsal for Zero-shot Action Recognition. In *ICCV*, pages 13638–13647, 2021. 7
- [9] Feng Cheng, Xizi Wang, Jie Lei, David Crandall, Mohit Bansal, and Gedas Bertasius. VindLU: A Recipe for Effective Video-and-Language Pretraining. *arXiv preprint arXiv:2212.05051*, 2022. 3
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *ICLR*, 2021. 3
- [11] Victor Escorcia, Fabian Caba Heilbron, Juan Carlos Niebles, and Bernard Ghanem. DAPs: Deep Action Proposals for Action Understanding. In *ECCV*, pages 768–784. Springer, 2016. 2
- [12] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale Vision Transformers. In *ICCV*, pages 6824–6835, 2021. 1, 8
- [13] Christoph Feichtenhofer. X3D: Expanding Architectures for Efficient Video Recognition. In *CVPR*, pages 203–213, 2020. 1, 2, 8
- [14] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. SlowFast Networks for Video Recognition. In *ICCV*, pages 6202–6211, 2019. 1, 2, 8
- [15] Christoph Feichtenhofer, Haoqi Fan, Bo Xiong, Ross Girshick, and Kaiming He. A Large-Scale Study on Unsupervised Spatiotemporal Representation Learning. In *CVPR*, pages 3299–3309, 2021. 1, 2
- [16] Junyu Gao, Tianzhu Zhang, and Changsheng Xu. I Know the Relationships: Zero-Shot Action Recognition via Two-Stream Graph Convolutional Networks and Knowledge Graphs. In *AAAI*, pages 8303–8311, 2019. 7
- [17] Chunhui Gu, Chen Sun, David A Ross, Carl Vondrick, Caroline Pantofaru, Yeqing Li, Sudheendra Vijayanarasimhan, George Toderici, Susanna Ricco, Rahul Sukthankar, et al. AVA: A Video Dataset of Spatio-temporally Localized Atomic Visual Actions. In *CVPR*, pages 6047–6056, 2018. 2
- [18] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary Object Detection via Vision and Language Knowledge Distillation. *ICLR*, 2021. 1, 2, 3
- [19] Ziyu Guo, Renrui Zhang, Longtian Qiu, Xianzheng Ma, Xupeng Miao, Xuming He, and Bin Cui. CALIP: Zero-Shot Enhancement of CLIP with Parameter-free Attention. *AAAI*, 2023. 3
- [20] Tengda Han, Weidi Xie, and Andrew Zisserman. Self-supervised Co-training for Video Representation Learning. *NeurIPS*, 33:5679–5690, 2020. 1, 2
- [21] Po-Yao Huang, Vasu Sharma, Hu Xu, Chaitanya Ryali, Haoqi Fan, Yanghao Li, Shang-Wen Li, Gargi Ghosh, Jitendra Malik, and Christoph Feichtenhofer. MAVIL: Masked Audio-Video Learners. *arXiv preprint arXiv:2212.08071*, 2022. 2
- [22] Jingwei Ji, Ranjay Krishna, Li Fei-Fei, and Juan Carlos Niebles. Action Genome: Actions as Composition of Spatio-temporal Scene Graph. In *CVPR*, pages 10236–10247, 2020. 2
- [23] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc 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. PMLR, 2021. 1, 3
- [24] Yunfan Jiang, Agrim Gupta, Zichen Zhang, Guanzhi Wang, Yongqiang Dou, Yanjun Chen, Li Fei-Fei, Anima Anandkumar, Yuke Zhu, and Linxi Fan. VIMA: General Robot Manipulation with Multimodal Prompts. *arXiv preprint arXiv:2210.03094*, 2022. 1, 3
- [25] Chen Ju, Tengda Han, Kunhao Zheng, Ya Zhang, and Weidi Xie. Prompting Visual-Language Models for Efficient Video Understanding. In *ECCV*, pages 105–124. Springer, 2022. 3
- [26] Kumara Kahatapitiya and Michael S Ryoo. Coarse-Fine Networks for Temporal Activity Detection in Videos. In *CVPR*, pages 8385–8394, 2021. 2
- [27] Siddharth Karamcheti, Suraj Nair, Annie S Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh, and Percy Liang. Language-Driven Representation Learning for Robotics. *arXiv preprint arXiv:2302.12766*, 2023. 3
- [28] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The Kinetics Human Action Video Dataset. *arXiv preprint arXiv:1705.06950*, 2017. 2, 6, 7, 8
- [29] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. HMDB: A large video database for human motion recognition. In *ICCV*, pages 2556–2563. IEEE, 2011. 2, 6, 7
- [30] 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 *CVPR*, pages 4804–4814, 2022. 7
- [31] Ji Lin, Chuang Gan, and Song Han. TSM: Temporal Shift Module for Efficient Video Understanding. In *ICCV*, pages 7083–7093, 2019. 1, 3, 7
- [32] Kevin Qinghong Lin, Alex Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Zhongcong Xu, Difei Gao, Rongcheng Tu, Wenzhe Zhao, Weijie Kong, et al. Egocentric Video-Language Pretraining. *NeurIPS*, 2022. 3
- [33] Ziyi Lin, Shijie Geng, Renrui Zhang, Peng Gao, Gerard de Melo, Xiaogang Wang, Jifeng Dai, Yu Qiao, and Hongsheng Li. Frozen CLIP Models are Efficient Video Learners. *arXiv preprint arXiv:2208.03550*, 2022. 2, 3, 4, 6, 7
- [34] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video Swin Transformer. In *CVPR*, pages 3202–3211, 2022. 2, 7
- [35] Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization. *ICLR*, 2019. 6
- [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. *Neurocomputing*, 508:293–304, 2022. 2, 3, 4, 8
- [37] Sachit Menon and Carl Vondrick. Visual Classification via Description from Large Language Models. *arXiv preprint arXiv:2210.07183*, 2022. 3
- [38] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple Open-Vocabulary Object Detection with Vision Transformers. *ECCV*, 2022. 1, 3
- [39] Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. Attention Bottlenecks for Multimodal Fusion. *NeurIPS*, 34:14200–14213, 2021. 1, 2
- [40] Bolin Ni, Houwen Peng, Minghao Chen, Songyang Zhang, Gaofeng Meng, Jianlong Fu, Shiming Xiang, and Haibin Ling. Expanding Language-Image Pretrained Models for General Video Recognition. In *ECCV*, pages 1–18. Springer, 2022. 2, 3, 4, 6, 7
- [41] David Nukrai, Ron Mokady, and Amir Globerson. Text-Only Training for Image Captioning using Noise-Injected CLIP. *EMNLP*, 2022. 3
- [42] Roni Paiss, Ariel Ephrat, Omer Tov, Shiran Zada, Inbar Mosseri, Michal Irani, and Tali Dekel. Teaching CLIP to Count to Ten. *arXiv preprint arXiv:2302.12066*, 2023. 3
- [43] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. ST-Adapter: Parameter-Efficient Image-to-Video Transfer Learning. *NeurIPS*, 2022. 7
- [44] AJ Piergiovanni and Michael Ryoo. Temporal Gaussian Mixture Layer for Videos. In *ICML*, pages 5152–5161. PMLR, 2019. 2
- [45] AJ Piergiovanni and Michael S Ryoo. Learning Latent Super-Events to Detect Multiple Activities in Videos. In *CVPR*, pages 5304–5313, 2018. 2
- [46] AJ Piergiovanni, Anelia Angelova, Alexander Toshev, and Michael S Ryoo. Evolving Space-Time Neural Architectures for Videos. In *ICCV*, pages 1793–1802, 2019. 8
- [47] Rui Qian, Tianjian Meng, Boqing Gong, Ming-Hsuan Yang, Huisheng Wang, Serge Belongie, and Yin Cui. Spatiotemporal Contrastive Video Representation Learning. In *CVPR*, pages 6964–6974, 2021. 1, 2
- [48] Rui Qian, Yeqing Li, Zheng Xu, Ming-Hsuan Yang, Serge Belongie, and Yin Cui. Multimodal Open-Vocabulary Video Classification via Pre-Trained Vision and Language Models. *arXiv preprint arXiv:2207.07646*, 2022. 3
- [49] Jie Qin, Li Liu, Ling Shao, Fumin Shen, Bingbing Ni, Jiaxin Chen, and Yunhong Wang. Zero-Shot Action Recognition with Error-Correcting Output Codes. In *CVPR*, pages 2833–2842, 2017. 7
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, pages 8748–8763. PMLR, 2021. 1, 3, 4, 6, 8
- [51] Kanchana Ranasinghe, Brandon McKinzie, Sachin Ravi, Yinfei Yang, Alexander Toshev, and Jonathon Shlens. Perceptual Grouping in Vision-Language Models. *arXiv preprint arXiv:2210.09996*, 2022. 3
- [52] Hanoona Rasheed, Muhammad Uzair Khattak, Muhammad Maaz, Salman Khan, and Fahad Shahbaz Khan. Fine-tuned CLIP Models are Efficient Video Learners. *arXiv preprint arXiv:2212.03640*, 2022. 3
- [53] Adria Recasens, Pauline Luc, Jean-Baptiste Alayrac, Luyu Wang, Florian Strub, Corentin Tallec, Mateusz Malinowski, Viorica Pătrăucean, Florent Alché, Michal Valko, et al. Broaden Your Views for Self-Supervised Video Learning. In *ICCV*, pages 1255–1265, 2021. 1, 2
- [54] Adria Recasens, Jason Lin, João Carreira, Drew Jaegle, Luyu Wang, Jean-baptiste Alayrac, Pauline Luc, Antoine Miech, Lucas Smaira, Ross Hemsley, et al. Zorro: the masked multimodal transformer. *arXiv preprint arXiv:2301.09595*, 2023. 2
- [55] Michael Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. TokenLearner: Adaptive Space-Time Tokenization for Videos. *NeurIPS*, 34:12786–12797, 2021. 2, 7
- [56] Michael S Ryoo, AJ Piergiovanni, Juhana Kangaspunta, and Anelia Angelova. AssembleNet++: Assembling Modality Representations via Attention Connections. In *ECCV*, pages 654–671. Springer, 2020. 2
- [57] Michael S Ryoo, AJ Piergiovanni, Mingxing Tan, and Anelia Angelova. AssembleNet: Searching for Multi-Stream Neural Connectivity in Video Architectures. *ICLR*, 2020. 1, 2, 8
- [58] Michael S Ryoo, Keerthana Gopalakrishnan, Kumara Kahatapitiya, Ted Xiao, Kanishka Rao, Austin Stone, Yao Lu, Julian Ibarz, and Anurag Arnab. Token Turing Machines. *arXiv preprint arXiv:2211.09119*, 2022. 2
- [59] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural Machine Translation of Rare Words with Subword Units. *ACL*, 2016. 3
- [60] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in Homes: Crowdsourcing Data Collection for Activity Under-

- standing. In *ECCV*, pages 510–526. Springer, 2016. [2](#), [6](#), [7](#), [8](#)
- [61] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. FLAVA: A Foundational Language And Vision Alignment Model. In *CVPR*, pages 15638–15650, 2022. [1](#), [3](#)
- [62] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. UCF101: A Dataset of 101 Human Actions Classes From Videos in The Wild. *arXiv preprint arXiv:1212.0402*, 2012. [6](#), [7](#)
- [63] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A Closer Look at Spatiotemporal Convolutions for Action Recognition. In *CVPR*, pages 6450–6459, 2018. [1](#), [2](#)
- [64] Michael Tschannen, Basil Mustafa, and Neil Houlsby. Image-and-Language Understanding from Pixels Only. *arXiv preprint arXiv:2212.08045*, 2022. [3](#)
- [65] Mengmeng Wang, Jiazheng Xing, and Yong Liu. Action-CLIP: A New Paradigm for Video Action Recognition. *arXiv preprint arXiv:2109.08472*, 2021. [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [66] Qian Wang and Ke Chen. Alternative Semantic Representations for Zero-Shot Human Action Recognition. In *ECML-PKDD*, pages 87–102. Springer, 2017. [7](#)
- [67] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local Neural Networks. In *CVPR*, pages 7794–7803, 2018. [2](#), [8](#)
- [68] Zhenhailong Wang, Manling Li, Ruochen Xu, Luwei Zhou, Jie Lei, Xudong Lin, Shuohang Wang, Ziyi Yang, Chenguang Zhu, Derek Hoiem, et al. Language Models with Image Descriptors are Strong Few-Shot Video-Language Learners. *arXiv preprint arXiv:2205.10747*, 2022. [2](#), [3](#)
- [69] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-Term Feature Banks for Detailed Video Understanding. In *CVPR*, pages 284–293, 2019. [8](#)
- [70] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-Term Feature Banks for Detailed Video Understanding. In *CVPR*, pages 284–293, 2019. [2](#)
- [71] Chao-Yuan Wu, Yanghao Li, Karttikeya Mangalam, Haoqi Fan, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. MeMViT: Memory-Augmented Multiscale Vision Transformer for Efficient Long-Term Video Recognition. In *CVPR*, pages 13587–13597, 2022. [2](#)
- [72] Wenhao Wu, Zhun Sun, and Wanli Ouyang. Revisiting Classifier: Transferring Vision-Language Models for Video Recognition. *AAAI*, 2023. [7](#)
- [73] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking Spatiotemporal Feature Learning For Video Understanding. *arXiv preprint arXiv:1712.04851*, 1(2):5, 2017. [1](#), [2](#)
- [74] 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. *EMNLP*, 2021. [2](#), [3](#)
- [75] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *CVPR*, pages 5288–5296, 2016. [2](#)
- [76] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. GroupViT: Semantic Segmentation Emerges from Text Supervision. In *CVPR*, pages 18134–18144, 2022. [3](#)
- [77] Xun Xu, Timothy M Hospedales, and Shaogang Gong. Multi-Task Zero-Shot Action Recognition with Prioritised Data Augmentation. In *ECCV*, pages 343–359. Springer, 2016. [7](#)
- [78] Hongwei Xue, Yuchong Sun, Bei Liu, Jianlong Fu, Ruihua Song, Houqiang Li, and Jiebo Luo. CLIP-ViP: Adapting Pre-trained Image-Text Model to Video-Language Representation Alignment. *arXiv preprint arXiv:2209.06430*, 2022. [3](#)
- [79] Shen Yan, Xuehan Xiong, Anurag Arnab, Zhichao Lu, Mi Zhang, Chen Sun, and Cordelia Schmid. Multiview Transformers for Video Recognition. In *CVPR*, pages 3333–3343, 2022. [7](#)
- [80] Shen Yan, Tao Zhu, Zirui Wang, Yuan Cao, Mi Zhang, Soham Ghosh, Yonghui Wu, and Jiahui Yu. Video-Text Modeling with Zero-Shot Transfer from Contrastive Captioners. *arXiv preprint arXiv:2212.04979*, 2022. [1](#), [3](#)
- [81] Antoine Yang, Arsha Nagrani, Paul Hongsuck Seo, Antoine Miech, Jordi Pont-Tuset, Ivan Laptev, Josef Sivic, and Cordelia Schmid. Vid2Seq: Large-Scale Pretraining of a Visual Language Model for Dense Video Captioning. *arXiv preprint arXiv:2302.14115*, 2023. [3](#)
- [82] 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. [1](#), [3](#)
- [83] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every Moment Counts: Dense Detailed Labeling of Actions in Complex Videos. *IJCV*, 126:375–389, 2018. [2](#)
- [84] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. CoCa: Contrastive Captioners are Image-Text Foundation Models. *arXiv preprint arXiv:2205.01917*, 2022. [3](#)
- [85] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A New Foundation Model for Computer Vision. *arXiv preprint arXiv:2111.11432*, 2021. [1](#), [3](#)
- [86] Andy Zeng, Adrian Wong, Stefan Welker, Krzysztof Choromanski, Federico Tombari, Aveek Purohit, Michael Ryoo, Vikas Sindhwani, Johnny Lee, Vincent Vanhoucke, et al. Socratic Models: Composing Zero-Shot Multimodal Reasoning with Language. *arXiv preprint arXiv:2204.00598*, 2022. [1](#), [2](#), [3](#)
- [87] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. LiT: Zero-Shot Transfer with Locked-image text Tuning. In *CVPR*, pages 18123–18133, 2022. [1](#), [2](#)
- [88] Bowen Zhang, Jiahui Yu, Christopher Fifty, Wei Han, Andrew M Dai, Ruoming Pang, and Fei Sha. Co-training Trans-

- former with Videos and Images Improves Action Recognition. *arXiv preprint arXiv:2112.07175*, 2021. 7
- [89] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning Video Representations from Large Language Models. *arXiv preprint arXiv:2212.04501*, 2022. 3
- [90] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to Prompt for Vision-Language Models. *IJCV*, 130(9):2337–2348, 2022. 3
- [91] Yi Zhu, Yang Long, Yu Guan, Shawn Newsam, and Ling Shao. Towards Universal Representation for Unseen Action Recognition. In *CVPR*, pages 9436–9445, 2018. 7