# Bidirectional Cross-Modal Knowledge Exploration for Video Recognition with Pre-trained Vision-Language Models

Wenhao Wu[1,2]    Xiaohan Wang[3]    Haipeng Luo[4]    Jingdong Wang[2]    Yi Yang[3]    Wanli Ouyang[5,1]

[1]The University of Sydney        [2]Baidu Inc.        [3]Zhejiang University
[4]University of Chinese Academy of Sciences        [5]Shanghai AI Laboratory

whwu.ucas@gmail.com

## Abstract

*Vision-language models (VLMs) pre-trained on large-scale image-text pairs have demonstrated impressive transferability on various visual tasks. Transferring knowledge from such powerful VLMs is a promising direction for building effective video recognition models. However, current exploration in this field is still limited. We believe that the greatest value of pre-trained VLMs lies in building a bridge between visual and textual domains. In this paper, we propose a novel framework called **BIKE**, which utilizes the cross-modal bridge to explore bidirectional knowledge: i) We introduce the Video Attribute Association mechanism, which leverages the* `Video-to-Text` *knowledge to generate textual auxiliary attributes for complementing video recognition. ii) We also present a Temporal Concept Spotting mechanism that uses the* `Text-to-Video` *expertise to capture temporal saliency in a parameter-free manner, leading to enhanced video representation. Extensive studies on six popular video datasets, including Kinetics-400 & 600, UCF-101, HMDB-51, ActivityNet and Charades, show that our method achieves state-of-the-art performance in various recognition scenarios, such as general, zero-shot, and few-shot video recognition. Our best model achieves a state-of-the-art accuracy of 88.6% on the challenging Kinetics-400 using the released CLIP model. The code is available at* https://github.com/whwu95/BIKE.

## 1. Introduction

In recent years, the remarkable success of large-scale pre-training in NLP (*e.g.*, BERT [9], GPT [4, 39], ERNIE [71] and T5 [40]) has inspired the computer vision community. Vision-language models (VLMs) leverage large-scale noisy image-text pairs with weak correspondence for contrastive learning (*e.g.*, CLIP [38], ALIGN [19], CoCa [67], Florence [68]), and demonstrate impressive transferability across a wide range of visual tasks.



(a) Traditional video recognition.

(b) Category embedding as classifier for video recognition.

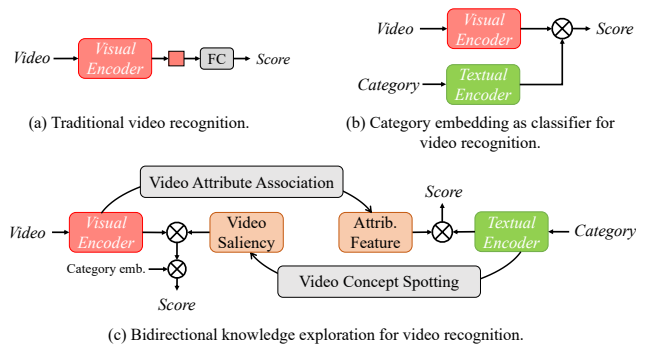(c) Bidirectional knowledge exploration for video recognition.

Figure 1. Illustration of the difference between our paradigm (c) with existing unimodality paradigm (a) and cross-modal paradigm (b). Please zoom in for the best view.

Naturally, transferring knowledge from such powerful pre-trained VLMs is emerging as a promising paradigm for building video recognition models. Currently, exploration in this field can be divided into two lines. As depicted in Figure 1(a), one approach [28, 36, 66] follows the traditional unimodal video recognition paradigm, initializing the video encoder with the pre-trained visual encoder of VLM. Conversely, the other approach [21, 35, 50, 60] directly transfers the entire VLM into a video-text learning framework that utilizes natural language (*i.e.*, class names) as supervision, as shown in Figure 1(b). This leads to an question: *have we fully utilized the knowledge of VLMs for video recognition?*

In our opinion, the answer is No. The greatest charm of VLMs is their ability to build a bridge between the visual and textual domains. Despite this, previous research employing pre-aligned vision-text features of VLMs for video recognition has only utilized unidirectional video-to-text matching. In this paper, we aim to facilitate bidirectional knowledge exploration through the cross-modal bridge for enhanced video recognition. With this in mind, we mine `Video-to-Text` and `Text-to-Video` knowledge by 1) generating textual information from the input video and

2) utilizing category descriptions to extract valuable video-related signals.

In the first `Video-to-Text` direction, a common practice for mining VLM knowledge is to embed the input video and category description into a pre-aligned feature space, and then select the category that is closest to the video, as illustrated in Figure 1(b), which serves as our baseline. One further question naturally arises: *Can we incorporate auxiliary textual information for video recognition?* To address this question, we introduce an **Video-Attributes Association** mechanism, which leverages the zero-shot capability of VLMs to retrieve the most relevant phrases from a pre-defined lexicon for the video. These phrases are considered potential "attributes" of the video and can predict the video category directly. For example, a video of someone kicking a soccer ball may be associated with relevant phrases such as "running on the grass", "juggling soccer ball" and "shooting goal". Surprisingly, using only the generated attributes, we can achieve 69% top-1 accuracy on the challenging Kinetics-400 dataset. Furthermore, these attributes provide additional information that the video visual signal may not capture, allowing us to build an *Attributes Recognition Branch* for video recognition.

In the second `Text-to-Video` direction, we believe that temporal saliency in videos can be leveraged to improve video representations. For instance, in a video with the category "kicking soccer ball", certain frames of kicking the ball should have higher saliency, while other frames that are unrelated to the category or background frames should have lower saliency. This insight motivates us to propose the **Video Concept Spotting** mechanism, which utilizes the cross-model bridge to generate category-dependent temporal saliency. In previous works [35, 50, 60], this intuitive exploration was disregarded. To be more specific, instead of treating each video frame equally, we use the correlation between each frame and the given concept (*e.g.*, category) as a measure of frame-level saliency. This saliency is then used to temporally aggregate the frames, resulting in a compact video representation.

In the light of the above explorations, we propose **BIKE**, a simple yet effective framework via **BI**directional cross-modal **K**nowledge **E**xploration for enhanced video recognition. Our **BIKE** comprises two branches: the *Attributes branch*, which utilizes the **Video-Attributes Association** mechanism to introduce auxiliary attributes for complementary video recognition, and the *Video branch*, which uses the **Video Concept Spotting** mechanism to introduce temporal saliency to enhance video recognition. To demonstrate the effectiveness of our **BIKE**, we conduct comprehensive experiments on popular video datasets, including Kinetics-400 [22] & 600 [6], UCF-101 [44], HMDB-51 [24], ActivityNet [5] and Charades [42]. The results show that our method achieves state-of-the-art performance in most scenarios, *e.g.*, general, zero-shot, and few-shot recognition.

Our main contributions can be summarized as follows:

- We propose a novel framework called **BIKE** that explores bidirectional knowledge from pre-trained vision-language models for video recognition.

- In the `Video-to-Text` direction, we introduce the **Video-Attributes Association** mechanism to generate extra attributes for complementary video recognition.

- In the `Text-to-Video` direction, we introduce the **Video Concept Spotting** mechanism to generate temporal saliency, which is used to yield the compact video representation for enhanced video recognition.

## 2. Methodology

An overview of our proposed **BIKE** is shown in Figure 2. We next elaborate on each component in more detail.

### 2.1. Preliminary: Video Recognition with VLM

In this section, we describe the typical cross-modal video recognition pipeline [21, 35, 50, 60] based on the pre-trained vision-language model (VLM). Given a video, we sample $T$ frames from the video as input $v$. We also have a collection of categories $C = \{c_1, c_2, \cdots, c_K\}$, where $K$ is the number of classes. The goal of the video recognition task is to classify the video $v$ into a category $c \in C$. Under the formulation of video recognition, the video $v$ is encoded with a vision encoder $f(\cdot|\theta_v)$ to obtain the video embedding $\mathbf{e_v}$, and the category $c$ is encoded with a text encoder $g(\cdot|\phi_c)$ to obtain the category embedding $\mathbf{e_c}$, where

$$\mathbf{e_v} = f(v|\theta_v), \mathbf{e_c} = g(c|\phi_c). \quad (1)$$

Finally, we obtain the similarity score $\mathcal{S}_V$ as follows:

$$\mathcal{S}_V = s(\mathbf{e_v}, \mathbf{e_c}), \quad (2)$$

where $s(\cdot, \cdot)$ is the cosine similarity function. The objective during training is to maximize $\mathcal{S}_V$ if $v$ and $c$ are matched, and minimize it in all other cases. During inference, we compute the score between the video embedding and each category embedding, and choose the category with the highest $\mathcal{S}_V$ as the top-1 prediction. The parameter $\theta_v$ and $\phi_c$ of the video encoder and text encoder are initialized with weights from the pre-trained VLM (*e.g.*, CLIP [38]). Throughout the rest of this work, we use the same notation.

### 2.2. Video-to-Text: Video-Attributes Association

First we focus on exploring `Video-to-Text` auxiliary signals. We present an *Attributes branch* as a complement to the regular *Video branch* in Sec. 2.1 for video recognition. **Pre-generated Attributes.** We begin by describing how to generate auxiliary attributes. As depicted in Figure 2(b), we

Figure 2. An overview of our **BIKE** for video recognition. (a) **BIKE** explores bidirectional cross-modal knowledge from the pre-trained vision-language model (*e.g.*, CLIP) to introduce auxiliary attributes and category-dependent temporal saliency for improved video recognition. **BIKE** comprises an auxiliary *Attributes branch* and a main *Video branch*. (b) In the `Video-to-Text` direction, we present the **Video-Attribute Association** mechanism, which retrieves semantically relevant phrases from a pre-defined lexicon as video attributes for the input video. These attributes are concatenated and combined with a textual prefix to form an attribute sentence for text recognition. (c) In the `Text-to-Video` direction, we present the **Video Concept Spotting** mechanism, which computes the similarity between video frames and a given category as a measure of temporal saliency to enhance video representation. $D$ is the dimension of embedding, $T$ is the number of frames, and $N$ is the number of words in the category name.

utilize the zero-shot capability of the VLM (*e.g.*, CLIP [38]) to identify the most relevant phases from a pre-defined lexicon as possible *"Attributes"* of the video. To achieve this, we first apply the CLIP's image encoder to the input video $V$ to extract frame-level features. These features are then combined using average pooling to yield a video embedding. Next, we feed each phase in the pre-defined lexicon into the CLIP's text encoder to produce a set of text embeddings. We then calculate the similarity between this video embedding and each text embedding, sort the results, and select the top few phrases as the *"Attributes"*. Once we have obtained the attributes, we employ a simple fusion method that concatenates them into a single attributes sentence $a$. We also add a manually-designed prompt as a prefix to the sentence, such as "This is a video about {}".

**Attributes Recognition.** As shown in Figure 2(a), the attributes sentence $a$ is encoded with a text encoder $g(\cdot|\phi_a)$ to produce the attribute embedding $\mathbf{e_a}$:

$$\mathbf{e_a} = g(a|\phi_a). \qquad (3)$$

We use this attribute embedding to perform *Attributes Recognition* by calculating the similarity $\mathcal{S}_A$ between the attribute embedding and category embeddings. Note that both the attribute sentence and categories are encoded using the same text encoder from CLIP. Interestingly, the Attributes branch can achieve a certain level of recognition performance (*e.g.*, ∼56%) without any extra training, even though it's a lightweight text recognition pipeline. During inference, we combine the well-trained Video branch with the plug-and-play Attributes branch using the following fusion equation:

$$\mathcal{S} = \lambda \mathcal{S}_V + (1 - \lambda)\mathcal{S}_A, \qquad (4)$$

where $\lambda$ is the fusion weight. Without any additional training, the *Attributes Recognition* surprisingly improve the video recognition performance, *e.g.*, 78.8% $\xrightarrow{+1.2\%}$ 80.0% on the challenging Kinetics-400. Naturally, the text encoder $g(\cdot|\phi_a)$ can be further trained in an end-to-end manner to improve the *Attributes branch* and provide a stronger complementary capability, *e.g.*, 78.8% $\xrightarrow{+2.6\%}$ 81.4%.

3

## 2.3. Text-to-Video: Video Concept Spotting

In Sec. 2.2, the `Video-to-Text` knowledge is employed to generate auxiliary attributes, thereby constructing a complementary *Attributes branch*. Naturally, we also conduct an exploration to leverage the `Text-to-Video` knowledge to enhance the standard *Video branch* for video recognition. Specifically, we propose the use of category-dependent temporal saliency to guide the temporal aggregation process, resulting in a compact video representation that enhances video recognition.

**Background.** To obtain a video representation based on a pre-trained image model, the typical pipeline involves two stages. First, we employ the image model to extract the spatial embedding of each frame. Next, the embeddings of these frames are temporally aggregated (*e.g.*, mean pooling) to yield a video-level representation.

**Parameter-Free Video Concept Spotting.** Mean pooling is a widely used technique to aggregate the frame embeddings and obtain the final video representation. Instead of treating each video frame equally as in mean pooling, we propose a parameter-free solution that utilizes the pre-aligned visual and textual semantics offered by the VLM (*e.g.*, CLIP [38]) to capture temporal saliency for video feature aggregation, as illustrated in Figure 2(c). To estimate temporal saliency, we employ word embeddings as the query to obtain finer word-to-frame saliency. Formally, the pre-trained VLM can encode each video or category name separately, and output two sets of embeddings: $\{\mathbf{v}_t \in \mathbb{R}^d | t = 1, 2, \cdots, T\}$ is a set of frame embeddings, where $T$ is the number of sampled frames, and $\{\mathbf{t}_n \in \mathbb{R}^d | n = 1, 2, \cdots, N\}$ is a set of word embeddings, where $N$ is the number of words in the class name. We calculate the similarity between each word and each frame to measure the fine-grained relevancy. After that, we perform a softmax operation to normalize the similarities for each frame, and then aggregate the similarities between a certain frame and different words to obtain a frame-level saliency.

$$\mathcal{S}_t = \frac{1}{N} \sum_{n=1}^{N} \frac{\exp(\mathbf{v}_t^{\mathsf{T}} \mathbf{t}_n / \tau)}{\sum_{t=1}^{T} \exp(\mathbf{v}_t^{\mathsf{T}} \mathbf{t}_n / \tau)}, t \in [1, T], n \in [1, N], \tag{5}$$

where $\tau$ is the temperature of this softmax function. *See Figure 3 for the visualization of temporal saliency.* Next, we utilize the temporal saliency to aggregate these frame embeddings as follows:

$$\mathbf{e}_\mathbf{v} = \sum_{t=1}^{T} \mathbf{v}_t \mathcal{S}_t, \tag{6}$$

$\mathbf{e}_\mathbf{v} \in \mathbb{R}^d$ is the final enhanced video representation.

## 2.4. Objectives of BIKE

We present the **BIKE** learning framework for video recognition, as depicted in Figure 2(a). Formally, our **BIKE** extracts feature representations $\mathbf{e}_\mathbf{v}$, $\mathbf{e}_\mathbf{a}$, and $\mathbf{e}_\mathbf{c}$ for a given video $v$, pre-generated attributes $a$, and category $c$ with the corresponding encoders $f(\cdot|\theta_v)$, $g(\cdot|\phi_a)$, and $g(\cdot|\phi_c)$. Model parameters $\theta_v$, $\theta_a$, and $\theta_c$ are initialized with the weights from the pre-trained VLM (*e.g.*, CLIP [38]). In this paper, we freeze the parameters of the pre-trained text encoder for $g(\cdot|\phi_c)$ and design extra manual prompts for the category $c$ and attributes sentence $a$.

During the training phase, our objective is to ensure that the video representation $\mathbf{e}_\mathbf{v}$ and the category representation $\mathbf{e}_\mathbf{c}$ are similar when they are related and dissimilar when they are not, and the same applies to the attributes-category pairs. Given a batch of $B$ quadruples $\{\mathbf{e}_{\mathbf{v}i}, \mathbf{e}_{\mathbf{a}i}, \mathbf{e}_{\mathbf{c}i} \equiv C[y_i], y_i\}_{i=1}^{B}$, where $C$ is the collection of $K$ categories indexed by $y_i \in [0, K-1]$ and $y_i$ is a label indicating the index of the category in the dataset, and $\mathbf{e}_{\mathbf{v}i}$, $\mathbf{e}_{\mathbf{a}i}$, $\mathbf{e}_{\mathbf{c}i}$ denote the $i$-th video embedding, attributes embedding, and category embedding, respectively. We follow the common practice [21,50] to consider the bidirectional learning objective and employ symmetric cross-entropy loss to maximize the similarity between matched *Video-Category* pairs and minimize the similarity for other pairs:

$$\mathcal{L}_{V2C} = -\frac{1}{B} \sum_{i}^{B} \frac{1}{|\mathcal{K}(i)|} \sum_{k \in \mathcal{K}(i)} \log \frac{\exp(s(\mathbf{e}_{\mathbf{c}i}, \mathbf{e}_{\mathbf{v}k})/\tau)}{\sum_{j}^{B} \exp(s(\mathbf{e}_{\mathbf{c}i}, \mathbf{e}_{\mathbf{v}j})/\tau)},$$

$$\mathcal{L}_{C2V} = -\frac{1}{B} \sum_{i}^{B} \frac{1}{|\mathcal{K}(i)|} \sum_{k \in \mathcal{K}(i)} \log \frac{\exp(s(\mathbf{e}_{\mathbf{c}k}, \mathbf{e}_{\mathbf{v}i})/\tau)}{\sum_{j}^{B} \exp(s(\mathbf{e}_{\mathbf{c}j}, \mathbf{e}_{\mathbf{v}i})/\tau)},$$

$$\mathcal{L}_V = \frac{1}{2}(\mathcal{L}_{V2C} + \mathcal{L}_{C2V}), \tag{7}$$

where $k \in \mathcal{K}(i) = \{k | k \in [1, B], y_k = y_i\}$, $s(\cdot, \cdot)$ is the cosine similarity, and $\tau$ refers to the temperature hyperparameter for scaling. Similarly, the loss for *Attributes branch* is formulated as:

$$\mathcal{L}_{A2C} = -\frac{1}{B} \sum_{i}^{B} \frac{1}{|\mathcal{K}(i)|} \sum_{k \in \mathcal{K}(i)} \log \frac{\exp(s(\mathbf{e}_{\mathbf{c}i}, \mathbf{e}_{\mathbf{a}k})/\tau)}{\sum_{j}^{B} \exp(s(\mathbf{e}_{\mathbf{c}i}, \mathbf{e}_{\mathbf{a}j})/\tau)},$$

$$\mathcal{L}_{C2A} = -\frac{1}{B} \sum_{i}^{B} \frac{1}{|\mathcal{K}(i)|} \sum_{k \in \mathcal{K}(i)} \log \frac{\exp(s(\mathbf{e}_{\mathbf{c}k}, \mathbf{e}_{\mathbf{a}i})/\tau)}{\sum_{j}^{B} \exp(s(\mathbf{e}_{\mathbf{c}j}, \mathbf{e}_{\mathbf{a}i})/\tau)},$$

$$\mathcal{L}_A = \frac{1}{2}(\mathcal{L}_{A2C} + \mathcal{L}_{C2A}). \tag{8}$$

The total loss $\mathcal{L}$ is the sum of $\mathcal{L}_V$ and $\mathcal{L}_A$:

$$\mathcal{L} = \mathcal{L}_V + \mathcal{L}_A. \tag{9}$$

For inference, we simply combine the similarity score of the two branches as Equation 4.

| Method | Venue | Input | Pre-training | Top-1(%) | Top-5(%) | Views | FLOPs | Param |
|---|---|---|---|---|---|---|---|---|
| NL I3D-101 [51] | CVPR'18 | $128\times224^2$ | ImageNet-1K | 77.7 | 93.3 | $10\times3$ | $359\times30$ | 61.8 |
| MVFNet$_{En}$ [56] | AAAI'21 | $24\times224^2$ | ImageNet-1K | 79.1 | 93.8 | $10\times3$ | $188\times30$ | - |
| TimeSformer-L [2] | ICML'21 | $96\times224^2$ | ImageNet-21K | 80.7 | 94.7 | $1\times3$ | $2380\times3$ | 121.4 |
| ViViT-L/16×2 [1] | ICCV'21 | $32\times320^2$ | ImageNet-21K | 81.3 | 94.7 | $4\times3$ | $3992\times12$ | 310.8 |
| VideoSwin-L [31] | CVPR'22 | $32\times384^2$ | ImageNet-21K | 84.9 | 96.7 | $10\times5$ | $2107\times50$ | 200.0 |
| *Methods with large-scale image pre-training* | | | | | | | | |
| ViViT-L/16×2 [1] | ICCV'21 | $32\times320^2$ | JFT-300M | 83.5 | 95.5 | $4\times3$ | $3992\times12$ | 310.8 |
| ViViT-H/16×2 [1] | ICCV'21 | $32\times224^2$ | JFT-300M | 84.8 | 95.8 | $4\times3$ | $8316\times12$ | 647.5 |
| TokenLearner-L/10 [41] | NeurIPS'21 | $32\times224^2$ | JFT-300M | 85.4 | 96.3 | $4\times3$ | $4076\times12$ | 450 |
| MTV-H [65] | CVPR'22 | $32\times224^2$ | JFT-300M | 85.8 | 96.6 | $4\times3$ | $3706\times12$ | - |
| CoVeR [70] | arXiv'21 | $16\times448^2$ | JFT-300M | 86.3 | - | $1\times3$ | - | - |
| CoVeR [70] | arXiv'21 | $16\times448^2$ | JFT-3B | 87.2 | - | $1\times3$ | - | - |
| *Methods with large-scale image-language pre-training* | | | | | | | | |
| CoCa ViT-giant [67] | arXiv'22 | $6\times288^2$ | JFT-3B+ALIGN-1.8B | 88.9 | - | - | - | 2100 |
| VideoPrompt ViT-B/16 [21] | ECCV'22 | $16\times224^2$ | WIT-400M | 76.9 | 93.5 | - | - | - |
| ActionCLIP ViT-B/16 [50] | arXiv'21 | $32\times224^2$ | WIT-400M | 83.8 | 96.2 | $10\times3$ | $563\times30$ | 141.7 |
| Florence [68] | arXiv'21 | $32\times384^2$ | FLD-900M | 86.5 | 97.3 | $4\times3$ | - | 647 |
| ST-Adapter ViT-L/14 [36] | NeurIPS'22 | $32\times224^2$ | WIT-400M | 87.2 | 97.6 | $3\times1$ | 8248 | - |
| AIM ViT-L/14 [66] | ICLR'23 | $32\times224^2$ | WIT-400M | 87.5 | 97.7 | $3\times1$ | 11208 | 341 |
| EVL ViT-L/14 [28] | ECCV'22 | $32\times224^2$ | WIT-400M | 87.3 | - | $3\times1$ | 8088 | - |
| EVL ViT-L/14 [28] | ECCV'22 | $32\times336^2$ | WIT-400M | 87.7 | - | $3\times1$ | 18196 | - |
| X-CLIP ViT-L/14 [35] | ECCV'22 | $16\times336^2$ | WIT-400M | 87.7 | 97.4 | $4\times3$ | $3086\times12$ | - |
| Text4Vis ViT-L/14 [60] | AAAI'23 | $32\times336^2$ | WIT-400M | 87.8 | 97.6 | $1\times3$ | $3829\times3$ | 230.7 |
| **BIKE ViT-L/14** | CVPR'23 | $16\times224^2$ | WIT-400M | 88.1 | 97.9 | $4\times3$ | $830\times12$ | 230 |
| | | $8\times336^2$ | | 88.3 | 98.1 | $4\times3$ | $932\times12$ | 230 |
| | | $32\times336^2$ | | **88.6** | **98.3** | $4\times3$ | $3728\times12$ | 230 |

Table 1. Comparisons with state-of-the-art methods on Kinetics-400. We report the FLOPs in inference phase."Views" indicates # temporal clip × # spatial crop. The magnitudes are Giga ($10^9$) and Mega ($10^6$) for FLOPs and Param.

# 3. Experiments

## 3.1. Setups

We conduct experiments on six widely used video benchmarks, *i.e.*, Kinetics-400 [22] & 600 [6], ActivityNet [5], Charades [42], UCF-101 [44] and HMDB-51 [24]. *See Supplementary for statistics of these datasets.*

**Training & Inference.** In our experiments, we adopt the visual encoder of CLIP [38] as the video encoder and use the textual encoder of CLIP for both the category and attributes encoders. To avoid conflict between the two branches, we first train the video encoder and then the attributes encoder. To prepare the video input, we sparsely sample $T$ (*e.g.*, 8, 16, 32) frames. We set the temperature hyperparameter $\tau$ to 0.01 for all training phases. *See Supplementary for detailed training hyperparameters.*

To trade off accuracy and speed, we consider two evaluation protocols. (1) *Single View*: We use only 1 clip per video and the center crop for efficient evaluation, as shown in Table 6. (2) *Multiple Views*: It is a common practice [7,14,56] to sample multiple clips per video with several spatial crops to get higher accuracy. For comparison with SOTAs, we use four clips with three crops ("4×3 Views") in Table 1.

## 3.2. Main Results

**Comparison with State-of-the-arts.** We present our results on **Kinetics-400** in Table 1 and compare our approach with SOTAs trained under various pre-training settings. Our approach outperforms regular video recognition methods while requiring significantly less computation, as shown in the upper table. We also demonstrate superiority over methods that use web-scale image pre-training, such as JFT-300M [45] and JFT-3B [69]. Our model performs better than all JFT-300M pre-trained methods, achieving a higher accuracy (+2.3%) than CoVeR [70]. Surprisingly, our method even outperforms the JFT-3B pre-trained model (**88.6%** *v.s.* 87.2%) despite the latter having almost 3 billion annotated images and a data scale 7.5×larger than ours. We further compare our method with others using web-scale image-language pre-training, such as CLIP [38] and Florence [68]. Despite Florence having a larger dataset (2×more data than the 400M image-text data used in CLIP), our approach still achieves a higher accuracy by 2.1%. Additionally, using only 8 frames and the same CLIP pre-

| Method | Top-1 | mAP |
|---|---|---|
| ListenToLook [17] | - | 89.9 |
| MARL [57] | 85.7 | 90.1 |
| DSANet [61] | - | 90.5 |
| TSQNet [62] | 88.7 | 93.7 |
| NSNet [63] | 90.2 | 94.3 |
| **BIKE ViT-L** | **94.7** | **96.1** |

Table 2. Comparisons with SOTAs on ActivityNet.

| Method | Frames | mAP |
|---|---|---|
| MultiScale TRN [73] | - | 25.2 |
| STM [20] | 16 | 35.3 |
| SlowFast R101 [14] | 16+64 | 42.5 |
| X3D-XL (312↑) [13] | 16 | 43.4 |
| ActionCLIP [50] | 32 | 44.3 |
| **BIKE ViT-L** | **16** | **50.4** |

Table 3. Comparisons on Multi-label video dataset Charades.

| Method | Shot | HMDB | UCF | ANet | K400 |
|---|---|---|---|---|---|
| VideoSwin [31] | 2 | 20.9 | 53.3 | - | - |
| VideoPrompt [21] | 5 | 56.6 | 79.5 | - | 58.5 |
| X-Florence [35] | 2 | 51.6 | 84.0 | - | - |
| **BIKE ViT-L** | 1 | 72.3 | 95.2 | 86.6 | 73.5 |
|  | 2 | **73.5** | **96.1** | **88.7** | **75.7** |
|  | 5 | **77.7** | **96.5** | **90.9** | **78.2** |

Table 4. Comparisons on few-shot action recognition across four video datasets.

| Method | UCF* / UCF | HMDB* / HMDB | ActivityNet*/ ActivityNet | Kinetics-600 |
|---|---|---|---|---|
| GA [34] | 17.3±1.1 / - | 19.3±2.1 / - | - | - |
| TS-GCN [16] | 34.2±3.1 / - | 23.2±3.0 / - | - | - |
| E2E [3] | 44.1 / 35.3 | 29.8 / 24.8 | 26.6 / 20.0 | - |
| DASZL [23] | 48.9±5.8 / - | - / - | - | - |
| ER [8] | 51.8±2.9 / - | 35.3±4.6 / - | - | 42.1±1.4 |
| ResT [26] | 58.7±3.3 / 46.7 | 41.1±3.7 / 34.4 | 32.5 / 26.3 | - |
| **BIKE ViT-L** | **86.6±3.4 / 80.8** | **61.4±3.6 / 52.8** | **86.2±1.0 / 80.0** | **68.5±1.2** |

Table 5. Comparisons on zero-shot video recognition. * denotes randomly selecting half of the test dataset's classes for evaluation, repeating the process ten times, and reporting the mean accuracy with standard deviation. For Kinetics-600, we adopt official code [8] to select the 220 new categories outside of Kinetics-400 for evaluation.

training, our model performs on par with the best results of other methods, such as EVL [28], X-CLIP [35], and Text4Vis [60]. When we use more frames as input, our method achieves a new state-of-the-art accuracy of 88.6% under the CLIP pre-training setting.

We also evaluate our method on the untrimmed video dataset, **ActivityNet-v1.3**, to verify its generalizability. We fine-tune the Kinetics-400 pre-trained model with 16 frames, and report the top-1 accuracy and mean average precision (mAP) using the official evaluation metrics. Our approach significantly outperforms recent SOTAs, as shown in Table 2. Furthermore, to demonstrate its effectiveness on smaller datasets, we also evaluate our method on **UCF-101** and **HMDB-51**, achieving top-1 accuracy of 98.8% and 83.1%, respectively. *We include the results in the Supplementary due to space limitations.*

**Multi-Label Video Recognition.** In addition to the single-label video recognition, we also evaluate our method on multi-label video recognition. We use the **Charades** dataset, which contains long-term activities with multiple actions, and utilize the Kinetics-400 pre-trained ViT-L backbone for training. The results are evaluated using the mAP metric. As shown in Table 3, our BIKE achieves the performance of 50.4% mAP, demonstrating its effectiveness in multi-label video classification.

**Few-Shot Video Recognition.** We demonstrate the few-shot recognition capability of our method, which refers to video recognition using only a few training samples. In this experiment, we scaled up the task to categorize all categories in the dataset with only a few samples per category for training. We used a CLIP pre-trained ViT-L/14 with 8 frames for few-shot video recognition, without further Kinetics-400 pre-training. The top-1 accuracy on four datasets is reported in Table 4. Our method shows remarkable transferability to diverse domain data in a data-poor situation. On UCF-101 and HMDB-51, our method outperforms VideoSwin [31] by 42.8% and 52.6%, respectively. In comparison with image-language pre-trained methods, our method outperforms VideoPrompt [21] and X-Florence [35] by 21.1% and 21.9% on HMDB-51, respectively. *See Supplementary for training details.*

**Zero-shot Video Recognition.** We further evaluate our method in an open-set setting. Table 5 presents the results of zero-shot evaluation on four video datasets using our Kinetics-400 pre-trained model (*i.e.*, ViT-L/14 with 8 frames). There are two major evaluation methods on UCF-101, HMDB-51, and ActivityNet: half-classes evaluation (marked as *) and full-classes evaluation. For fair comparison, we present the results under the half-classes evaluation protocol, which has been widely used in previous works [3, 8, 26, 34]. Additionally, we provide results on the entire dataset for more challenging and realistic accuracy evaluation. *See Supplementary for further details on evaluation protocols.* Our method exhibits strong cross-

| Video branch | $g(\cdot\|\phi_c)$ | Top-1(%) |
|---|---|---|
| Baseline: Mean Pool | 🔓 | 76.8 |
| + Video Concept Spotting | 🔓 | 78.5 (+1.7) |
| + (Technique) Transf | 🔓 | 78.7 (+1.9) |
| + Frozen label encoder | 🔒 | **78.9 (+2.1)** |

(a) The effectiveness of temporal saliency. 🔓 means finetuning category encoder $g(\cdot\|\phi_c)$. Transf is the temporal transformer.

| VCS Source | Recognition Source | Top-1 |
|---|---|---|
| Word Emb. | Word Emb. | 78.1 |
| [CLS] Emb. | [CLS] Emb. | 74.7 |
| Word Emb. | [CLS] Emb. | **78.5** |

(b) Different category embeddings are used for Video Concept Spotting (VCS) and recognition.

| Attributes | Category | Top-1 |
|---|---|---|
| ✗ | ✗ | 46.2 |
| ✓ | ✗ | 51.2 |
| ✓ | ✓ | **56.6** |

(c) The effects of the textual prompt in *Attributes recognition branch* (w/o training).

| #Attributes | A | V+A |
|---|---|---|
| 3 | 53.4 | 79.9 |
| 5 | 56.6 | **80.0** |
| 7 | 57.1 | 79.7 |

(d) Study on different number of attributes (w/o training).

| Training | A | V $\xrightarrow{+\Delta\%}$ V+A |
|---|---|---|
| ✗ | 56.6 | 78.9 $\xrightarrow{+1.1\%}$ 80.0 |
| ✓ | **69.6** | 78.9 $\xrightarrow{+2.5\%}$ **81.4** |

(e) The impact of *Attributes branch*. ✓ means fine-tuning the attributes encoder.

| | V $\xrightarrow{+\Delta\%}$ V+A |
|---|---|
| Baseline | 76.8 $\xrightarrow{+2.4\%}$ 79.2 |
| Ours | 78.9 $\xrightarrow{+2.5\%}$ **81.4** |

(f) The effects of *Attributes branch* to complement *Video branch*.

| Lexicon | V $\xrightarrow{+\Delta\%}$ V+A |
|---|---|
| IN-1K | 78.9 $\xrightarrow{+1.4\%}$ 80.3 |
| K400 | 78.9 $\xrightarrow{+2.5\%}$ **81.4** |

(g) Study on the impact of different lexicon.

| Method | $T$ | Backbone | Top-1(%) |
|---|---|---|---|
| VideoPrompt [21] | 16 | ViT-B/32 | 76.9 |
| ActionCLIP [50] | 8 | ViT-B/32 | 78.4 |
| BIKE (Ours) | 8 | ViT-B/32 | **81.4 (+3.0)** |

(h) Comparison with CLIP-based methods using single-view inference. $T$ is the number of frames.

| Backbone | Baseline → V → V+A | V* → V*+A |
|---|---|---|
| ViT-B/32 | 76.8 $\xrightarrow{+2.1\%}$ 78.9 $\xrightarrow{+2.5\%}$ 81.4 | 80.2 $\xrightarrow{+1.7\%}$ 81.9 |
| ViT-B/16 | 79.9 $\xrightarrow{+2.2\%}$ 82.1 $\xrightarrow{+1.1\%}$ 83.2 | 83.2 $\xrightarrow{+0.7\%}$ 83.9 |
| ViT-L/14 | 85.2 $\xrightarrow{+0.8\%}$ 86.4 $\xrightarrow{+0.1\%}$ 86.5 | 87.4 $\xrightarrow{+0\%}$ 87.4 |

(i) Component-by-component evaluation of our approach using various backbones. Models are fed 8 frames, where * stands for multiple view inference.

Table 6. Ablation studies on Kinetics-400. Models use ViT-B/32 as the backbone, and 8 frames as input, unless otherwise specified. We report top-1 accuracy (%) for a single clip input with 224×224 spatial size. The **V** and **A** abbreviations are used for the *Video recognition branch* and *Attributes recognition branch*, respectively. We refer to ImageNet-1K and Kinetics-400 as IN-1K and K400, respectively.

dataset generalization ability and outperforms classic zero-shot video recognition methods.

### 3.3. Ablation Studies

In this section, we provide extensive ablations to demonstrate our method with the instantiation in Table 6.

**The Effect of Temporal Saliency.** We investigate the impact of our proposed *Video Concept Spotting* (VCS) mechanism on the performance of the *Video branch*, as shown in Table 6a. We start with a baseline that uses mean pooling to aggregate the features of all frames, without considering temporal saliency. We observe that equipping the baseline with VCS can improve the accuracy by +1.7%. We then introduce a multi-layer (*e.g.*, 6-layers) Transformer encoder with position embedding for sequence features, commonly used in previous methods, and find that it provides an additional 0.2% performance boost. Moreover, freezing the category encoder not only reduces training parameters but also slightly improves performance (+0.2%).

**Exploration of Category Embedding for Temporal Saliency and Classification.** As mentioned in Section 2.3, CLIP's textual encoder can generate two types of embeddings: the [CLS] embedding for the entire sentence and the word embedding for each word. Therefore, we can encode the category into these two types of embeddings. The category embedding has two roles in our method: 1) it serves as

a query to determine the temporal saliency, and 2) it calculates similarity with video representation to produce recognition results. We demonstrate the results for three different combinations in Table 6b. We find that the global [CLS] embedding performs better than the word-level embedding for final recognition, but the word-level embedding is necessary for temporal saliency.

**Prompt Engineering and Number of Attributes.** For both attributes and categories in the *Attributes recognition* branch, we manually define a prompt, *i.e.*, "This is a video about {}". The results in Table 6c show that the prompt significantly improves accuracy, even without training the attributes encoder. Furthermore, in Table 6d, we observe that the number of attributes has little effect on the performance of the *Attributes recognition* and two-branch recognition.

**The Impact of *Attributes branch*.** Table 6e shows that without any training, the *Attributes branch* can be plug-and-played on the *Video branch* to improve the recognition performance. After training the attributes encoder, the *Attributes branch* further boosts performance by an impressive 2.5% on the fusion result. Additionally, we find that the *Attributes branch* can also improve the baseline when fused with it, as shown in Table 6f. By combining the *VCS* and the *Attributes branch*, we can achieve a remarkable improvement of 4.6% on the baseline.

**Attributes Generation with Different Lexicons.** In

Sec. 2.2, we use a pre-defined lexicon to obtain attributes. In Table 6g, we explore the impact of different lexicons. We used ImageNet-1K, an image dataset that covers 1000 object categories, as our lexicon to search for potential object attributes. According to the results, this can increase the performance of the *Attributes branch* by 1.4%. We found that using the 400 categories of Kinetics-400 as the lexicon can further improve the results.

**Comparison with CLIP-Based Methods.** Table 6h presents a comparison between our method and two CLIP-based approaches, VideoPrompt [21] and ActionCLIP [50], both trained with contrastive loss. Despite using fewer frames, our method achieves higher Top-1 accuracy than VideoPrompt. Moreover, using the same ViT-B/32 backbone, our approach outperforms ActionCLIP by **3.0%**.

**More Evaluation with Different Backbones.** Table 6i presents a comprehensive evaluation of the applicability of our method using larger backbones. Our observations are as follows: 1) Despite the greatly improved performance of the baseline with larger backbones, our *VCS* mechanism still provides consistent, additional gains. This demonstrates the continued necessity of `Text-to-Video` saliency knowledge for large models. 2) As the absolute accuracy of the *Video branch* increases, the complementing effect of the *Attributes branch* gradually weakens. We conjecture that with larger models, richer representations are learned, leading to reduced bias in learned representations and an increased correlation with the *Attributes branch*, resulting in a reduction in complementary information. 3) Multiple-view evaluation involving more video clips leads to increased performance, reducing the bias of the model itself. For models with a top-1 accuracy of 87.4%, the *Attributes branch* is unable to provide supplementary knowledge. Therefore, the *Attributes branch* is not utilized in our ViT-L/14 models presented in Sec. 3.2.

### 3.4. Visualization

Figure 3 illustrates the temporal saliency generated by Video Concept Spotting mechanism, highlighting the frame that is most relevant to the category. We also demonstrate the complementarity of the auxiliary attributes generated by our Video-Attribute Association mechanism with the video branch. *See more qualitative results in Supplementary.*

## 4. Related Works

**Video Recognition.** Convolutional networks have been the standard backbone architecture in video recognition for a long time. Early works focused on jointly learning spatial and temporal context through parallel branches [14, 15, 43, 49, 52, 54]. Later works developed plug-and-play temporal modules [25, 30, 32, 37, 46, 48, 56, 61, 64] for 2D CNN backbones to improve temporal modeling. Some works also designed dynamic inference mechanisms [53, 55, 57, 58, 62, 63]



Figure 3. Visualization of (**Top**) temporal saliency and (**Bottom**) attributes. Please zoom in for the best view.

for efficient video recognition. Recently, Vision Transformers [10, 18, 29] has emerged as a new trend in image recognition backbones. Transformers have also been adopted for video recognition, such as TimeSFormer [2], ViViT [1], VideoSwin [31], and MViT [11].

**Transferring CLIP Models for Video Recognition.** CLIP [38] provides good practice in learning the coordinated vision-language models using large-scale image and text pairs. The pre-trained model can learn powerful visual representations aligned with rich linguistic semantics. Initially, some works [12, 33, 59, 72] propose to directly use CLIP for video-text retrieval. Later, a few works also explore the use of CLIP models for video recognition [21, 28, 35, 36, 50, 60, 66], they can be broadly categorized into two lines. The first line [28, 36, 66] follows the unimodal transferring paradigm, where the image encoder of CLIP is used as a strong initialization for the video encoder. The second line [21, 35, 50, 60] provides cross-model learning baselines that directly extend CLIP to video-label matching for video recognition. However, these studies only briefly tap into the knowledge from CLIP. In contrast, our work aims to further explore the bidirectional cross-modal knowledge from CLIP to enhance the cross-model baseline. Our approach introduces auxiliary attributes in the `Video-to-Text` direction and category-dependent temporal saliency in the `Text-to-Video` direction, resulting in a more effective and interpretable video recognition.

## 5. Conclusion

In this work, we introduce a novel two-stream framework called **BIKE** that leverages bidirectional cross-modal knowledge from CLIP models to enhance video recognition. Our approach involves the *Attributes branch*, which utilizes the **Attributes-Category Association** mechanism to generate attributes for auxiliary recognition, and the *Video branch*, which uses the **Video Concept Spotting** mechanism to generate temporal saliency and produce a more compact video representation. Our approach is evaluated on six video datasets, and the experimental results demonstrate its effectiveness.

# References

[1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, pages 6836–6846, 2021. 5, 8

[2] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, pages 813–824. PMLR, 2021. 5, 8

[3] 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. 6, 12

[4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 33:1877–1901, 2020. 1

[5] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, 2015. 2, 5

[6] Joao Carreira, Eric Noland, Andras Banki-Horvath, Chloe Hillier, and Andrew Zisserman. A short note about kinetics-600. *arXiv preprint arXiv:1808.01340*, 2018. 2, 5

[7] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017. 5, 15

[8] Shizhe Chen and Dong Huang. Elaborative rehearsal for zero-shot action recognition. In *ICCV*, pages 13638–13647, 2021. 6, 13

[9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 1

[10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 8

[11] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, pages 6824–6835, 2021. 8

[12] Bo Fang, Chang Liu, Yu Zhou, Min Yang, Yuxin Song, Fu Li, Weiping Wang, Xiangyang Ji, Wanli Ouyang, et al. Uatvr: Uncertainty-adaptive text-video retrieval. *arXiv preprint arXiv:2301.06309*, 2023. 8

[13] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *CVPR*, pages 203–213, 2020. 6

[14] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2019. 5, 6, 8

[15] Christoph Feichtenhofer, Axel Pinz, and Richard Wildes. Spatiotemporal residual networks for video action recognition. In *NeurIPS*, 2016. 8

[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*, volume 33, pages 8303–8311, 2019. 6

[17] Ruohan Gao, Tae-Hyun Oh, Kristen Grauman, and Lorenzo Torresani. Listen to look: Action recognition by previewing audio. In *CVPR*, 2020. 6

[18] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. *NeurIPS*, 34, 2021. 8

[19] 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

[20] Boyuan Jiang, MengMeng Wang, Weihao Gan, Wei Wu, and Junjie Yan. Stm: Spatiotemporal and motion encoding for action recognition. In *ICCV*, pages 2000–2009, 2019. 6, 15

[21] 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. 1, 2, 4, 5, 6, 7, 8

[22] 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, 5

[23] Tae Soo Kim, Jonathan Jones, Michael Peven, Zihao Xiao, Jin Bai, Yi Zhang, Weichao Qiu, Alan Yuille, and Gregory D Hager. Daszl: Dynamic action signatures for zero-shot learning. In *AAAI*, volume 35, pages 1817–1826, 2021. 6

[24] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, 2011. 2, 5

[25] Yan Li, Bin Ji, Xintian Shi, Jianguo Zhang, Bin Kang, and Limin Wang. Tea: Temporal excitation and aggregation for action recognition. In *CVPR*, pages 909–918, 2020. 8

[26] Chung-Ching Lin, Kevin Lin, Lijuan Wang, Zicheng Liu, and Linjie Li. Cross-modal representation learning for zero-shot action recognition. In *CVPR*, pages 19978–19988, 2022. 6

[27] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, 2019. 15

[28] 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. In *ECCV*, pages 388–404. Springer, 2022. 1, 5, 6, 8

[29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 8

[30] Zhaoyang Liu, Donghao Luo, Yabiao Wang, Limin Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Tong Lu. Teinet: Towards an efficient architecture for video recognition. In *AAAI*, pages 11669–11676, 2020. 8

[31] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, pages 3202–3211, 2022. 5, 6, 8

[32] Zhaoyang Liu, Limin Wang, Wayne Wu, Chen Qian, and Tong Lu. Tam: Temporal adaptive module for video recognition. In *ICCV*, pages 13708–13718, 2021. 8

[33] 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 and captioning. *Neurocomputing*, 508:293–304, 2022. 8

[34] Ashish Mishra, Vinay Kumar Verma, M Shiva Krishna Reddy, S Arulkumar, Piyush Rai, and Anurag Mittal. A generative approach to zero-shot and few-shot action recognition. In *WACV*, pages 372–380. IEEE, 2018. 6

[35] 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. 1, 2, 5, 6, 8

[36] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning. In *NeurIPS*, 2022. 1, 5, 8

[37] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatiotemporal representation with pseudo-3d residual networks. In *ICCV*, 2017. 8

[38] 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, 2, 3, 4, 5, 8

[39] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 1

[40] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020. 1

[41] Michael Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. Tokenlearner: Adaptive space-time tokenization for videos. *NeurIPS*, 34:12786–12797, 2021. 5

[42] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *ECCV*, pages 510–526. Springer, 2016. 2, 5

[43] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014. 8

[44] 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. 2, 5

[45] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *ICCV*, pages 843–852, 2017. 5

[46] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *CVPR*, 2018. 8, 15

[47] Limin Wang, Wei Li, Wen Li, and Luc Van Gool. Appearance-and-relation networks for video classification. In *CVPR*, 2018. 15

[48] Limin Wang, Zhan Tong, Bin Ji, and Gangshan Wu. Tdn: Temporal difference networks for efficient action recognition. In *CVPR*, pages 1895–1904, 2021. 8, 15

[49] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016. 8

[50] Mengmeng Wang, Jiazheng Xing, and Yong Liu. Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472*, 2021. 1, 2, 4, 5, 6, 7, 8

[51] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018. 5

[52] Xiaohan Wang, Linchao Zhu, Heng Wang, and Yi Yang. Interactive prototype learning for egocentric action recognition. In *ICCV*, pages 8168–8177, 2021. 8

[53] Xiaohan Wang, Linchao Zhu, Fei Wu, and Yi Yang. A differentiable parallel sampler for efficient video classification. *ACM Transactions on Multimedia Computing, Communications and Applications*, 19(3):1–18, 2023. 8

[54] Xiaohan Wang, Linchao Zhu, Yu Wu, and Yi Yang. Symbiotic attention for egocentric action recognition with object-centric alignment. *IEEE TPAMI*, 2020. 8

[55] Yulin Wang, Zhaoxi Chen, Haojun Jiang, Shiji Song, Yizeng Han, and Gao Huang. Adaptive focus for efficient video recognition. In *ICCV*, pages 16249–16258, 2021. 8

[56] Wenhao Wu, Dongliang He, Tianwei Lin, Fu Li, Chuang Gan, and Errui Ding. Mvfnet: Multi-view fusion network for efficient video recognition. In *AAAI*, 2021. 5, 8, 15

[57] Wenhao Wu, Dongliang He, Xiao Tan, Shifeng Chen, and Shilei Wen. Multi-agent reinforcement learning based frame sampling for effective untrimmed video recognition. In *ICCV*, 2019. 6, 8

[58] Wenhao Wu, Dongliang He, Xiao Tan, Shifeng Chen, Yi Yang, and Shilei Wen. Dynamic inference: A new approach toward efficient video action recognition. In *Proceedings of CVPR Workshops*, pages 676–677, 2020. 8

[59] Wenhao Wu, Haipeng Luo, Bo Fang, Jingdong Wang, and Wanli Ouyang. Cap4video: What can auxiliary captions do for text-video retrieval? In *CVPR*, 2023. 8

[60] Wenhao Wu, Zhun Sun, and Wanli Ouyang. Revisiting classifier: Transferring vision-language models for video recognition. In *AAAI*, 2023. 1, 2, 5, 6, 8

[61] Wenhao Wu, Yuxiang Zhao, Yanwu Xu, Xiao Tan, Dongliang He, Zhikang Zou, Jin Ye, Yingying Li, Mingde Yao, Zichao Dong, et al. Dsanet: Dynamic segment aggregation network for video-level representation learning. In *ACM MM*, pages 1903–1911, 2021. 6, 8

[62] Boyang Xia, Zhihao Wang, Wenhao Wu, Haoran Wang, and Jungong Han. Temporal saliency query network for efficient video recognition. In *ECCV*, pages 741–759. Springer, 2022. 6, 8

[63] Boyang Xia, Wenhao Wu, Haoran Wang, Rui Su, Dongliang He, Haosen Yang, Xiaoran Fan, and Wanli Ouyang. Nsnet: Non-saliency suppression sampler for efficient video recognition. In *ECCV*, pages 705–723. Springer, 2022. 6, 8

[64] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *ECCV*, 2018. 8, 15

[65] 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. 5

[66] Taojiannan Yang, Yi Zhu, Yusheng Xie, Aston Zhang, Chen Chen, and Mu Li. Aim: Adapting image models for efficient video understanding. In *ICLR*, 2023. 1, 5, 8

[67] 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. 1, 5

[68] 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, 5

[69] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *CVPR*, pages 12104–12113, 2022. 5

[70] Bowen Zhang, Jiahui Yu, Christopher Fifty, Wei Han, Andrew M Dai, Ruoming Pang, and Fei Sha. Co-training transformer with videos and images improves action recognition. *arXiv preprint arXiv:2112.07175*, 2021. 5

[71] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. Ernie: Enhanced language representation with informative entities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451, 2019. 1

[72] Shuai Zhao, Linchao Zhu, Xiaohan Wang, and Yi Yang. Centerclip: Token clustering for efficient text-video retrieval. *The 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022. 8

[73] Bolei Zhou, Alex Andonian, Aude Oliva, and Antonio Torralba. Temporal relational reasoning in videos. In *ECCV*, 2018. 6

# Bidirectional Cross-Modal Knowledge Exploration for Video Recognition with Pre-trained Vision-Language Models

## *Supplementary Material*

Wenhao Wu[1,2]    Xiaohan Wang[3]    Haipeng Luo[4]    Jingdong Wang[2]    Yi Yang[3]    Wanli Ouyang[5,1]

[1]The University of Sydney        [2]Baidu Inc.        [3]Zhejiang University
[4]University of Chinese Academy of Sciences        [5]Shanghai AI Laboratory

whwu.ucas@gmail.com

In this appendix, we provide additional details and results for our approach. Specifically, §A contains further *details* on the training process (§A.1), attributes branch (§A.2), zero-shot evaluation (§A.3), statistics of video datasets (§A.4), visual encoder architectures (§A.5), and Distributed InfoNCE (§A.6). In §B, we present additional *results*, including comparisons on UCF-101 and HMDB-51 (§B.1) and more visualizations (§B.2).

## A. Implementation Details

### A.1. Training details

**Regular Video Recognition**. We present our approach for regular video recognition in Table A.1, sharing the same training recipe for all video datasets, including Kinetics-400, ActivityNet, Charades, HMDB-51, and UCF-101.

**Few-shot Video Recognition**. We repeat the samples to maintain the same number of iterations as the regular counterpart. For instance, if the model is trained on Kinetics-400 with around 900 iterations per epoch for the general setting, we repeat the sample to maintain the same number of iterations for few-shot settings. We train few-shot models for 2 epochs on Kinetics-400 and 10 epochs on other video datasets, *i.e.*, ActivityNet, HMDB-51, and UCF-101, while keeping other settings the same as in Table A.1.

**Zero-shot Video Recognition**. We use the Kinetics-400 pre-trained models to perform cross-dataset recognition without additional training on other datasets such as ActivityNet, HMDB-51, UCF-101, and Kinetics-600.

### A.2. Attributes Branch

To improve the quality of auxiliary attributes, we pre-generate them using CLIP ViT-L/14 with 8 frames. We employ the text encoder architecture of CLIP ViT-B/32 as our attribute encoder. To integrate the *Attributes branch* with the *Video branch*, we set $\lambda$ to 0.6 for the *Video branch* with ViT-B and $\lambda$ to 0.8 for the *Video branch* with ViT-L.

| Setting | Value |
|---|---|
| *Training Hyperparameter* | |
| Batch size | 256 |
| Vocabulary size | 49408 |
| Training epochs | 30 (ViT-B), 20 (ViT-L) |
| Optimizer | AdamW |
| Learning rate (Base) | 5e-5, cosine |
| Learning rate (CLIP layers) | 5e-6, cosine |
| Weight decay | 0.2 |
| Linear warm-up epochs | 5 |
| Adam $\beta_1,\beta_2$ | 0.9, 0.999 |
| *Augmentation* | |
| Resize | RandomSizedCrop |
| Crop size | 224 (Default) |
| Random Flip | 0.5 |
| Random Gray scale | 0.2 |

Table A.1. Default training recipe for video recognition.

### A.3. Evaluation Protocols for Zero-shot Recognition

We employ our Kinetics-400 pre-trained models to evaluate on other datasets. For UCF-101, HMDB-51, and ActivityNet, we adopt two major evaluation protocols as described in [3]:

1. **Half-Classes Evaluation:** To ensure comparability with previous works, we randomly select half of the test dataset's classes - 50 for UCF, 25 for HMDB, and 100 for ActivityNet - and evaluate on the selected subset. We repeat this process ten times and average the results for each test dataset. We refer to this setting as UCF[*], HMDB[*] and ActivityNet[*].

2. **Full-Classes Evaluation:** This evaluation setting involves directly evaluating on the full dataset to return more realistic accuracy scores.

| Model | Embedding dimension | Input resolution | Vision Transformer | | | Text Transformer | | |
|---|---|---|---|---|---|---|---|---|
| | | | layers | width | heads | layers | width | heads |
| ViT-B/32 | 512 | 224 | 12 | 768 | 12 | 12 | 512 | 8 |
| ViT-B/16 | 512 | 224 | 12 | 768 | 12 | 12 | 512 | 8 |
| ViT-L/14 | 768 | 224 | 24 | 1024 | 16 | 12 | 768 | 12 |
| ViT-L/14-336px | 768 | 336 | 24 | 1024 | 16 | 12 | 768 | 12 |

Table A.2. CLIP-ViT hyperparameters

For Kinetics-600, we follow [8] to choose the 220 new categories outside of Kinetics-400 in Kinetics-600 for evaluation. We use the three splits provided by [8] and sample 160 categories for evaluation from the 220 categories in Kinetics-600 for each split. We report the mean accuracy of the three splits as the final accuracy.

### A.4. Statistics of Video Datasets

We describe the video datasets used in our experiments:
**Kinetics-400** is a large-scale video dataset that includes 240,000 training videos and 20,000 validation videos across 400 different human action categories. Each video in the dataset is a 10-second clip of an action moment, annotated from raw YouTube videos.
**Kinetics-600** is an extension of Kinetics-400, consisting of approximately 480,000 videos from 600 action categories. The videos are divided into 390,000 for training, 30,000 for validation, and 60,000 for testing. In this paper, we use its test set for zero-shot evaluation.
**UCF-101** is an action recognition dataset that contains 13,320 videos from 101 realistic action categories, collected from YouTube.
**HMDB-51** is a collection of realistic videos from various sources, including movies and web videos. The dataset comprises 7,000 video clips from 51 action categories.
**ActivityNet-v1.3** is a large-scale untrimmed video benchmark that contains 19,994 untrimmed videos of 5 to 10 minutes from 200 activity categories.
**Charades** is a video dataset designed for action recognition and localization tasks. It contains over 10,000 short video clips of people performing daily activities, and consists of 157 action categories.

### A.5. Encoder Architectures

In this paper, we provide the complete architecture details of the visual encoder and textual encoders. The CLIP-ViT architectures are shown in Table A.2.

### A.6. Distributed InfoNCE

Instead of Data-Parallel Training (DP), which is single-process, multi-thread, and only works on a single machine, Distributed Data-Parallel Training (DDP) is a widely adopted single-program multiple-data training paradigm for single- and multi-machine training. Due to GIL contention across threads, per-iteration replicated model, and additional overhead introduced by scattering inputs and gathering outputs, DP is usually slower than DDP even on a single machine. Hence, we develop the Distributed InfoNCE based on DDP for large batch size and fast training.

The core of the Distributed InfoNCE implementation is batch gathering, which enables us to calculate the NM×NM similarity matrix across M GPUs for InfoNCE loss. Without batch gathering, each GPU only computes a local N×N matrix where N≪NM. This means that the cosine similarity and the InfoNCE loss would only be calculated for the pairs within a single GPU, and their gradients would be later averaged and synced. That's obviously not what we want.

The batch gathering technique allows each GPU to hold N vision features and perform a matrix product with NM text features, resulting in an N×NM matrix. This computation is distributed (*i.e.*, sharded) across M GPUs, and we have calculated NM×NM similarities across the GPUs in total. The loss we employ is symmetric, and the same process is applied *w.r.t.* text inputs. Algorithm 1 provides an example pseudocode to help understand the process.

## B. More Results

### B.1. Comparisons on UCF-101 and HMDB-51

In this section, we evaluate the performance of our method on the UCF-101 and HMDB-51 datasets to demonstrate its capacity for generalization to smaller datasets. We finetune our models on these two datasets using the pre-trained ViT-L model on Kinetics-400 and report the accuracy on split one. We use 16 frames as inputs and train for 30 epochs. Table A.3 shows that our model has strong transferability, achieving a mean class accuracy of 98.8% on UCF-101 and 83.1% on HMDB-51.

### B.2. More Qualitative Results

We present additional visualizations of the *Temporal Saliency* generated by our Video Concept Spotting mechanism in Figure A.1. In Figure A.2, we also showcase more visualizations of the *Generated Attributes* produced by our Video-Attribute Association mechanism using two different lexicons.
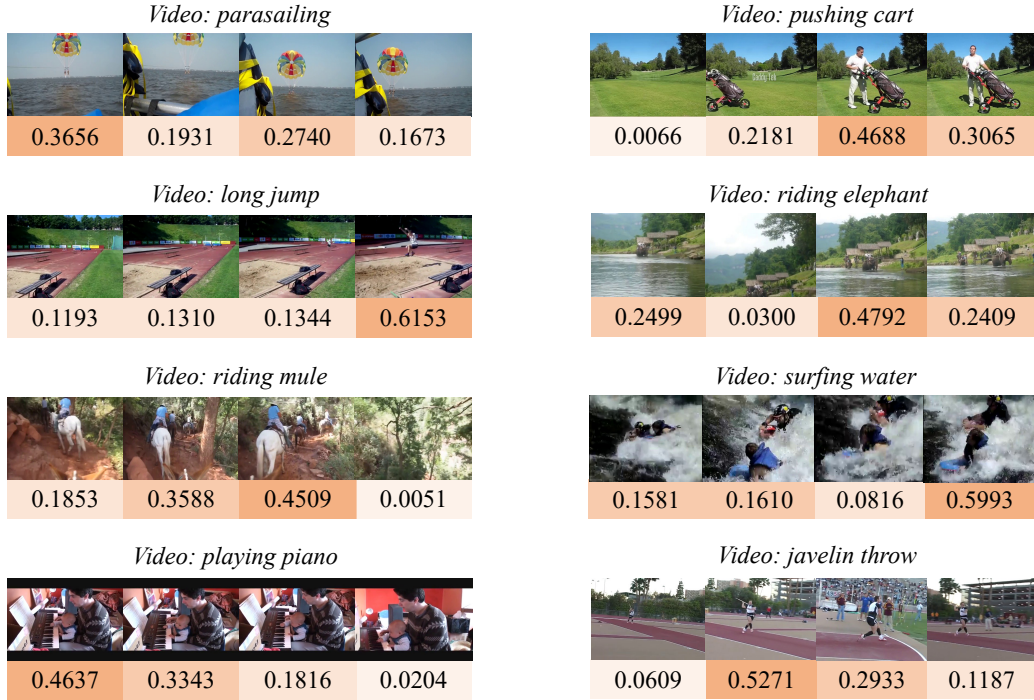
Figure A.1. Visualization of temporal saliency from our **Video Concept Spotting** mechanism. Please zoom in for best view.
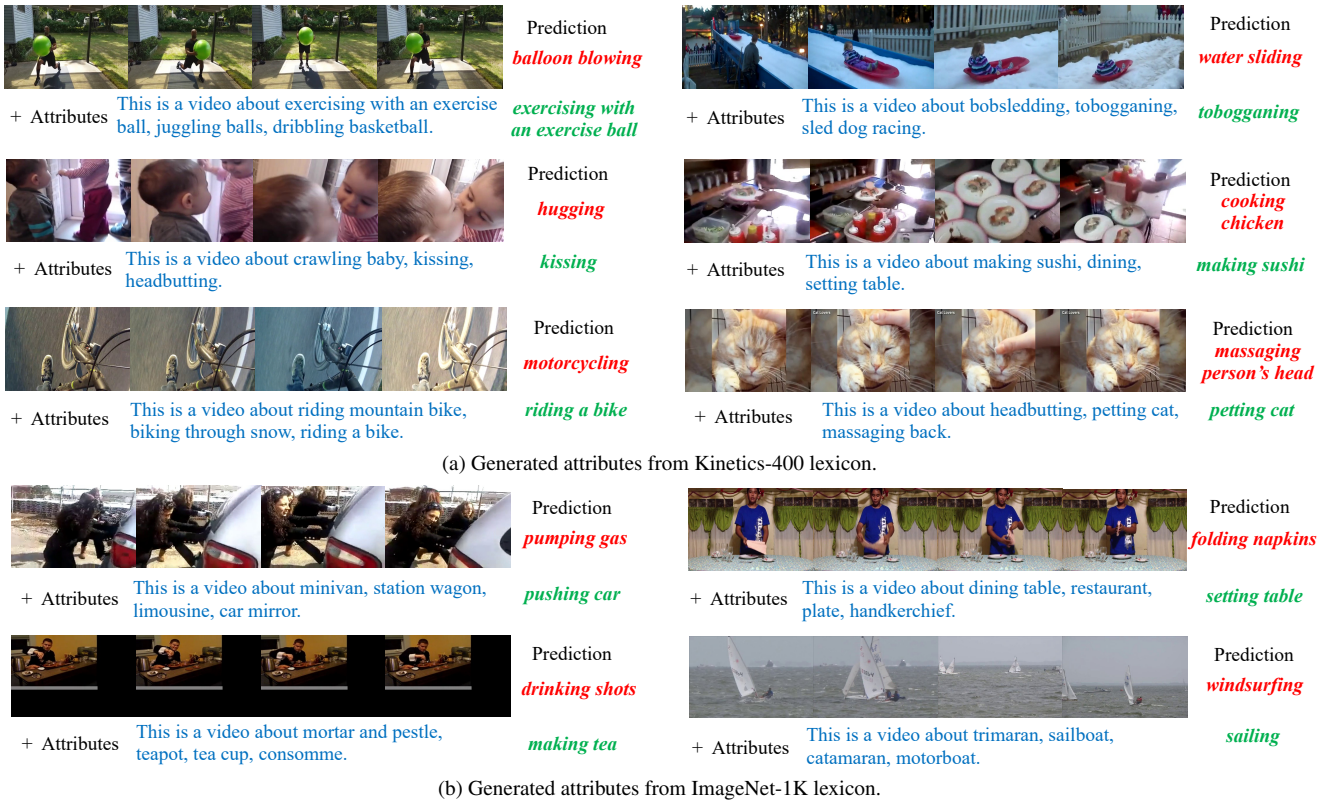


(a) Generated attributes from Kinetics-400 lexicon.



(b) Generated attributes from ImageNet-1K lexicon.

Figure A.2. Visualization of the attribute sentence generated by the **Video-Attribute Association** mechanism that corrected the original incorrect prediction to the correct one.

**Algorithm 1** Numpy-like Pseudocode of Distributed InfoNCE for our *Video branch*

```
1    # category_encoder: encoder network for category input
2    # video_encoder: encoder network for video input
3    # V: minibatch of video inputs
4    # T: minibatch of category inputs
5    # N: the local batch size of each GPU, e.g.,16
6    # M: the number of GPUs, e.g.,8
7    # N * M: the global batch size for multi-gpu training, e.g.,128
8
9    # extract feature representations of each modality
10   local_vision_features = video_encoder(V) # shape: [N, embed_dim]
11   local_text_features = category_encoder(T) # shape: [N, embed_dim]
12
13   # normalization
14   local_vision_features = l2_normalize(local_vision_features, axis=1)
15   local_text_features = l2_normalize(local_text_features, axis=1)
16
17   # batch_gather is a function gathering and concatenating the tensors across GPUs.
18   all_vision_features = batch_gather(local_vision_features) # shape: [N * M, embed_dim]
19   all_text_features = batch_gather(local_text_features) # shape: [N * M, embed_dim]
20
21   # scaled pairwise cosine similarities
22   # shape = [N, N * M]
23   logits_per_vision = logit_scale * local_vision_features @ all_text_features.t()
24   # shape = [N, N * M]
25   logits_per_text = logit_scale * local_text_features @ all_vision_features.t()
26
27   # The logits are then used as inputs for N*M-way (e.g., 128-way) classification,
28   # resulting in a loss value corresponding to N inputs in each GPU.
29   # Then Distributed Data Parallel mechanism takes care of averaging these across GPUs,
30   # which becomes equivalent to calculating the loss over NMxNM (e.g.,128x128) similarities.
31
```

| Method | UCF-101 | HMDB-51 |
|---|---|---|
| ARTNet [47] | 94.3% | 70.9% |
| I3D [7] | 95.6% | 74.8% |
| R(2+1)D [46] | 96.8% | 74.5% |
| S3D-G [64] | 96.8% | 75.9% |
| TSM [27] | 95.9% | 73.5% |
| STM [20] | 96.2% | 72.2% |
| MVFNet [56] | 96.6% | 75.7% |
| TDN [48] | 97.4% | 76.4% |
| Ours ViT-L | **98.8%** | 82.2% |
| Ours ViT-L (336↑) | 98.6% | **83.1%** |

Table A.3. Top-1 accuracy on UCF-101 and HMDB-51 achieved by different methods which are transferred from their **Kinetics Pre-trained** models with RGB modality.