

# Zero-Shot Temporal Action Detection via Vision-Language Prompting

Sauradip Nag<sup>1,2</sup>, Xiatian Zhu<sup>1,3</sup>, Yi-Zhe Song<sup>1,2</sup>, and Tao Xiang<sup>1,2</sup>

<sup>1</sup> CVSSP, University of Surrey, UK

<sup>2</sup> iFlyTek-Surrey Joint Research Centre on Artificial Intelligence, UK

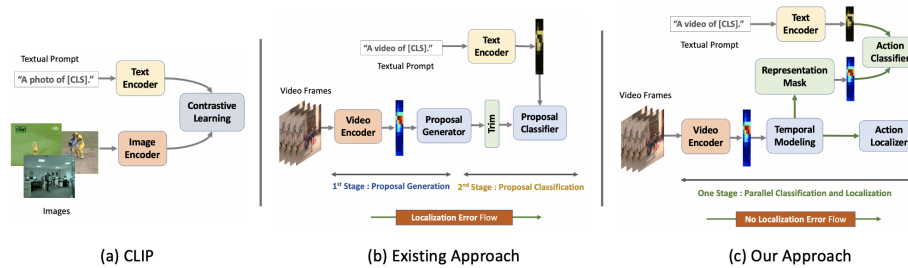
<sup>3</sup> Surrey Institute for People-Centred Artificial Intelligence, University of Surrey, UK  
{s.nag,xiatian.zhu,y.song,t.xiang}@surrey.ac.uk

**Abstract.** Existing temporal action detection (TAD) methods rely on large training data including segment-level annotations, limited to recognizing previously seen classes alone during inference. Collecting and annotating a large training set for each class of interest is costly and hence unscalable. Zero-shot TAD (ZS-TAD) resolves this obstacle by enabling a pre-trained model to recognize any unseen action classes. Meanwhile, ZS-TAD is also much more challenging with significantly less investigation. Inspired by the success of zero-shot image classification aided by vision-language (ViL) models such as CLIP, we aim to tackle the more complex TAD task. An intuitive method is to integrate an off-the-shelf proposal detector with CLIP style classification. However, due to the *sequential* localization (*e.g.*, proposal generation) and classification design, it is prone to localization error propagation. To overcome this problem, in this paper we propose a novel *zero-Shot Temporal Action detection model via Vision-Language prompting* (STALE). Such a novel design effectively eliminates the dependence between localization and classification by breaking the route for error propagation in-between. We further introduce an interaction mechanism between classification and localization for improved optimization. Extensive experiments on standard ZS-TAD video benchmarks show that our STALE significantly outperforms state-of-the-art alternatives. Besides, our model also yields superior results on supervised TAD over recent strong competitors. The PyTorch implementation of STALE is available on <https://github.com/sauradip/STALE>.

**Keywords:** Zero-shot transfer, Temporal action localization, Language supervision, Task adaptation, Detection, Dense prediction.

## 1 Introduction

The introduction of large pretrained Visual-Language (ViL) models (*e.g.*, CLIP [36] and ALIGN [16]) has surged recently the research attempts on *zero-shot transfer* to diverse downstream tasks via *prompting*. Central in this research line is a favourable ability of *synthesizing the classification weights* from natural language which can describe almost all the classes of interest, along with a



**Fig. 1.** Illustration of the (a) standard vision-language model CLIP [36]. (b) Existing ZS-TAD method [17] suffers from the intrinsic localization error propagation problem, due to the *sequential* localization (e.g., proposal generation) and classification design. (c) We solve this problem by designing a STALE model with a *parallel* localization and classification architecture.

strong image encoder trained with an objective of aligning a massive number of potentially noisy image-text pairs in a common feature space.

Whilst existing works mostly focus on image recognition tasks [54,12], how to capitalize the knowledge of such ViL models for natural video understanding tasks (e.g., temporal action detection [20,47,4,27,26,28,29]) remains unsolved largely. There are several challenges to be solved. **First**, as those *public accessible* large ViL models are often pretrained with big image data, temporal structured information is absent which is limited for video representation. **Second**, ViL models favor classification type of tasks *without built-in detection capability* in formulation. That is, their optimization assumes implicitly limited background *w.r.t* the foreground content with the training data. In contrast, natural videos often come with a large, variable proportion of background, and detecting the content of interest (e.g., human actions) is critical [1]. It is non-trivial to overcome these obstacles. On using ViL models for temporal action detection (TAD), indeed a recent attempt [17] is made. Specifically, [17] presents a two-stage video model: generating many action proposals with an off-the-shelf pre-trained proposal detector (e.g., BMN [20]), followed by proposal classification. This method has obvious limitations: (1) It is unable to adapt the localization module (i.e., proposal detection) which is pretrained and frozen throughout. (2) As a result, the compatibility between the localization and classification is limited, further leading to a localization error propagation problem due to the sequential localization and classification pipeline.

In this paper, to overcome the aforementioned limitations, we formulate a novel one-stage zero-shot temporal action localization architecture characterized with *parallel* classification and localization. This parallel head design naturally eliminates error propagation. Importantly, we introduce a learnable localization module based on a notion of *representation masking* – which is class-agnostic and hence generalizable to unseen classes. It can be optimized end-to-end, together with the classification component powered by a ViL model for zero-shot transfer, consequently solving the component compatibility problem. For im-

proving cross-modal (from image to video) task adaptation, we further design an inter-stream alignment regularization along with the self-attention mechanism. We term our method ***Zero-Shot Temporal Action Detection Model via Vision-Language Prompting*** (STALE).

**Contributions.** (1) We investigate the under-studied yet critical problem of how to capitalize large pretrained ViL models for zero-shot temporal action localization (ZS-TAD) in untrimmed videos. (2) We present a novel one-stage model, STALE, featured with a parallel classification and localization design interleaved by a learnable class-agnostic representation masking component for enabling zero-shot transfer to unseen classes. For enhancing cross-modal task adaptation, we introduce an inter-stream alignment regularization in the Transformer framework. (3) Extensive experiments on standard ZS-TAD video benchmarks show that our STALE outperforms state-of-the-art alternative methods, often by a large margin. Besides, our model can be also applied to the fully supervised TAD setting and achieve superior performance in comparison to recent supervised competitors.

## 2 Related Works

**Vision-language models** There have been a series of works on the interaction of computer vision and natural language processing fields, *e.g.*, text-to-image retrieval [43], image caption [46], visual question answering [2], and so on. Among these works, vision language (ViL) pre-training has attracted growing attention during the past few years [19,23]. As a milestone, Radford et al. [36] devise a large-scale pretraining ViL model, named CLIP, trained with a contrastive learning strategy on 400 million image-text pairs. It shows impressive zero-shot transferable ability over 30 classification datasets. Since then, many follow-ups have been proposed, including improved training strategy (*e.g.*, CoOp [54], CLIP-Adapter [12], Tip-adapter [50]). In video domains, similar idea has also been explored for transferable representation learning [24], text based action localization [33]. CLIP has also been used very recently in action recognition (*e.g.*, ActionCLIP [42]) and TAD [17]. Conceptually, all these methods use CLIP in a classification perspective. Commonly a two-stage pipeline is applied, where the first stage needs to crop the foreground followed by aligning the foreground using CLIP in the second stage. However, this strategy often suffers from the error propagation problem, that is, the errors from first stage would flow into the second stage and potentially amplify. In this work, we tackle this challenge by designing a one-stage ZS-TAD architecture featured with parallel classification and localization.

**Temporal action detection** Substantial progress has been made in TAD. Inspired by object detection in static images [38], R-C3D [45] uses anchor boxes by following the design of proposal generation and classification. With a similar model design, TURN [11] aggregates local features to represent snippet-level features for temporal boundary regression and classification. SSN [52] decomposes an action instance into three stages (starting, course, and ending) and employs

structured temporal pyramid pooling to generate proposals. BSN [21] predicts the start, end and actionness at each temporal location and generates proposals with high start and end probabilities. The actionness was further improved in BMN [20] via additionally generating a boundary-matching confidence map for improved proposal generation. GTAN [22] improves the proposal feature pooling procedure with a learnable Gaussian kernel for weighted averaging. G-TAD [47] learns semantic and temporal context via graph convolutional networks for more accurate proposal generation. BSN++ [40] further extends BMN with a complementary boundary generator to capture rich context. CSA [39] enriches the proposal temporal context via attention transfer. Recently, VSGN [51] improves short-action localization using a cross-scale multi-level pyramidal architecture. Commonly, the existing TAD models mostly adopt a 2-stage *sequential* localization and classification architecture. This would cause the localization error propagation problem particularly in low-data setting such as ZS-TAD. Our STALE is designed to address this limitation by designing a single stage model, thereby removing the dependence between localization and classification and cutting off the error propagation path.

**Zero-shot temporal action detection** Zero-shot learning (ZSL) is designed to recognize new classes that are not seen during training [44]. The idea is to learn shared knowledge from prior information and then transfer that knowledge from seen classes to unseen classes [30,35]. Visual attributes (*e.g.*, color, shape, and any properties) are the typical forms of prior information. For example, Lampert et al. [18] prelearned the attribute classifiers independently to accomplish ZSL on unseen classes, while Parikh et al. [32] learned relative attributes. Despite promising results on ZSL, attribute-based methods have poor scalability because the attributes need to be manually defined. Semantic embeddings of seen and unseen concepts, another type of prior information, can solve this scalability problem [48]. They are generally learned in an unsupervised manner such as Word2Vec [13] or GloVe [34]. Zhang et al. [49] firstly applied zero-shot learning on TAD using Word2Vec. Very recently, EffPrompt [17] used image-text pretraining from CLIP [36] for ZS-TAD. However, due to a two-stage design, this method also has the error propagation problem, in addition to its inability of learning the action localization module. We address all these limitations by introducing a new one-stage ZS-TAD architecture.

### 3 Methodology

We aim to efficiently steer an image-based ViL model (CLIP [36]) to tackle dense video downstream tasks such as Zero-Shot Temporal Action Detection (ZS-TAD) in untrimmed videos. This is essentially a model adaptation process with the aim to leverage the rich semantic knowledge from large language corpus.

#### 3.1 Preliminaries: Visual-Language Pre-training

The key capability of CLIP is to align the embedding spaces of visual and language data [36]. It consists of two encoders, namely an image encoder (*e.g.*,

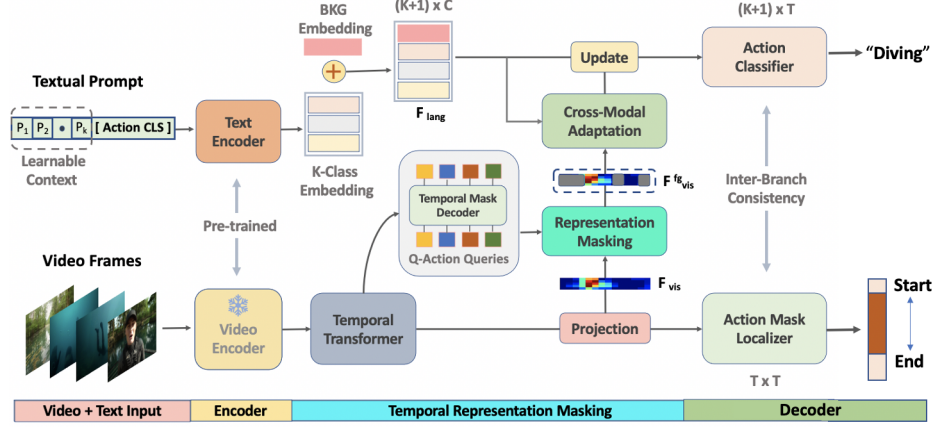
ResNet [14] or ViT [9]) and a text encoder (*e.g.*, Transformer [41]). To learn rich transferable semantic knowledge, CLIP leverages 400 million image-text pairs during training. To exploit CLIP’s knowledge for downstream classification tasks, an effective approach is to construct a set of text prompts with a template such as “a photo of [CLS].”, where [CLS] can be replaced by any class name of interest. Given an image, one can then use CLIP to compute the similarity scores between this image and the text prompts in the embedding space and take the class with the highest score as the prediction. Recently, a couple of works [54, 12] have demonstrated that CLIP can obtain strong classification performance with few or even zero training examples per class. We raise an interesting question: ***Whether the impressive ability of CLIP can be transferred to more complex vision tasks like dense prediction?***

This extended transfer could be intrinsically nontrivial. *Firstly*, how to leverage the visual-language pre-trained model in dense prediction tasks is a barely studied problem especially in zero-shot setup [37]. A simple method is to only use the image encoder of CLIP. However, we argue that the language priors with the pretrained text encoder are also of great importance and should be leveraged together. *Secondly*, transferring the knowledge from CLIP to dense prediction is more difficult than classification tasks, due to the substantial task discrepancy involved. The pretraining focuses on global representation learning of both images and texts, which is incompatible for local pixel-level outputs as required in downstream tasks. RegionCLIP [53] recently solves this problem in a 2-stage design including class-agnostic masking generalizable to unseen classes, followed by CLIP style classification. EffPrompt [17] similarly tackles a dense video understanding task TAD. Nonetheless, this approach suffers from the notorious localization error propagation challenge.

### 3.2 Language Guided Temporal Action Detection

To solve the aforementioned issues, we propose a language guided proposal-free framework. It can better leverage the language-priors of the pre-trained CLIP model. In the following sections, we start by describing the problem scenario and notations followed by introducing the idea for model adaptation through prompt learning. We then discuss how masking can help maintain the zero-shot transfer property in a one-stage design. Lastly, we discuss about the refinement of boundaries.

**Problem definition** We assume a dataset  $D$  with the training  $D_{train} = \{V_i, \psi_i\}_{i=1}^N$  and validation  $D_{val}$  sets. Each untrimmed training video  $V_i$  is labeled with temporal segmentation  $\Psi_i = \{(\psi_j, \xi_j, y_j)\}_{j=1}^{M_i}$  where  $\psi_j/\xi_j$  denote the start/end time,  $y_j$  is the action category, and  $M_i$  is the action instance number. Respectively,  $y_j$  refers to one of the training ( $D_{train}$ ) action labels in the text format for recognition, *e.g.*,  $y_j$  replaces the [CLS] token in a sentence of “a photo of [CLS]”. We consider both closed-set and open-set scenarios. In the closed-set scenario, the action categories for training and evaluation are identical, *i.e.*,



**Fig. 2. Overview of the proposed *Zero-Shot Temporal Action Detection via Vision-Language Prompting* (STALE) method.** Given an untrimmed video  $V$ , (a) we first extract a sequence of  $T$  snippet features with a pre-trained frozen video encoder and conduct self-attention learning using temporal embedding to obtain the snippet embedding  $E$  with a global context. (b) For each snippet embedding, we then predict a classification score  $P$  with the classification stream by masking the foreground feature and aligning with the text encoder embedding to obtain a classifier output  $P$ . The other branch of snippet embedding is used by action mask classifier to obtain a foreground mask  $M$  in parallel, (c) both of which are further used for consistency refinement at the feature level.

$D_{train} = D_{val}$ . While in the open-set case, the action categories for training and evaluation are disjoint, *i.e.*,  $D_{train} \cap D_{val} = \phi$ .

**Visual language embedding** Given a varying length untrimmed video  $V$ , following the standard practice [47, 20] we first sample a very large  $T$  equidistantly distributed temporal snippets (points) over the entire length.

*Visual embedding:* To extract features from the video snippets we use a frozen pre-trained video encoder (*e.g.*, a I3D [6], CLIP [36]) to extract RGB  $X_r \in \mathbb{R}^{d \times T}$  and optical flow features  $X_o \in \mathbb{R}^{d \times T}$  at the snippet level, where  $d$  denotes the feature dimension. We then concatenate them as  $E = [X_r; X_o] \in \mathbb{R}^{2d \times T}$ . Each snippet is a short sequence of (*e.g.*, 16 in our case) consecutive frames. While  $F$  contains local spatio-temporal information, it lacks a global context critical for TAD. We hence leverage the self-attention mechanism [41] to learn the global context. Formally, we set the input (**query**, **key**, **value**) of a multi-head Transformer encoder  $\mathcal{T}()$  as the features  $(F, F, F)$  (Fig. 2). Positional encoding is not applied as it is found to be detrimental (see the **supplementary material**). The final video snippet embedding is then obtained as

$$F_{vis} = \mathcal{T}(E) \in \mathbb{R}^{C \times T}, \quad (1)$$

with  $C$  being the embedding dimension.

*Textual encoding:* For textual embedding, we use a standard CLIP pre-trained Transformer [36] with learnable prompt, similar to [54], and as opposed to hand-crafted prompt used in CLIP [36]. By making it learnable, textual contexts can now achieve better transferability in downstream classification tasks by directly optimizing the contexts using back-propagation. The input to the text encoder, formed as:

$$F_{lan} = [G_p; G_e^k] \text{ where } G_p \in \mathbb{R}^{\hat{N} \times C'}, \quad (2)$$

are the learnable textual contexts, with  $\hat{N}$  as the length of the contexts. The embedding  $G_e^k \in \mathbb{R}^{C'}$  represents the textual embedding from CLIP vocabulary for each of the name of  $k$ -th class. The textual embedding for background class is not directly available from CLIP vocabulary, but necessary in TAD. To solve this issue, we learn a specific background embedding, denoted as  $G_e^{bg} \in \mathbb{R}^{C'}$ . We append this to the action class embedding  $G_e^k$ , making  $F_{lan}$  an embedding with  $K + 1$  classes. The background embedding  $G_e^{bg}$  is initialized randomly.

**Class agnostic representation masking** We introduce a novel *class agnostic representation masking* concept for enabling the usage of a ViL model for ZS-TAD. This is conceptually inspired by mask-transformer [8] with focus on a totally different problem (image segmentation *without* the aid of ViL model). Specifically, given the snippet embedding  $F_{vis}$  per video and  $N_z$  mask queries particularly introduced in this work, we leverage a transformer decoder [5] to generate  $N_z$  latent embeddings. Then we pass each latent embedding through a masking projection layer to obtain a mask embedding for each segment as  $B_q \in \mathbb{R}^{q \times C}$  where  $q$  indexes a query. A binary mask prediction *w.r.t* each query can be then calculated as:

$$L_q = \sigma(B_q * F_{vis}) \in \mathbb{R}^{q \times T}, \quad (3)$$

where  $\sigma$  is sigmoid activation. As such, each snippet location is associated with  $q$  queries. To choose the optimal query per location, we deploy a tiny MLP to weigh these queries in a location specific manner. This is realized by learning a weight vector  $W_q \in \mathbb{R}^{1 \times q}$  as:

$$\hat{L} = \sigma(W_q * L_q + b_q) \in \mathbb{R}^T. \quad (4)$$

We then binarize this mask at a threshold  $\theta_{bin}$  and select the foreground mask, denoted by  $\hat{L}_{bin}$ . To obtain the foreground features  $F_{vis}^{fg}$ , we use  $\hat{L}_{bin}$  to retrieve the snippet embedding  $F_{vis}$ . Given the binary nature of this foreground feature mask  $\hat{L}_{bin}$ , our representation masking can be first optimized on seen action classes and further generalize to unseen classes.

**Vision-language cross-modal adaption** Intuitively, integrating the descriptions of visual contexts is likely to enrich the text representation. For example, “a video of a man *playing kickball* in a big park” is richer than “a video of a man



playing kickball”. This motivates us to investigate how to use visual contexts to refine the text features. Specifically, we leverage the contextual-level visual feature to guide the text feature to adaptively explore informative regions of a video. For cross-attention, we adopt the standard architecture of the transformer [41]. Concretely, the cross attention module consists of a self-attention layer, a co-attention layer and a feed forward network formed as:

$$\hat{E}_c = \mathcal{T}_c(F_{lan}, F_{vis}^{fg}, F_{vis}^{fg}), \quad (5)$$

where  $\mathcal{T}_c$  is a transformer layer, taking  $F_{lan}$  as the **query**, and  $F_{vis}^{fg}$  as the **key** and **value**. This module encourages the text features to find most related visual clues across the foreground snippets. We then update the text features through a residual connection:

$$\hat{F}_{lan} = F_{lan} + \alpha \hat{E}_c, \quad (6)$$

where  $\alpha \in \mathbb{R}^C$  is a learnable parameter to control the scaling of the residual.  $\alpha$  is initialized with very small values (e.g.,  $10^{-3}$ ) to maximally preserve the language priors of text features.

**Parallel classification and mask prediction** Our TAD head is featured with parallel classification and mask prediction as detailed below.

**(I) Contextualized vision-language classifier:** In the standard training process of CLIP [36], the global feature is normally used during contrastive alignment. In general, it estimates the snippet-text score pairs by taking the average pooling of snippet features and then uses it with the language features. However, this formulation is unsuitable for dense classification tasks like TAD where each temporal snippet needs to be assigned with a class label. Under this consideration, we instead use the updated textual features  $\hat{F}_{lan} \in \mathbb{R}^{(K+1) \times C}$  and the masked foreground feature  $F_{vis}^{fg} \in \mathbb{R}^{T \times C}$  as follows:

$$\mathcal{P} = \hat{F}_{lan} * (F_{vis}^{fg})^T, \quad (7)$$

where  $\mathcal{P} \in \mathbb{R}^{(K+1) \times T}$  represents the classification output, where each snippet location  $t \in T$  is assigned with a probability distribution  $p_t \in \mathbb{R}^{(K+1) \times 1}$ . Note,  $l_2$  normalized along the channel dimension is applied prior to Eq. (7).

**(II) Action mask localizer:** In parallel to the classification stream, this stream predicts 1-D masks of action instances across the whole temporal span of the video. Since the 1-D masks are conditioned on the temporal location  $t$ , we exploit dynamic convolution [7]. This is because, in contrast to standard convolution, dynamic filters allow to leverage separate network branches to generate a filter at each snippet location. As a consequence, the dynamic filters can learn the context of the action (background) instances at each snippet location individually. More specifically, given the  $t$ -th snippet  $F_{vis}(t)$ , it outputs a 1-D mask vector  $m_t = [q_1, \dots, q_T] \in \mathbb{R}^{T \times 1}$  with each element  $q_i \in [0, 1] (i \in [1, T])$  indicating foreground probability of  $i$ -th snippet. This is implemented by a stack of three 1-D dynamic convolution layers  $H_m$  as follows:

$$\mathcal{M} = \text{sigmoid}(H_m(F_{vis})), \quad (8)$$



where  $t$ -th column of  $\mathcal{M}$  is the temporal mask prediction by  $t$ -th snippet. More details on dynamic convolution formulation is given in **supplementary material**.

### 3.3 Model Training and Inference

**Label assignment** To train our one-stage STALE, the ground-truth needs to be arranged into the designed format. Concretely, given a training video with temporal intervals and class labels, we label all the snippets of a single action instance with the same action class. All the snippets off from action intervals are labeled as background. For an action snippet of a particular instance in the class stream, we assign the video-length binary instance mask at the same snippet location in the action mask stream. Each mask is action instance specific. All snippets of a specific action instance share the same mask. Refer to **supplementary material** for more details.

**Learning objective** The classification stream is composed of a simple cross-entropy loss. For a training snippet, we denote  $y \in \mathbb{R}^{(K+1) \times T}$  the ground-truth class label, and  $p \in \mathcal{P}$  the classification output. We compute the classification loss as:

$$L_c = \text{CrossEntropy}(p, y). \quad (9)$$

For the segmentation mask branch, we combine a weighted cross entropy and binary dice loss [25]. Formally, for a snippet location, we denote  $\mathbf{m} \in \mathbb{R}^{T \times 1}$  the predicted segmentation mask, and  $\mathbf{g} \in \mathbb{R}^{T \times 1}$  the ground-truth mask. The loss for the segmentation mask branch is formulated as:

$$L_m = \beta_{fg} \sum_{t=1}^T \mathbf{g}(t) \log(\mathbf{m}(t)) + \beta_{bg} \sum_{t=1}^T (1 - \mathbf{g}(t)) \log(1 - \mathbf{m}(t)) \\ + \lambda_2 \left( 1 - \frac{\mathbf{m}^\top \mathbf{g}}{\sum_{t=1}^T (\mathbf{m}(t)^2 + \mathbf{g}(t)^2)} \right), \quad (10)$$

where  $\beta_{fg}/\beta_{bg}$  is the inverse of foreground/background snippet’s proportion. We set the loss trade-off coefficient  $\lambda_2 = 0.4$ .

We further impose a 1-D action completeness loss formed by binary cross entropy (BCE). It can penalize the foreground masking output  $\hat{L} \in \mathbb{R}^{T \times 1}$ . Given a ground truth one-hot foreground mask  $\hat{g} \in \mathbb{R}^{T \times 1}$ , we design the loss to model the completeness of foreground as follows:

$$L_{comp} = -\left( \hat{g} * \log(\hat{L}) + \log(\hat{g}) * \hat{L} \right). \quad (11)$$

*Inter-branch consistency* In our STALE, there is structural consistency in terms of foreground between the class and mask labels by design. To leverage this consistency for improved optimization, we formulate the consistency loss as:

$$L_{const} = 1 - \text{cosine}(\hat{F}_{clf}, \hat{F}_{mask}), \quad (12)$$

where  $\hat{F}_{clf} = \text{topk}(\text{argmax}((P_{bin} * E_p)[ : K, : ]))$  is the features obtained from the top scoring foreground snippets obtained from the thresholded classification

output  $P_{bin} := \eta(P - \theta_c)$  with  $\theta_c$  the threshold and  $E_p$  obtained by passing the embedding  $E$  into a 1D conv layer for matching the dimension of  $P$ . The top scoring features from the mask output  $M$  are obtained similarly as:  $\hat{F}_{mask} = \text{topk}(\sigma(1D\text{Pool}(E_m * M_{bin})))$  where  $M_{bin} := \eta(M - \theta_m)$  is a binarization of mask-prediction  $M$ ,  $E_m$  is obtained by passing the embedding  $E$  into a 1D conv layer for matching the dimension of  $M$ , and  $\sigma$  is sigmoid activation.

**Overall objective** The overall objective loss function of our STALE is defined as:  $L = L_c + L_m + L_{comp} + L_{const}$ . This loss is used to train the model end-to-end, whilst leaving out the pre-trained video encoder frozen due to the GPU memory constraint.

**Model inference** At test time, we generate action instance predictions for each test video by the classification  $\mathbf{P}$  and mask  $\mathbf{M}$  predictions. For  $\mathbf{P}$ , we only consider the snippets whose class probabilities are greater than  $\theta_c$  and select top scoring snippets. For each such top scoring action snippet, we then obtain the temporal mask by thresholding the  $t_i$ -th column of  $\mathbf{M}$  using the localization threshold  $\Theta$ . To produce sufficient candidates, we use a set of thresholds  $\Theta = \{\theta_i\}$ . For each candidate, we compute a confidence score  $s$  by multiplying the classification and max-mask scores. SoftNMS [3] is finally applied to obtain top scoring results.

## 4 Experiments

**Datasets** We conduct extensive experiments on two popular TAD benchmarks. (1) ActivityNet-v1.3 [4] has 19,994 videos from 200 action classes. We follow the standard setting to split all videos into training, validation and testing subsets in ratio of 2:1:1. (2) THUMOS14 [15] has 200 validation videos and 213 testing videos from 20 categories with labeled temporal boundary and action class

**Implementation details** For fair comparison with existing TAD works we use Kinetics pre-trained I3D [6] as the video encoder for both ActivityNet and THUMOS. We also use the two-stream features as used in [17] for fair-comparison. For comparing with CLIP based TAD baselines, we also adopt the image and text encoders from pre-trained CLIP (ViT-B/16+Transformer). For model-adaptation, the visual encoder is kept frozen, the trainable parts are textual prompt embeddings, text encoder, temporal embedding module, temporal masking modules and TAD decoder heads. For the CLIP encoders, the video frames are pre-processed to  $224 \times 224$  spatial resolution, and the maximum number of textual tokens is 77 (following the original CLIP design). Each video’s feature sequence  $F$  is rescaled to  $T = 100/256$  snippets for ActivityNet/THUMOS using linear interpolation. Our model is trained for 15 epochs using Adam with learning rate of  $10^{-4}/10^{-5}$  for ActivityNet/THUMOS respectively.

### 4.1 Comparative Results

**Zero-shot action detection Setting** In this section, we evaluate on the open-set scenario where  $D_{train} \cap D_{val} = \phi$ , i.e., action categories for training and

Table 1. Results of Zero-Shot Action Detection

Train Split	Methods	THUMOS14						ActivityNet v1.3			
		0.3	0.4	0.5	0.6	0.7	Avg	0.5	0.75	0.95	Avg
75% Seen 25% Unseen	B-II	28.5	20.3	17.1	10.5	6.9	16.6	32.6	18.5	5.8	19.6
	B-I	33.0	25.5	18.3	11.6	5.7	18.8	35.6	20.4	2.1	20.2
	EffPrompt	39.7	31.6	23.0	14.9	7.5	23.3	37.6	22.9	3.8	23.1
	<b>STALE</b>	<b>40.5</b>	<b>32.3</b>	<b>23.5</b>	<b>15.3</b>	<b>7.6</b>	<b>23.8</b>	<b>38.2</b>	<b>25.2</b>	<b>6.0</b>	<b>24.9</b>
50% Seen 50% Unseen	B-II	21.0	16.4	11.2	6.3	3.2	11.6	25.3	13.0	3.7	12.9
	B-I	27.2	21.3	15.3	9.7	4.8	15.7	28.0	16.4	1.2	16.0
	EffPrompt	37.2	29.6	<b>21.6</b>	<b>14.0</b>	<b>7.2</b>	21.9	32.0	19.3	2.9	19.6
	<b>STALE</b>	<b>38.3</b>	<b>30.7</b>	21.2	13.8	7.0	<b>22.2</b>	<b>32.1</b>	<b>20.7</b>	<b>5.9</b>	<b>20.5</b>

testing are disjoint. We follow the setting and dataset splits proposed by [17]. More specifically, we initiate two evaluation settings on THUMOS14 and ActivityNet1.3: (1) Training with 75% action categories and testing on the left 25% action categories; (2) Training with 50% categories and testing on the left 50% categories. To ensure statistical significance, we conduct 10 random samplings to split categories for each setting, following [17].

**Competitors** As there is no open-source implementation for [17], we use their same reported baselines. More specifically, (1) One baseline using BMN [20] as proposal generator and CLIP [36] with hand crafted prompt. This is the same baseline as reported in [17]. This serves as a 2-stage TAD baseline using CLIP. We term it B-I. (2) One comparative CLIP based TAD EffPrompt [17]. (3) One CLIP + TAD model for one-stage baseline: As ZS-TAD is a relatively new problem, we need to implement the competitors by extending existing TAD methods using CLIP by ourselves. We select a variant of CLIP for dense prediction task [37] using the existing CLIP pre-trained weights. We term this baseline, DenseCLIP (w/ CLIP Pretrained Image Encoder) + TAD, as B-II. The text-encoders however are identical for both the baselines using CLIP pre-training weights. However, we could not compare with the earlier zero-shot TAD method ZS-TAD [49] due to unavailability of code and no common data-split between [49] and [17].

**Performance** The ZS-TAD results are reported in Table 1. With 50% labeled data, our STALE surpasses both 1-stage and 2-stage baselines, as well as the CLIP based TAD method [17] by a good margin on ActivityNet. This suggests that our representation masking is able to perform the zero-shot transfer better than conventional 2-stage designs. This also indicates that the localization-error propagation hurts in low training data regime. Also noted that the one-stage baseline (B-II) performs worst among all other baselines. This suggests that a class-agnostic masking is necessary to achieve strong performance in ZS-TAD. The performance of our model however drops on THUMOS14 on stricter metrics, possibly due to the mask decoder suffering from foreground imbalance. We

**Table 2. Comparison with state-of-the-art on closed-set setting**

Methods	Mode	Encoder Backbone	THUMOS14						ActivityNet v1.3			
			0.3	0.4	0.5	0.6	0.7	Avg	0.5	0.75	0.95	Avg
TALNet	RGB+Flow	I3D	53.2	48.5	42.8	33.8	20.8	39.8	38.2	18.3	1.3	20.2
GTAN	RGB+Flow	P3D	57.8	47.2	38.8	-	-	-	52.6	34.1	8.9	34.3
MUSES	RGB+Flow	I3D	68.9	64.0	56.9	46.3	31.0	53.4	50.0	34.9	6.5	34.0
VSGN	RGB+Flow	I3D	66.7	60.4	52.4	41.0	30.4	50.1	52.3	36.0	8.3	35.0
Context-Loc	RGB+Flow	I3D	68.3	63.8	54.3	41.8	26.2	-	56.0	35.2	3.5	34.2
BU-TAL	RGB+Flow	I3D	53.9	50.7	45.4	38.0	28.5	43.3	43.5	33.9	9.2	30.1
B-III	RGB+Flow	I3D	68.3	62.3	51.9	38.8	23.7	-	47.2	30.7	8.6	30.8
<b>STALE</b>	RGB+Flow	I3D	<b>68.9</b>	<b>64.1</b>	<b>57.1</b>	<b>46.7</b>	<b>31.2</b>	<b>52.9</b>	<b>56.5</b>	<b>36.7</b>	<b>9.5</b>	<b>36.4</b>
TALNet	RGB	I3D	42.6	-	31.9	-	14.2	-	-	-	-	-
A2Net	RGB	I3D	45.0	40.5	31.3	19.9	10.0	29.3	39.6	25.7	2.8	24.8
B-I	RGB	CLIP	36.3	31.9	25.4	17.8	10.4	24.3	28.2	18.3	3.7	18.2
B-II	RGB	CLIP	57.1	49.1	40.4	31.2	23.1	40.2	51.5	33.3	6.6	32.7
EffPrompt	RGB	CLIP	50.8	44.1	35.8	25.7	15.7	34.5	44.0	27.0	5.1	27.3
<b>STALE</b>	RGB	CLIP	<b>60.6</b>	<b>53.2</b>	<b>44.6</b>	<b>36.8</b>	<b>26.7</b>	<b>44.4</b>	<b>54.3</b>	<b>34.0</b>	<b>7.7</b>	<b>34.3</b>

observe similar trend with 75% labeled data, our approach is again superior than all other competitors on both datasets. It is also interesting to note that the one-stage baseline (B-II) has a larger performance gap with the two-stage baseline (B-I), around 3.1% avg mAP on ActivityNet. This however reduces to 0.6% when the number of labeled data increase, suggesting that the one-stage baseline has potential to improve with more data.

**Closed-set action detection Setting** Closed-set action detection refers to the common setting, where the model is trained and evaluated on videos of the same action categories, *i.e.*,  $D_{train} = D_{val}$ . For a fair comparison, we use the same dataset splits as in the literature.

**Competitors** We considered the following methods for comparison. (1) Seven representative TAD methods with I3D encoder backbone; (2) One CLIP+TAD method EffPrompt [17]; (3) One two-stage CLIP based TAD baseline B-I; (4) One single-stage CLIP based TAD baseline B-II; (5) We also created another single-stage baseline by replacing CLIP-Pretrained encoder with Kinetics pre-trained Video-Encoder (*e.g.*, I3D). We term it B-III.

**Performance** From the results of Table 2, we observe that with more labeled data, our approach also surpasses existing TAD approaches often by a large margin. This is consistent over both the datasets. Thus the text-embedding is indeed helpful for our design. We also observe that our method performs better by almost similar margin when we use different feature backbone (*e.g.*, CLIP). Thus it proves our design is feature agnostic. Another key observation is that our single-stage baseline (B-II) performs significantly better than [17] by at least 5% in avg mAP for both the datasets. Thus justifies that our parallel design is better alternative for CLIP based approaches.

**Table 3.** Analysis of localization error propagation on ActivityNet with 75% split. GT: Ground-Truth.

Metric	mAP	
	0.5	Avg
Baseline-I(B-I)		
GT proposals	56.2	47.1
Predicted proposals	34.8	19.9
<b>STALE</b>		
GT masks	53.6	42.3
Predicted masks	38.2	24.9

**Table 4.** Analysis of Representation-Masking on 75% seen split on ActivityNet

Masking Method	# Queries	mAP	
		0.5	Avg
No Mask	-	21.7	10.9
GT-Mask	-	54.3	35.8
1-D CNN	-	33.5	20.1
Maskformer [8]	5	37.5	24.4
	<b>20</b>	<b>38.2</b>	<b>24.9</b>
	100	39.0	25.1

## 4.2 Ablation

**Localization error propagation analysis** To examine the effect of localization error propagation with previous TAD models, we design a proof-of-concept experiment by measuring the performance drop between ground-truth proposals and pseudo proposals. Due to unavailability of training code in [17], we carefully re-created B-I following the details in [17]. For our STALE model, we contrast ground-truth and output masks. This experiment is tested on ActivityNet with 75% label split. Table 3 shows that the proposal based TAD baseline suffers almost double performance degradation from localization (*i.e.*, proposal) error due to its sequential localization and classification design. This verifies the advantage of STALE’s parallel design.

**Necessity of representation masking** To validate the role of *representation masking* in generalization of the classifier, we perform experiments in the 75% split setting. Firstly, we compare our approach by removing the Maskformer-Decoder [8] from the pipeline of STALE and passing the temporal feature  $F_{vis}$  directly into the cross-modal adaptation module for classification. As shown in Table 4, we observe a sharp drop of 14% in avg mAP, justifying that foreground features are indeed necessary to align the text embedding. This problem is also profound in DenseCLIP baseline (B-II) as illustrated in Table 1. We also tested another alternative to masking, *e.g.*, vanilla 1-D CNN. It is evident that Maskformer-Decoder benefits from learning the query embedding over vanilla 1-D CNN. We also observe that increasing the number of queries improves the overall localization performance but comes with high memory footprint. Further, due to memory constraints of query-matching step of the decoder output during classification, we learn a MLP to reduce the cost. We verify that class-agnostic representation-masking has much stronger generalization in localizing unseen action categories. Interestingly, we observe that the maximum STALE can achieve avg mAP of 35.8% indicating a room for further improvement.

**Importance of text encoder** We ablate the effect of text encoder fine-tuning. Note the video encoder is frozen due to memory constraint. We use CLIP [36]

pre-trained text transformer (denoted as CLIP-text) in this experiment. It can be observed from Table 5 that using a text-encoder is indeed important for the overall performance. Further, fine-tuning the text encoder is also effective due to the large domain gap between CLIP pre-training and TAD task.

**Effect of temporal modeling** Recall that we use a multi-head Transformer (w/o positional encoding) for temporal modeling in STALE. We evaluate this design choice by comparing (I) a 1D CNN with 3 dilation rates (1, 3, 5) each with 2 layers, and (II) a multi-scale Temporal Convolutional Network MS-TCN [10]. Each CNN design substitutes the default Transformer while remaining all the others. Table 6 shows that the Transformer is clearly superior to both CNN alternatives. This suggests that our default design captures stronger contextual learning capability even in low-data setting like ZS-TAD.

**Table 5.** Importance of *text encoder* using CLIP vocabulary on 75% train split.

Text encoder	Fine-tune	mAP	
		0.5	Avg
No encoder	-	6.9	11.5
CLIP-text	✗	36.9	22.8
CLIP-text	✓	38.2	24.9

**Table 6.** Transformer vs. CNN on ActivityNet under 75% seen label setting.

Network	mAP	
	0.5	Avg
1D CNN	29.0	19.3
MS-TCN	33.5	21.4
Transformer	<b>38.2</b>	<b>24.9</b>

## 5 Limitation

Our proposed STALE is based upon a visual-language model pre-trained using large web image-text pair data. This could potentially incur few limitations. First, there may exist some unconstrained bias rooted with the training data. Second, there might be substantial domain gap against the target video data in fine-tuning (*e.g.*, no temporal knowledge) [31]. However, this could be naturally solved when video based visual-language models become available.

## 6 Conclusion

In this work, we have proposed a novel *Zero-Shot Temporal Action Detection via Vision-Language Prompting* (STALE) model for the under-studied yet practically useful zero-shot temporal action detection (ZS-TAD). It is characterized by a *parallel* localization (mask generation) and classification architecture designed to solve the localization error propagation problem with conventional ZS-TAD models. For improved optimization, we further introduced an inter-branch consistency regularization to exploit their structural relationships. Extensive experiments on ActivityNet and THUMOS have demonstrated that our STALE yields state-of-the-art performance under both zero-shot and supervised settings.

## References

1. Alwassel, H., Heilbron, F.C., Escorcia, V., Ghanem, B.: Diagnosing error in temporal action detectors. In: ECCV. pp. 256–272 (2018)
2. Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Zitnick, C.L., Parikh, D.: Vqa: Visual question answering. In: Proceedings of the IEEE international conference on computer vision. pp. 2425–2433 (2015)
3. Bodla, N., Singh, B., Chellappa, R., Davis, L.S.: Soft-nms–improving object detection with one line of code. In: Proceedings of the IEEE international conference on computer vision. pp. 5561–5569 (2017)
4. Caba Heilbron, F., Escorcia, V., Ghanem, B., Carlos Niebles, J.: Activitynet: A large-scale video benchmark for human activity understanding. In: CVPR. pp. 961–970 (2015)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020)
6. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
7. Chen, Y., Dai, X., Liu, M., Chen, D., Yuan, L., Liu, Z.: Dynamic convolution: Attention over convolution kernels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11030–11039 (2020)
8. Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems* **34** (2021)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: ICLR (2020)
10. Farha, Y.A., Gall, J.: Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In: CVPR. pp. 3575–3584 (2019)
11. Gao, J., Yang, Z., Chen, K., Sun, C., Nevatia, R.: Turn tap: Temporal unit regression network for temporal action proposals. In: ICCV (2017)
12. Gao, P., Geng, S., Zhang, R., Ma, T., Fang, R., Zhang, Y., Li, H., Qiao, Y.: Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544* (2021)
13. Goldberg, Y., Levy, O.: word2vec explained: deriving mikolov et al.’s negative-sampling word-embedding method. *arXiv preprint arXiv:1402.3722* (2014)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
15. Idrees, H., Zamir, A.R., Jiang, Y.G., Gorban, A., Laptev, I., Sukthankar, R., Shah, M.: The thumos challenge on action recognition for videos “in the wild”. *Computer Vision and Image Understanding* **155**, 1–23 (2017)
16. Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: International Conference on Machine Learning. pp. 4904–4916. PMLR (2021)
17. Ju, C., Han, T., Zheng, K., Zhang, Y., Xie, W.: A simple baseline on prompt learning for efficient video understanding (2022)
18. Lampert, C.H., Nickisch, H., Harmeling, S.: Attribute-based classification for zero-shot visual object categorization. *IEEE transactions on pattern analysis and machine intelligence* **36**(3), 453–465 (2013)



19. Lei, J., Li, L., Zhou, L., Gan, Z., Berg, T.L., Bansal, M., Liu, J.: Less is more: Clipbert for video-and-language learning via sparse sampling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7331–7341 (2021)
20. Lin, T., Liu, X., Li, X., Ding, E., Wen, S.: Bmn: Boundary-matching network for temporal action proposal generation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3889–3898 (2019)
21. Lin, T., Zhao, X., Su, H., Wang, C., Yang, M.: BSN: Boundary sensitive network for temporal action proposal generation. In: ECCV (2018)
22. Long, F., Yao, T., Qiu, Z., Tian, X., Luo, J., Mei, T.: Gaussian temporal awareness networks for action localization. In: CVPR (2019)
23. Lu, J., Batra, D., Parikh, D., Lee, S.: Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems* **32** (2019)
24. Miech, A., Alayrac, J.B., Smaira, L., Laptev, I., Sivic, J., Zisserman, A.: End-to-end learning of visual representations from uncurated instructional videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9879–9889 (2020)
25. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth international conference on 3D vision (3DV). pp. 565–571. IEEE (2016)
26. Nag, S., Zhu, X., Song, Y.Z., Xiang, T.: Temporal action localization with global segmentation mask transformers (2021)
27. Nag, S., Zhu, X., Song, Y.z., Xiang, T.: Proposal-free temporal action detection via global segmentation mask learning. In: ECCV (2022)
28. Nag, S., Zhu, X., Song, Y.z., Xiang, T.: Semi-supervised temporal action detection with proposal-free masking. In: ECCV (2022)
29. Nag, S., Zhu, X., Xiang, T.: Few-shot temporal action localization with query adaptive transformer. *arXiv preprint arXiv:2110.10552* (2021)
30. Niu, L., Cai, J., Veeraraghavan, A., Zhang, L.: Zero-shot learning via category-specific visual-semantic mapping and label refinement. *IEEE Transactions on Image Processing* **28**(2), 965–979 (2018)
31. Pan, J., Lin, Z., Zhu, X., Shao, J., Li, H.: Parameter-efficient image-to-video transfer learning. *arXiv preprint arXiv:2206.13559* (2022)
32. Parikh, D., Grauman, K.: Relative attributes. In: 2011 International Conference on Computer Vision. pp. 503–510. IEEE (2011)
33. Paul, S., Mithun, N.C., Roy-Chowdhury, A.K.: Text-based localization of moments in a video corpus. *IEEE Transactions on Image Processing* **30**, 8886–8899 (2021)
34. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP). pp. 1532–1543 (2014)
35. Qin, J., Liu, L., Shao, L., Shen, F., Ni, B., Chen, J., Wang, Y.: Zero-shot action recognition with error-correcting output codes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2833–2842 (2017)
36. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
37. Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: Denseclip: Language-guided dense prediction with context-aware prompting. *arXiv preprint arXiv:2112.01518* (2021)

38. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: towards real-time object detection with region proposal networks. *TPAMI* **39**(6), 1137–1149 (2016)
39. Sridhar, D., Quader, N., Muralidharan, S., Li, Y., Dai, P., Lu, J.: Class semantics-based attention for action detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 13739–13748 (2021)
40. Su, H., Gan, W., Wu, W., Qiao, Y., Yan, J.: Bsn++: Complementary boundary regressor with scale-balanced relation modeling for temporal action proposal generation. *arXiv preprint arXiv:2009.07641* (2020)
41. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *arXiv preprint arXiv:1706.03762* (2017)
42. Wang, M., Xing, J., Liu, Y.: Actionclip: A new paradigm for video action recognition. *arXiv preprint arXiv:2109.08472* (2021)
43. Wang, Z., Liu, X., Li, H., Sheng, L., Yan, J., Wang, X., Shao, J.: Camp: Cross-modal adaptive message passing for text-image retrieval. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5764–5773 (2019)
44. Xian, Y., Lampert, C.H., Schiele, B., Akata, Z.: Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE transactions on pattern analysis and machine intelligence* **41**(9), 2251–2265 (2018)
45. Xu, H., Das, A., Saenko, K.: R-c3d: Region convolutional 3d network for temporal activity detection. In: *ICCV* (2017)
46. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: *International conference on machine learning*. pp. 2048–2057. PMLR (2015)
47. Xu, M., Zhao, C., Rojas, D.S., Thabet, A., Ghanem, B.: G-tad: Sub-graph localization for temporal action detection. In: *CVPR* (2020)
48. Xu, X., Hospedales, T., Gong, S.: Semantic embedding space for zero-shot action recognition. In: *2015 IEEE International Conference on Image Processing (ICIP)*. pp. 63–67. IEEE (2015)
49. Zhang, L., Chang, X., Liu, J., Luo, M., Wang, S., Ge, Z., Hauptmann, A.: Zstad: Zero-shot temporal activity detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 879–888 (2020)
50. Zhang, R., Fang, R., Gao, P., Zhang, W., Li, K., Dai, J., Qiao, Y., Li, H.: Tip-adapter: Training-free clip-adapter for better vision-language modeling. *arXiv preprint arXiv:2111.03930* (2021)
51. Zhao, C., Thabet, A.K., Ghanem, B.: Video self-stitching graph network for temporal action localization. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 13658–13667 (2021)
52. Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., Lin, D.: Temporal action detection with structured segment networks. In: *ICCV* (2017)
53. Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., et al.: Regionclip: Region-based language-image pretraining. *arXiv preprint arXiv:2112.09106* (2021)
54. Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. *arXiv preprint arXiv:2109.01134* (2021)