

# OadTR: Online Action Detection with Transformers

Xiang Wang<sup>1</sup> Shiwei Zhang<sup>2\*</sup> Zhiwu Qing<sup>1</sup> Yuanjie Shao<sup>1</sup> Zhengrong Zuo<sup>1</sup>  
Changxin Gao<sup>1</sup> Nong Sang<sup>1\*</sup>

<sup>1</sup>Key Laboratory of Image Processing and Intelligent Control,  
School of Artificial Intelligence and Automation,  
Huazhong University of Science and Technology, China

<sup>2</sup>Alibaba Group, China

{wxiang,qzw,shaoyuanjie,zhrzuo,cgao,nsang}@hust.edu.cn, zhangjin.zsw@alibaba-inc.com

## Abstract

Most recent approaches for online action detection tend to apply Recurrent Neural Network (RNN) to capture long-range temporal structure. However, RNN suffers from non-parallelism and gradient vanishing, hence it is hard to be optimized. In this paper, we propose a new encoder-decoder framework based on Transformers, named OadTR, to tackle these problems. The encoder attached with a task token aims to capture the relationships and global interactions between historical observations. The decoder extracts auxiliary information by aggregating anticipated future clip representations. Therefore, OadTR can recognize current actions by encoding historical information and predicting future context simultaneously. We extensively evaluate the proposed OadTR on three challenging datasets: HDD, TVSeries, and THUMOS14. The experimental results show that OadTR achieves higher training and inference speeds than current RNN based approaches, and significantly outperforms the state-of-the-art methods in terms of both mAP and mcAP. Code is available at <https://github.com/wangxiang1230/OadTR>.

## 1. Introduction

The purpose of online action detection is to correctly identify ongoing actions from streaming videos without any access to the future. Recently, this task has received increasing attention due to its great potential of diverse application prospects in real life, such as autonomous driving [24], video surveillance [45], anomaly detection [37,38], etc. The crucial challenge of this task is that we need to detect the actions at the moment that video frames arrive with inadequate observations. To solve the problem, it is important to learn the long-range temporal dependencies.

\* Corresponding authors.

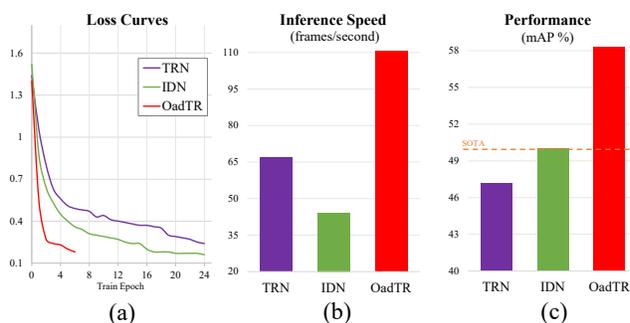


Figure 1: Comparison between OadTR and the state-of-the-art online action detection methods (*i.e.*, TRN [54] and IDN [15]): (a) Comparison of the training speeds; (b) Comparison of the inference speeds; (c) Comparison of the performance on the challenge THUMOS14 dataset.

Current approaches tend to apply RNN to model the temporal dependencies and have achieved impressive improvements [11, 13, 15, 16, 28, 54]. Typically, Information Discrimination Network (IDN) [15] designs a RNN like architecture to encode long-term historical information, and then conduct action recognition at current moment. However, RNN like architectures have the problems of non-parallelism and gradient vanishing [33, 39]. Thus it is hard to optimize the architectures, which may result in an unsatisfactory performance. This is a challenging problem for current approaches, but much less effort has been paid to solve it. To further improve the performance, we need to design a new efficient and easily-optimized framework. For this purpose, we propose to apply Transformers [48]. Transformers possess the strong power of long-range temporal modeling by the self-attention module, and have achieved remarkable performance in both natural language processing [12, 48, 57] and various vision tasks [14, 58]. Existing works have proved that Transformers have a better convergence than RNN architectures [23, 26], and they are also computationally efficient. The above properties of Trans-

formers can naturally provide alternative scheme for the online action detection task.

The above observations motivate this work. In particular, we propose a carefully designed framework, termed OadTR, by introducing the power of Transformers to the online action detection task, as illustrated in Figure 2. The proposed OadTR is an encoder-decoder architecture which can simultaneously learn long-range historical relationships and future information to classify current action. The first step is to extract clip-level feature sequence from a given video by a standard CNN. We then embed a task token to the clip-level feature sequence and input them to the encoder module. By this means, the output of the task token can encode the global temporal relationships among the historical observations. In contrast, the decoder is designed to predict the actions that may take place in the future moments. Finally, we concatenate the outputs of both the task token and decoder to detect the online actions. We compare OadTR with other RNN based approaches in Figure 1, which shows that the proposed OadTR is both efficient and effective. To further demonstrate the effectiveness of OadTR, we conduct a large number of experiments on three public datasets, including HDD [41], TVSeries [10], and THUMOS14 [22], and achieve significant improvements in terms of both mAP and mcAP metrics.

Summarily, we make the following three contributions:

- To the best of our knowledge, we are the first to bring Transformers into online action detection task and propose a new framework, *i.e.*, OadTR;
- We specially design the encoder and decoder of OadTR which can aggregate long-range historical information and future anticipations to improve online action detection;
- We conduct extensive experiments, and the results demonstrate the proposed OadTR significantly outperforms state-of-the-art methods. The massive and comprehensive ablation studies can further dissect the underlying properties of OadTR.

## 2. Related Work

In this section, we will review the related methods of our approach as follows:

**Online Action Detection.** Given a live video stream, online action detection aims to identify the actions that are taking place, even if only part of the actions can be observed. De Geest *et al.* [10] explicitly introduced online action detection task for the first time and proposed TVSeries dataset. After that, they also proposed a two-stream feedback network with LSTM [20] to model temporal structure [11]. RED [16] designs a Reinforced Encoder-Decoder network

and a module to encourage making the right decisions as early as possible. Several recent methods [17, 44] focus on detecting action starts and minimizing the time delay of identifying the start point of an action. IDN [15] directly manipulates the GRU cell [8] to model the relations between past information and an ongoing action. Inspired by that humans often identify current actions by considering the future [9], TRN [54] uses LSTM to predict future information recursively and combine it with past observations to identify actions. In this paper, we also introduce future information to assist in identifying current actions, but in parallel. Note that the aforementioned methods adopt RNN to model input action sequences, which are inefficient and lack of interaction between features, resulting in poor modeling capabilities for long-term dependence.

**Temporal Action Detection.** The goal of temporal action detection is to locate the start time points and end time points of all action instances in the untrimmed video. One-stage methods [27, 31, 50] draw on the SSD [35] method in object detection and design end-to-end action detection networks with multi-layer feature pyramid structures. Two-stage methods [7, 53] adopt the Faster-RCNN [42] architecture, including proposal generation subnet and proposal classification subnet. Most recent methods focus on generating high-quality temporal action proposals. These methods [2, 29, 30, 32, 40, 51, 55] usually locate temporal boundaries with high probabilities, then combine these boundaries as proposals and retrieve proposals by evaluating the confidence of whether a proposal contains an action within its region. However, the above approaches assume the entire input video can be observed, which is not available in the online task.

**Transformers.** Since the success of the Transformer-based models in the natural language processing field [12, 48, 57], there are many attempts to explore the feasibility of Transformers in vision tasks. DETR [5] and its variants [58, 61] effectively remove the need for many hand-designed components like non-maximum suppression procedures and anchor generation by adopting Transformers. ViT [14] divides an image into  $16 \times 16$  patches and feeds them into a standard Transformer’s encoder. There are also some attempts in semantic segmentation [52, 59], lane shape prediction [34], video frame synthesis [36], *etc.* To the best of our knowledge, we are the first to introduce Transformers into the online action detection task. In particular, unlike the original auto-regressive Transformer, OadTR adopts a non-auto-regressive Transformer to generate sequences in parallel to improve efficiency.

## 3. Methodology

In this section, we will first introduce the problem definition and then present the proposed OadTR detailedly.

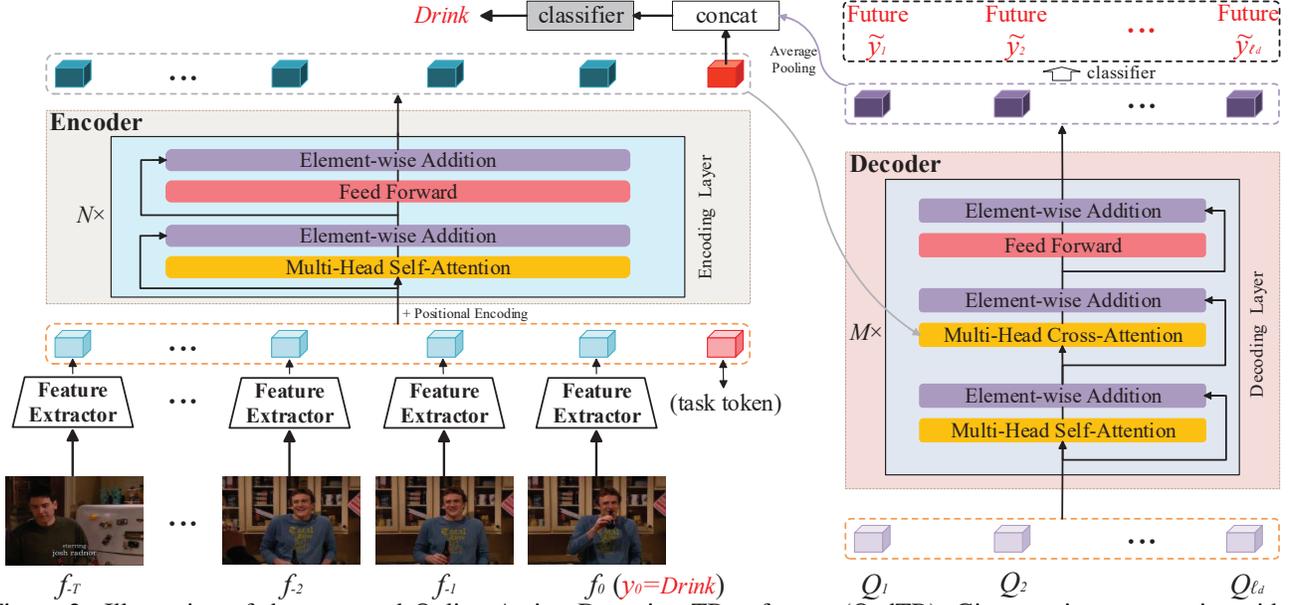


Figure 2: Illustration of the proposed Online Action Detection TRansformer (OadTR). Given an input streaming video  $\mathbf{V} = \{f_t\}_{t=-T}^0$ , a task token is attached to the visual features output by the feature extraction network. Then the token feature sequence is input into the standard Transformer’s encoder to model long-range historical temporal dependencies. Afterward, the decoder of OadTR anticipates the future context information in parallel. Finally, the predicted future context are involved in classifying the current action. Note that OadTR, including the encoder and decoder, is an end-to-end parallel framework.

### 3.1. Problem description

Given a video stream that may contain multiple actions, the goal of the task is to identify the actions currently taking place in real-time. We denote  $\mathbf{V} = \{f_t\}_{t=-T}^0$  as the input streaming video, which needs to classify the current frame chunk  $f_0$ . We use  $y_0$  to represent the action category of the current frame chunk  $f_0$ , and  $y_0 \in \{0, 1, \dots, C\}$ , where  $C$  is the total number of the action categories and index 0 denotes the background category.

### 3.2. OadTR

To explore the potential benefits of the Transformer, we introduce the power of self-attention into the online action detection task. In this section, we present OadTR, a encoder-decoder conformation. OadTR adopts an attention mechanism to capture long-range contextual information in the temporal dimension of features. The schematic diagram of OadTR is shown in Figure 2.

#### 3.2.1 Encoder

Given a streaming video  $V = \{f_t\}_{t=-T}^0$ , the feature extractor [49] extracts a 1D feature sequence by collapsing the spatial dimensions. Then an additional linear projection layer further maps each vectorized frame chunk feature into a  $D$ -dimensional feature space and  $F = \{token_t\}_{t=-T}^0 \in \mathbb{R}^{(T+1) \times D}$  denotes the resulting token sequence.

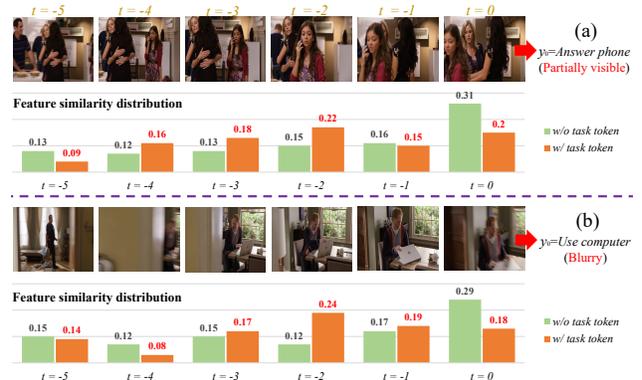


Figure 3: Comparison of similarity distribution between the classification features (*i.e.*, features before sending to classifier) and the input features sequence  $\tilde{F}$ . Note that *w/o task token*: the output classification features correspond to the token of the  $f_0$  input; *w/ task token*: the output classification features correspond to the task token.

In the encoder, we extend a learnable  $token_{class} \in \mathbb{R}^D$  to the embedded feature sequence  $F$  and get the combined token feature sequence  $\tilde{F} = Stack(\{token_t\}_{t=-T}^0, token_{class}) \in \mathbb{R}^{(T+2) \times D}$ . Note that  $token_{class}$  is used to learn global discriminative features related to the online action detection task. Intuitively, if there is no  $token_{class}$  here, the final feature representation obtained by other tokens will inevitably be biased towards this specified token as a whole, and thus

cannot be used to represent this learning task (*i.e.*, *w/o* task token in Figure 3). In contrast, the semantic embedding of  $token_{class}$  can be obtained by adaptively interacting with other tokens in the encoder, which is more suitable for feature representations (*i.e.*, *w/* task token in Figure 3). We will further confirm the necessity of  $token_{class}$  in Sec 4.3.

Since there is no frame order information in the encoder, we need to embed position encoding additionally. Position encoding can take two forms: sinusoidal inputs and trainable embeddings. We add position encoding  $E_{pos} \in \mathbb{R}^{(T+2) \times D}$  to the token sequence (*i.e.*, element-wise addition) to retain positional information:

$$X_0 = \tilde{F} + E_{pos} \quad (1)$$

In this way, positional information can be kept despite the orderless self-attention.

Multi-head self-attention (MSA) is the core component of the Transformer. Intuitively, the idea behind self-attention is that each token can interact with other tokens and can learn to gather useful semantic information more effectively, which is very suitable for capturing long-range dependencies. We compute the dot products of the query with all keys and apply a softmax function to obtain the weights on the values. Generally, the formula for self-attention is defined as:

$$X' = \text{Norm}(X_0) \quad (2)$$

$$\text{Attention}(Q_i; K_i; V_i) = \text{softmax}\left(\frac{Q_i K_i^T}{\sqrt{d_k}}\right) V_i \quad (3)$$

$$H_i = \text{Attention}(Q_i; K_i; V_i) \quad (4)$$

where  $Q_i = X' \mathbf{W}_i^q$ ,  $K_i = X' \mathbf{W}_i^k$  and  $V_i = X' \mathbf{W}_i^v$  are linear layers applied on input sequence, and  $\mathbf{W}_i^q, \mathbf{W}_i^k, \mathbf{W}_i^v \in \mathbb{R}^{D \times \frac{D}{N_{head}}}$ . Note that queries, keys, and values are all vectors, and  $N_{head}$  is the number of heads.  $\frac{1}{\sqrt{d_k}}$  is a scaling factor, and  $d_k$  is typically set to  $D$ . The scaling factor can make the training more stable and accelerate convergence. Subsequently, the outputs of heads  $H_1, H_2, \dots, H_{N_{head}}$  are concatenated and feed into a linear layer. The formula is as follows:

$$\hat{H} = \text{Stack}(H_1, H_2, \dots, H_{N_{head}}) \mathbf{W}_d \in \mathbb{R}^{(T+2) \times D} \quad (5)$$

where  $\mathbf{W}_d$  is a linear projection. Multi-head self-attention allows the encoder to focus on multiple different patterns, which is beneficial to improve the robustness and capacity of the encoder.

Subsequently, it is followed by a two-layered feed-forward network (FFN) with GELU [19] activations. Meanwhile, layernorm [1] and residual connections [18] are also applied. The final multiple formulas can be expressed as:

$$\hat{H} = \text{MSA}(\text{Norm}(X_0)) \quad (6)$$

$$m'_1 = \hat{H} + X_0 \quad (7)$$

$$m_n = \text{FFN}(\text{Norm}(m'_{n-1})) + m'_{n-1} \quad (8)$$

$$m'_n = \text{MSA}(\text{Norm}(m_{n-1})) + m_{n-1} \quad (9)$$

where  $n = 1, 2, \dots, N$ ,  $N$  is the number of encoding layers and  $m_N \in \mathbb{R}^{(T+2) \times D}$  denotes the final output feature representations of the encoder. For the convenience of explanation, we use  $m_N^{token} \in \mathbb{R}^D$  to signify the output representation of the encoder corresponding to the task token.

### 3.2.2 Decoder

When a person is watching a movie, he will not only remember the past, but also make predictions about what will happen in the near future [3]. Consequently, the decoder of OadTR makes use of the observation of past information to predict the actions that will occur in the near future, so as to better learn more discriminative features.

Prediction query  $Q_i \in \mathbb{R}^{D'}$ ,  $i = 1, 2, \dots, \ell_d$  is also learnable, where  $D'$  is the number of query channels. The difference with the original Transformer [48] is that our decoder decodes the  $\ell_d$  prediction queries in parallel at each decoding layer. The decoder is allowed to utilize semantic information from the encoder through the encoder-decoder cross-attention mechanism.

Here, we use  $\tilde{Q}_i \in \mathbb{R}^{D'}$ ,  $i = 1, 2, \dots, \ell_d$  to represent the sequence output after the decoder.

### 3.2.3 Training

In OadTR, we mainly use the encoder to identify the current frame chunk  $f_0$ , and the decoder to predict the coming future. At the same time, the prediction results are used as auxiliary information to better recognize the action.

For the classification task of the current frame chunk, we first concatenate the task-related features in the encoder with the pooled predicted features in the decoder. Then the resulting features go through a full connection layer and a softmax operation for action classification:

$$\tilde{Q} = \text{Avg-pool}(\tilde{Q}_1, \tilde{Q}_2, \dots, \tilde{Q}_{\ell_d}) \quad (10)$$

$$p_0 = \text{softmax}(\text{Concat}[m_N^{token}, \tilde{Q}] \mathbf{W}_c) \quad (11)$$

where  $\mathbf{W}_c$  represents full connection layer parameters for classification, and  $p_0 \in \mathbb{R}^{C+1}$ .

In addition to the estimated current action, OadTR also outputs predicted features for the next  $\ell_d$  time steps. Since future information is available during offline training, in order to ensure that a good feature expression is learned, we also conduct supervised training on future prediction features:

$$\tilde{p}^i = \text{softmax}(\tilde{Q}_i \mathbf{W}'_c), i = 1, 2, \dots, \ell_d \quad (12)$$

Method	Reference	Input	mAP (%)
CNN [10]	ECCV'16	Sensors	22.7
LSTM [41]	CVPR'18		23.8
ED [16]	BMVC'17		27.4
TRN [54]	ICCV'19		29.2
<b>OadTR</b>	-		<b>29.8</b>

Table 1: Comparison between our OadTR and other state-of-the-art online action detection methods on the HDD [41] dataset in terms of mAP (%).

Method	Reference	Input	mcAP (%)
RED [16]	BMVC'17	TSN-Anet	79.2
TRN [54]	ICCV'19		83.7
IDN [15]	CVPR'20		84.7
<b>OadTR</b>	-	-	<b>85.4</b>
IDN [15]	CVPR'20	TSN-Kinetics	86.1
<b>OadTR</b>	-		<b>87.2</b>

Table 2: Comparison between our OadTR and other state-of-the-art online action detection methods on the TVSeries [10] dataset in terms of mcAP (%). Note that we use the same two-stream features for fair comparison.

Method	Reference	Setting	mAP (%)
CNN [47]	ICLR'15	Offline	34.7
CNN [46]	NIPS'14		36.2
LRCN [13]	CVPR'15		39.3
MultiLSTM [56]	IJCV'18		41.3
CDC [43]	CVPR'17		44.4
RED [16]	BMVC'17	Online (TSN-Anet)	45.3
TRN [54]	ICCV'19		47.2
IDN [15]	CVPR'20		50.0
<b>OadTR</b>	-	-	<b>58.3</b>
IDN [15]	CVPR'20	Online (TSN-Kinetics)	60.3
<b>OadTR</b>	-		<b>65.2</b>

Table 3: Performance comparison on the THUMOS14 [22] dataset in terms of mAP (%). OadTR, IDN [15], TRN [54], and RED [16] use the same two-stream features.

Therefore, the final joint training loss is:

$$Loss = CE(p_0, y_0) + \lambda \sum_{i=1}^{\ell_d} CE(\tilde{p}_i, \tilde{y}_i) \quad (13)$$

where CE is the cross entropy loss,  $\tilde{y}_i$  is the actual action category for the next step  $i$  and  $\lambda$  is a balance coefficient, set to 0.5 in the experiment.

## 4. Experiments

In this section, we evaluate the proposed OadTR on three benchmark datasets: HDD [41], TVSeries [10], and THUMOS14 [22] for online action detection. First, we compare the results between our OadTR and start-of-the-art methods. Then, we conduct more detailed ablation studies to evaluate the effectiveness of OadTR.

### 4.1. Dataset and setup

**HDD.** This dataset includes approximately 104 hours of driving actions collected in the San Francisco Bay Area, and there are a total of 11 action categories. The dataset consists

of 137 sections and provides various non-visual sensors collected by the vehicle's controller area network bus. We use 100 sections for training and 37 sections for evaluation.

**TVSeries.** TVSeries contains six popular TV series, about 150 minutes for each, about 16 hours in total. The dataset totally includes 30 actions, and every action occurs at least 50 times in the dataset. TVSeries contains many unconstrained perspectives and a wide variety of backgrounds.

**THUMOS14.** This dataset has 1010 validation videos and 1574 testing videos with 20 classes. For the online action detection task, there are 200 validation videos and 213 testing videos labeled with temporal annotations. As in the previous works [15, 54], we train our model on the validation set and evaluate on the test set.

**Implementation details.** For feature extractor, following previous works [15, 16, 54], we adopt the two-stream network [49] (3072-dimensions) pre-trained on ActivityNet v1.3 [4] (TSN-Anet), where spatial and temporal sub-networks adopt ResNet-200 [18] and BN-Inception [21] separately. For fair quantitative comparison with [15], we also conduct experiments with the same TSN features (4096-dimensions) pre-trained on Kinetics [6] (TSN-Kinetics).

In terms of training, we implement our proposed OadTR in PyTorch and conduct all experiments with Nvidia V100 graphics cards. Without those bells and whistles, we use Adam [25] for optimization, the batch size is set to 128, the learning rate is set to 0.0001, and weight decay is 0.0005. Unless otherwise specified, we set  $T$  to 31 for HDD dataset and 63 for both TVSeries and THUMOS14 dataset.

**Evaluation metric.** To evaluate the performance of OadTR, following the previous methods [10, 15, 16, 54], we report per-frame mean Average Precision (mAP) on HDD and THUMOS14 dataset, and per-frame mean calibrated Average Precision [10] (mcAP) on TVSeries. The mAP is widely used, which needs to average the average precision (AP) of each action class. The calibrated Average Precision (cAP) can be formulated as:

$$cPrec = \frac{TP}{TP + \frac{FP}{w}} \quad (14)$$

$$cAP = \frac{\sum_k cPrec(k) \times I(k)}{\sum TP} \quad (15)$$

where  $I(k)$  is equal to 1 if frame  $K$  is a TP. The coefficient  $w$  is the ratio between negative and positive frames.

### 4.2. Compared with state-of-the-art methods

To evaluate the performance, we compare our proposed OadTR and other state-of-the-art methods [15, 16, 54] on HDD, TVSeries, and THUMOS14 datasets. Noted that we use the same network parameter settings (e.g.,  $N = 3$ ,  $M = 5$ ) when comparing different datasets. As illus-

Method	Portion of action									
	0%-10%	10%-20%	20%-30%	30%-40%	40%-50%	50%-60%	60%-70%	70%-80%	80%-90%	90%-100%
CNN [10]	61.0	61.0	61.2	61.1	61.2	61.2	61.3	61.5	61.4	61.5
LSTM [10]	63.3	64.5	64.5	64.3	65.0	64.7	64.4	64.4	64.4	64.3
FV-SVM [10]	67.0	68.4	69.9	71.3	73.0	74.0	75.0	75.4	76.5	76.8
TRN [54]	78.8	79.6	80.4	81.0	81.6	81.9	82.3	82.7	82.9	83.3
IDN [15]	<b>80.6</b>	81.1	81.9	82.3	82.6	82.8	82.6	82.9	83.0	83.9
<b>OadTR</b>	79.5	<b>83.9</b>	<b>86.4</b>	<b>85.4</b>	<b>86.4</b>	<b>87.9</b>	<b>87.3</b>	<b>87.3</b>	<b>85.9</b>	<b>84.6</b>
IDN (Kinetics) [15]	<b>81.7</b>	81.9	83.1	82.9	83.2	83.2	83.2	83.0	83.3	86.6
<b>OadTR (Kinetics)</b>	81.2	<b>84.9</b>	<b>87.4</b>	<b>87.7</b>	<b>88.2</b>	<b>89.9</b>	<b>88.9</b>	<b>88.8</b>	<b>87.6</b>	<b>86.7</b>

Table 4: Performance comparison for different portions of actions on the TVSeries dataset in terms of mcAP (%). Note that the corresponding portions of actions are only used to compute mcAP after detecting current actions on all frame chunks in an online manner.

# Method	mAP (%)	# Method	mcAP (%)	# Method	mAP (%)
LSTM [41]	23.8	LSTM [15]	80.9	LSTM [15]	46.3
Encoder-only (Baseline)	28.7	Encoder-only (Baseline)	84.8	Encoder-only (Baseline)	55.8
Baseline + TT	29.0	Baseline + TT	85.0	Baseline + TT	56.9
Baseline + DE	29.2	Baseline + DE	85.1	Baseline + DE	56.7
Baseline + TT + DE (OadTR)	<b>29.8</b>	Baseline + TT + DE (OadTR)	<b>85.4</b>	Baseline + TT + DE (OadTR)	<b>58.3</b>

(a) Ablation study of the effectiveness of our proposed components on the HDD dataset.

(b) Ablation study of the effectiveness of our proposed components on the TVSeries dataset.

(c) Ablation study of the effectiveness of our proposed components on the THUMOS14 dataset.

# Position encoding	mAP (%)	# Head	mAP (%)	# Dim	mAP (%)	# Generalization	m(c)AP (%)
No Position	28.8	Head = 1	57.4	Dim = 128	56.4	Vanilla Transformer (HDD)	29.8
Fixed Position	29.3	Head = 2	57.7	Dim = 256	57.1	Sparse Transformer (HDD)	29.6
Learned Position	<b>29.8</b>	Head = 4	<b>58.3</b>	Dim = 512	57.3	Vanilla Transformer (TVSeries)	85.4
		Head = 8	57.7	Dim = 1024	<b>58.3</b>	Sparse Transformer (TVSeries)	85.0
		Head = 16	57.8	Dim = 1536	57.7	Vanilla Transformer (THUMOS14)	58.3
		Head = 32	57.4	Dim = 2048	57.6	Sparse Transformer (THUMOS14)	58.1

(d) Ablation study of different position encoding manners on the HDD dataset.

(e) Ablation study of the head number on the THUMOS14 dataset.

(f) Ablation study of the query dimensions on the THUMOS14 dataset.

(g) Generalization evaluation. The results indicate the generalization and superiority of our model design.

Table 5: Ablation studies.

trated in Table 1, our OadTR achieves state-of-the-art performance and improves mAP from 29.2% to 29.8% on the HDD dataset, demonstrating that our OadTR can achieve an overall performance promotion of online action detection. It could be attributed to that our OadTR introduces Transformers to obtain global historical information and the future context efficiently.

Table 2 compares recent online action detection methods on the TVSeries dataset. To ensure a fair comparison, we adopt the same video features. The results signify that our OadTR can achieve good performance in the case of different video feature inputs. The reason for the better results of TSN-Kinetics may be that the categories of Kinetics are more diverse and contains much common generalizable presentations.

As shown in Table 3, we also conduct a comprehensive comparison on the THUMOS14 dataset. Specifically, performance is improved by 8.3% (50.0% vs. 58.3%) under TSN-Anet feature input and 4.9% (60.3% vs. 65.2%) under TSN-Kinetics feature input. The above results indicate the effectiveness of our proposed OadTR.

When only a fraction of each action is considered, we compared OadTR with previous methods. Table 4 shows that our OadTR significantly outperforms state-of-the-art methods [15, 54] at most time stages. Specifically, this indicates the superiority of OadTR in recognizing actions at early stages as well as all stages. It could be attributed to the ability of our OadTR to efficiently model temporal dependencies.

### 4.3. Ablation studies

To facilitate our analysis of the model, we take the encoder without task token as our baseline. We further conduct detailed ablation studies to evaluate different components of the proposed framework, include the following:

**Encoder-only (Baseline):** We adopt the encoder in the original Transformer [48] and apply it directly to the online action detection task. Note that compared with our OadTR’s encoder, the task token is missing, and the classifier is applied to the last output representation (*i.e.*, corresponding to  $f_0$ ) of the Transformer’s encoder.

**Baseline + TT:** We incorporate *Baseline* and the task token

Dataset	Task	Decoder steps ( $\ell_d$ )			
		2	4	8	16
HDD	Online Action Detection	28.3	28.4	<b>29.8</b>	28.8
	Action Anticipation	27.2	26.2	23.0	16.9
TVSeries	Online Action Detection	85.3	85.2	<b>85.4</b>	85.2
	Action Anticipation	81.8	80.4	77.8	74.3
THUMOS14	Online Action Detection	57.4	57.8	<b>58.3</b>	58.0
	Action Anticipation	53.5	51.0	45.9	40.7

Table 6: Online action detection and action anticipation results of our proposed OadTR with decoder steps  $\ell_d = 2, 4, 8, 16$ .

Method	Dataset	Encoding layers ( $N, M = 5$ )					
		1	2	3	4	5	6
OadTR	HDD	28.8	29.0	<b>29.8</b>	29.5	28.4	28.4
	TVSeries	85.4	<b>85.5</b>	85.4	85.3	84.8	85.0
	THUMOS14	57.5	56.9	<b>58.3</b>	57.4	56.9	56.6

(a) Online action detection results of different encoding layers. Note that we fixed  $M = 5$  for simplicity.

Method	Dataset	Decoding layers ( $M, N = 3$ )					
		1	2	3	4	5	6
OadTR	HDD	27.7	28.0	29.5	29.5	<b>29.8</b>	28.6
	TVSeries	84.8	84.8	85.2	85.4	85.4	<b>85.5</b>
	THUMOS14	57.4	57.9	56.9	57.2	<b>58.3</b>	56.6

(b) Online action detection results of different decoding layers. Note that we fixed  $N = 3$  for simplicity.

Table 7: Ablation study of encoding layers  $N$  and decoding layers  $M$  using TSN-Anet features.

(TT) together, which is our OadTR’s encoder. This method adds a task-related token, and we use ablation experiments to illustrate the necessity of this token.

**Baseline + DE:** In this method, we add the decoder (DE) of the prediction task in OadTR to the baseline method to test and verify the function of the decoder.

**Baseline + TT + DE (OadTR):** This is the method we propose in this paper, adding task token and decoder to the baseline method together.

In Table 5 (a-c), we report the performance comparison experiments of the above methods on HDD, TVSeries, and THUMOS14 datasets. The results of *Baseline + TT* demonstrate that using the additional task token is helpful for action classification. The results of *Baseline + DE* explain the remarkable power of the auxiliary prediction task. When combined with the above two improvements, our OadTR (*i.e.*, *Baseline + TT + DE*) achieves the best results on three datasets. Specifically, when compared with the baseline method, our OadTR approach has improved by 1.1%, 0.6%, and 2.5% on HDD, TVSeries, and THUMOS14 datasets, respectively.

**Importance of position encoding.** To demonstrate the importance of using position encoding, we organized some comparative experiments. As shown in Table 5 (d), position encoding is necessary, and the learnable position encoding achieves the best result.

**Effect of head number.** Multi-head self-attention is a criti-

cal component. Here we study the impact of different head numbers of the decoder on performance. We can find that a optimal result is achieved when the number of heads is 4 (Table 5 (e)).

**Effect of query dimensions.** We conduct experiments to examine how different query dimensions affect online action detection performance. As shown in Table 5 (f), when the number of feature dimensions is relatively small (*e.g.*, 128), the model capacity is limited and the performance is relatively poor. As the number of feature dimensions gradually increases, the model capacity increases, and performance improves. However, when it exceeds a specific value (*e.g.*, 1024), over-fitting may occur.

**Generalizability.** To further investigate the generalization of our OadTR for many Transformer variants, we replace the standard Transformer [48] with Sparse Transformer [60]. As illustrated in Table 5 (g), we can find that the performance is still good after replacement. Generally, Sparse Transformer can reduce computational consumption but leads to a little performance degradation.

**Effect of decoder step count.** The step size used to predict the future also has an impact on performance. In Table 6, we compare four step sizes (*i.e.*, 2, 4, 8, and 16), and the results denote that  $steps = 8$  achieves the best results on the three datasets.

**Effect of encoding layers  $N$  and decoding layers  $M$ .** To further study the impact of different encoding and decoding layers on performance, additional experiments are conducted, and results are shown in Table 7. In most cases, the best results are achieved when  $N = 3, M = 5$ . However, there will also be fluctuations, such as on the TVSeries dataset.

**Feature aggregation type.** We also conduct experiments to explore different types that aggregate future and current features. We can notice that Avg-pool is better than Max-pool (Table 9). The reason may be that the predicted deep semantic representations of different time steps all have a specific promotion effect on the current classification. Simultaneously, the results of *w/o encoder* also indicate the necessity of the encoder to learn discriminative features.

#### 4.4. Action anticipation

In our proposed OadTR, we introduce predicted future information to identify current actions. To evince the accuracy of our predictions, we also conduct experiments to compare with other methods. Table 8 demonstrates that OadTR outperforms the current state-of-the-art method [54] by a large margin. In particular, the performance of OadTR is 2.1% higher than TRN [54] on the TVSeries dataset and 7.0% higher on THUMOS14. Furthermore, the performance of OadTR can also be further improved when pre-training on Kinetics.

Method	Time predicted into the future (seconds)								Avg
	0.25s	0.5s	0.75s	1.0s	1.25s	1.5s	1.75s	2.0s	
ED [16]	78.5	78.0	76.3	74.6	73.7	72.7	71.7	71.0	74.5
RED [16]	79.2	78.7	77.1	75.5	74.2	73.0	72.0	71.2	75.1
TRN [54]	79.9	78.4	77.1	75.9	74.9	73.9	73.0	72.3	75.7
<b>OadTR</b>	81.9	80.6	79.4	78.2	77.1	76.0	75.2	74.3	<b>77.8</b>
<b>OadTR (Kinetics)</b>	84.1	82.6	81.3	80.1	78.9	77.7	76.7	75.7	<b>79.1</b>

(a) Results on the TVSeries dataset in terms of mcAP (%).

Method	Time predicted into the future (seconds)								Avg
	0.25s	0.5s	0.75s	1.0s	1.25s	1.5s	1.75s	2.0s	
ED [16]	43.8	40.9	38.7	36.8	34.6	33.9	32.5	31.6	36.6
RED [16]	45.3	42.1	39.6	37.5	35.8	34.4	33.2	32.1	37.5
TRN [54]	45.1	42.4	40.7	39.1	37.7	36.4	35.3	34.3	38.9
<b>OadTR</b>	50.2	49.3	48.1	46.8	45.3	43.9	42.4	41.1	<b>45.9</b>
<b>OadTR (Kinetics)</b>	59.8	58.5	56.6	54.6	52.6	50.5	48.6	46.8	<b>53.5</b>

(b) Results on the THUMOS14 dataset in terms of mAP (%).

Table 8: Action anticipation results of our OadTR compared to state-of-the-art methods using the same two-stream features.

Method	Dataset	Settings		
		Max-pool	Avg-pool	w/o encoder
OadTR	HDD	29.2	<b>29.8</b>	26.1
	TVSeries	85.1	<b>85.4</b>	80.8
	THUMOS14	57.9	<b>58.3</b>	53.5

Table 9: Comparison of different fusion methods and the necessity of the encoder feature encoding.

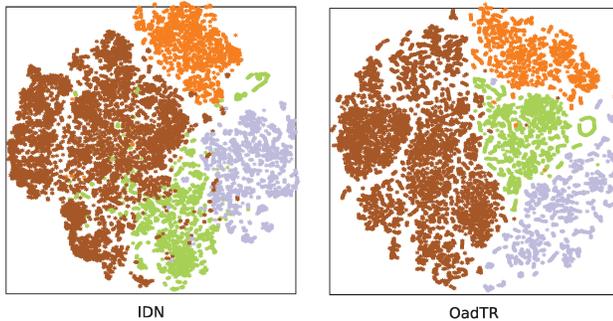


Figure 4: The t-SNE visualization of the classification embedding logits. Different colors correspond to different action categories from the THUMOS14 dataset. The mutual color correspondences include: **CliffDiving**, **HighJump**, **PoleVault** and **Shotput**. Better view in colored PDF.

#### 4.5. Qualitative evaluation

For better analysis, we visualize the classification results in Figure 4. Obviously, by visualizing all the test samples of the four action categories, we can observe that our OadTR has better separability compared with the current state-of-the-art IDN [15]. Further, we show OadTR’s multi-head attention visualization results in Figure 5, which indicates the importance of the multi-head design to learn more complex and comprehensive relations between different neighboring frame chunks.

## 5. Conclusion

In this paper, we have proposed a new online action detection framework built upon Transformers, termed OadTR.



Figure 5: Attention visualization maps. They indicate how much attention is paid to parts of the input streaming video.

In contrast to existing RNN based methods that process the sequence one by one recursively and hard to be optimized, we aim to design a direct end-to-end parallel network. OadTR can recognize current actions by encoding historical information and predicting future context simultaneously. Extensive experiments are conducted and verify the effectiveness of OadTR. Particularly, OadTR achieves higher training and inference speeds than current RNN based approaches, and obtains significantly better performance compared to the state-of-the-art methods. In the future, we will extend our OadTR model to more tasks such as action recognition, spatio-temporal action detection, *etc.*

## Acknowledgment

This work is supported by the National Natural Science Foundation of China under grant 61871435, Fundamental Research Funds for the Central Universities no.2019kfyXKJC024, and 111 Project on Computational Intelligence and Intelligent Control under Grant B18024.

## References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [2] Yueran Bai, Yingying Wang, Yunhai Tong, Yang Yang, Qiyue Liu, and Junhui Liu. Boundary content graph neural network for temporal action proposal generation. In *ECCV*, pages 121–137. Springer, 2020. 2
- [3] Andreja Bubic, D Yves Von Cramon, and Ricarda I Schubotz. Prediction, cognition and the brain. *Frontiers in human neuroscience*, 4:25, 2010. 4
- [4] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015. 5
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229. Springer, 2020. 2
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. 5
- [7] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018. 2
- [8] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 2
- [9] Andy Clark. Whatever next? predictive brains, situated agents, and the future of cognitive science. *Behavioral and brain sciences*, 36(3):181–204, 2013. 2
- [10] Roeland De Geest, Efstratios Gavves, Amir Ghodrati, Zhenyang Li, Cees Snoek, and Tinne Tuytelaars. Online action detection. In *ECCV*, pages 269–284. Springer, 2016. 2, 5, 6
- [11] Roeland De Geest and Tinne Tuytelaars. Modeling temporal structure with lstm for online action detection. In *WACV*, pages 1549–1557. IEEE, 2018. 1, 2
- [12] 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, 2
- [13] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, pages 2625–2634, 2015. 1, 5
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 1, 2
- [15] Hyunjun Eun, Jinyoung Moon, Jongyoul Park, Chanho Jung, and Changick Kim. Learning to discriminate information for online action detection. In *CVPR*, pages 809–818, 2020. 1, 2, 5, 6, 8
- [16] Jiyang Gao, Zhenheng Yang, and Ram Nevatia. Red: Reinforced encoder-decoder networks for action anticipation. In *BMVC*, 2017. 1, 2, 5, 8
- [17] Mingfei Gao, Mingze Xu, Larry S Davis, Richard Socher, and Caiming Xiong. Startnet: Online detection of action start in untrimmed videos. In *ICCV*, pages 5542–5551, 2019. 2
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 4, 5
- [19] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [21] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 5
- [22] Yu-Gang Jiang, Jingen Liu, A Roshan Zamir, George Toderici, Ivan Laptev, Mubarak Shah, and Rahul Sukthankar. Thumos challenge: Action recognition with a large number of classes, 2014. 2, 5
- [23] Shigeki Karita, Nanxin Chen, Tomoki Hayashi, Takaaki Hori, Hirofumi Inaguma, Ziyang Jiang, Masao Someki, Nelson Enrique Yalta Soplín, Ryuichi Yamamoto, Xiaofei Wang, et al. A comparative study on transformer vs rnn in speech applications. In *ASRU*, pages 449–456. IEEE, 2019. 1
- [24] Jinkyu Kim, Teruhisa Misu, Yi-Ting Chen, Ashish Tawari, and John Canny. Grounding human-to-vehicle advice for self-driving vehicles. In *CVPR*, pages 10591–10599, 2019. 1
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [26] Naihan Li, Shujie Liu, Yanqing Liu, Sheng Zhao, and Ming Liu. Neural speech synthesis with transformer network. In *AAAI*, volume 33, pages 6706–6713, 2019. 1
- [27] Xin Li, Tianwei Lin, Xiao Liu, Wangmeng Zuo, Chao Li, Xiang Long, Dongliang He, Fu Li, Shilei Wen, and Chuang Gan. Deep concept-wise temporal convolutional networks for action localization. In *ACMMM*, pages 4004–4012, 2020. 2
- [28] Yanghao Li, Cuiling Lan, Junliang Xing, Wenjun Zeng, Chunfeng Yuan, and Jiaying Liu. Online human action detection using joint classification-regression recurrent neural networks. In *ECCV*, pages 203–220. Springer, 2016. 1
- [29] Chuming Lin, Jian Li, Yabiao Wang, Ying Tai, Donghao Luo, Zhipeng Cui, Chengjie Wang, Jilin Li, Feiyue Huang, and Rongrong Ji. Fast learning of temporal action proposal via dense boundary generator. In *AAAI*, volume 34, pages 11499–11506, 2020. 2
- [30] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, pages 3889–3898, 2019. 2

- [31] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In *ACMMM*, pages 988–996, 2017. [2](#)
- [32] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *ECCV*, pages 3–19, 2018. [2](#)
- [33] Zachary C Lipton, John Berkowitz, and Charles Elkan. A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*, 2015. [1](#)
- [34] Ruijin Liu, Zejian Yuan, Tie Liu, and Zhiliang Xiong. End-to-end lane shape prediction with transformers. In *WACV*, pages 3694–3702, 2021. [2](#)
- [35] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, pages 21–37. Springer, 2016. [2](#)
- [36] Zhouyong Liu, Shun Luo, Wubin Li, Jingben Lu, Yufan Wu, Chunguo Li, and Luxi Yang. Convtransformer: A convolutional transformer network for video frame synthesis. *arXiv preprint arXiv:2011.10185*, 2020. [2](#)
- [37] Guansong Pang, Cheng Yan, Chunhua Shen, Anton van den Hengel, and Xiao Bai. Self-trained deep ordinal regression for end-to-end video anomaly detection. In *CVPR*, pages 12173–12182, 2020. [1](#)
- [38] Hyunjong Park, Jongyoun Noh, and Bumsub Ham. Learning memory-guided normality for anomaly detection. In *CVPR*, pages 14372–14381, 2020. [1](#)
- [39] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *ICML*, pages 1310–1318. PMLR, 2013. [1](#)
- [40] Zhiwu Qing, Haisheng Su, Weihao Gan, Dongliang Wang, Wei Wu, Xiang Wang, Yu Qiao, Junjie Yan, Changxin Gao, and Nong Sang. Temporal context aggregation network for temporal action proposal refinement. In *CVPR*, pages 485–494, 2021. [2](#)
- [41] Vasilii Ramanishka, Yi-Ting Chen, Teruhisa Misu, and Kate Saenko. Toward driving scene understanding: A dataset for learning driver behavior and causal reasoning. In *CVPR*, pages 7699–7707, 2018. [2](#), [5](#), [6](#)
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence (TPAMI)*, 39(6):1137–1149, 2016. [2](#)
- [43] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-deconvolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, pages 5734–5743, 2017. [5](#)
- [44] Zheng Shou, Junting Pan, Jonathan Chan, Kazuyuki Miyazawa, Hassan Mansour, Anthony Vetro, Xavier Giroi Nieto, and Shih-Fu Chang. Online action detection in untrimmed, streaming videos-modeling and evaluation. In *ECCV*, volume 1, page 5, 2018. [2](#)
- [45] Tianmin Shu, Dan Xie, Brandon Rothrock, Sinisa Todorovic, and Song Chun Zhu. Joint inference of groups, events and human roles in aerial videos. In *CVPR*, pages 4576–4584, 2015. [1](#)
- [46] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *NIPS*, 27:568–576, 2014. [5](#)
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [5](#)
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. [1](#), [2](#), [4](#), [6](#), [7](#)
- [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*, pages 20–36. Springer, 2016. [3](#), [5](#)
- [50] Xiang Wang, Changxin Gao, Shiwei Zhang, and Nong Sang. Multi-level temporal pyramid network for action detection. In *PRCV*, pages 41–54. Springer, 2020. [2](#)
- [51] Xiang Wang, Shiwei Zhang, Zhiwu Qing, Yuanjie Shao, Changxin Gao, and Nong Sang. Self-supervised learning for semi-supervised temporal action proposal. In *CVPR*, pages 1905–1914, 2021. [2](#)
- [52] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Masayoshi Tomizuka, Kurt Keutzer, and Peter Vajda. Visual transformers: Token-based image representation and processing for computer vision. *arXiv preprint arXiv:2006.03677*, 2020. [2](#)
- [53] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *ICCV*, pages 5783–5792, 2017. [2](#)
- [54] Mingze Xu, Mingfei Gao, Yi-Ting Chen, Larry S Davis, and David J Crandall. Temporal recurrent networks for online action detection. In *ICCV*, pages 5532–5541, 2019. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [55] Mengmeng Xu, Chen Zhao, David S Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *CVPR*, pages 10156–10165, 2020. [2](#)
- [56] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision (IJCV)*, 126(2-4):375–389, 2018. [5](#)
- [57] Manzil Zaheer, Guru Guruganesh, Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *arXiv preprint arXiv:2007.14062*, 2020. [1](#), [2](#)
- [58] Minghang Zheng, Peng Gao, Xiaogang Wang, Hongsheng Li, and Hao Dong. End-to-end object detection with adaptive clustering transformer. *arXiv preprint arXiv:2011.09315*, 2020. [1](#), [2](#)
- [59] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*, 2021. [2](#)
- [60] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond

efficient transformer for long sequence time-series forecasting. In *AAAI*, 2021. [7](#)

- [61] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [2](#)