Proposal-Free Temporal Action Detection via Global Segmentation Mask Learning

Sauradip Nag^{1,2}, Xiatian Zhu^{1,3}, Yi-Zhe Song^{1,2}, and Tao Xiang^{1,2}

¹ CVSSP, University of Surrey, UK

² iFlyTek-Surrey Joint Research Centre on Artificial Intelligence, UK
 ³ Surrey Institute for People-Centred Artificial Intelligence, University of Surrey, UK

 ${\tt s.nag, xiatian.zhu, y.song, t.xiang} @ {\tt surrey.ac.uk} \\$

Abstract. Existing temporal action detection (TAD) methods rely on generating an overwhelmingly large number of proposals per video. This leads to complex model designs due to proposal generation and/or perproposal action instance evaluation and the resultant high computational cost. In this work, for the first time, we propose a proposal-free Temporal Action detection model via Global Segmentation mask (TAGS). Our core idea is to learn a *global segmentation mask* of each action instance jointly at the full video length. The TAGS model differs significantly from the conventional proposal-based methods by focusing on global temporal representation learning to directly detect local start and end points of action instances without proposals. Further, by modeling TAD holistically rather than locally at the individual proposal level, TAGS needs a much simpler model architecture with lower computational cost. Extensive experiments show that despite its simpler design, TAGS outperforms existing TAD methods, achieving new state-of-the-art performance on two benchmarks. Importantly, it is $\sim 20 \times$ faster to train and $\sim 1.6 \times$ more efficient for inference. Our PyTorch implementation of TAGS is available at https://github.com/sauradip/TAGS.

1 Introduction

Temporal action detection (TAD) aims to identify the temporal interval (*i.e.*, the start and end points) and the class label of all action instances in an untrimmed video [16,5]. All existing TAD methods rely on *proposal generation* by either regressing predefined anchor boxes [42,8,15,22] (Fig. 1(a)) or directly predicting the start and end times of proposals [18,4,19,46,27,44,45] (Fig. 1(b)). Centered around proposals, existing TAD methods essentially take a local view of the video and focus on each individual proposal for action instance temporal refinement and classification. Such an approach thus suffers from several fundamental limitations: (1) An excessive (sometimes exhaustive) number of proposals are usually required for good performance. For example, BMN [18] generates ~ 5000 proposals per video by exhaustively pairing start and end points predicted. Generating and evaluating such a large number of proposals means high computational costs



Fig. 1: All existing TAD methods, no matter whether (a) anchor-based or (b) anchor-free, all need to generate action proposals. Instead, (c) our *global segmentation mask model* (TAGS) is proposal-free.

for both training and inference. (2) Once the proposals are generated, the subsequent modeling is local to each individual proposal. Missing global context over the whole video can lead to sub-optimal detection.

In this work, for the first time, we address these limitations by proposing a proposal-free TAD model. Our model, termed TAGS, learns a *global segmenta*tion mask of action instances at the full video length (Fig. 1(c)). By modeling TAD globally rather than locally, TAGS not only removes the need for proposal generation, and the associated design and computational complexity, it is also more effective. Concretely, instead of predicting the start/end points of each action instance, TAGS learns to predict an action segmentation mask of an entire video. Such a mask represents the *global* temporal structure of all action instances in the video; TAGS is thus intrinsically global context-aware.

Taking a proposal-free approach to TAD, our TAGS has a simpler model architecture design than existing methods. Specifically, it takes each local snippet (*i.e.*, a short sequence of consecutive frames of a video) as a predictive unit. That is, taking as input a snippet feature representation for a given video, TAGS directly outputs the target action segmentation mask as well as class label concurrently. To facilitate global context modeling, we leverage self-attention [36] for capturing necessary video-level inter-snippet relationship. Once the mask is generated, simple foreground segment classification follows to produce the final TAD result. To facilitate global segmentation mask learning, we further introduce a novel boundary focused loss that pays more attention to temporal boundary regions, and leverage mask predictive redundancy and inter-branch consistency for prediction enhancement. During inference, once the masks and class labels are predicted, top-scoring segments with refined boundary can then be selected via non-maximal suppression (NMS) to produce the final TAD result.

We make the following **contributions**. (I) We present a novel proposal-free TAD model based on *global segmentation mask* (TAGS) learning. To the best of our knowledge, this is the first model that eliminates the need for proposal generation/evaluation. As a result, it has a much simpler model design with a lower computational cost than existing alternatives. (II) We improve TAD feature representation learning with global temporal context using self-attention leading to

context-aware TAD. (III) To enhance the learning of temporal boundary, we propose a novel boundary focused loss function, along with mask predictive redundancy and inter-branch consistency. (IV) Extensive experiments show that the proposed TAGS method yields new state-of-the-art performance on two TAD datasets (ActivityNet-v1.3 and THUMOS'14). Importantly, our method is also significantly more efficient in both training/inference. For instance, it is $20/1.6 \times$ faster than G-TAD [46] in training and inference respectively.

2 Related Works

Although all existing TAD methods use proposals, they differ in how the proposals are generated.

Anchor-based proposal learning methods These methods generate proposal based on a pre-determined set of anchors. Inspired by object detection in static images [30], R-C3D [42] proposes to use anchor boxes. It follows the structure of proposal generation and classification in design. With similar model design, TURN [15] aggregates local features to represent snippet-level features, which are then used for temporal boundary regression and classification. Later, GTAN [22] improves the proposal feature pooling procedure with a learnable Gaussian kernel for weighted averaging. PBR-Net [20] improves the detection performance using a pyramidal anchor based detection and fine-grained refinement using frame-level features. G-TAD [46] learns semantic and temporal context via graph convolutional networks for better proposal generation. MUSES [21] further improves the performance by handling intra-instance variations caused by shot change. VSGN [51] focuses on short-action detection in a cross-scale multi-level pyramidal architecture. Note that these anchor boxes are often exhaustively generated so are high in number.

Anchor-free proposal learning methods Instead of using pre-designed and fixed anchor boxes, these methods directly learn to predict temporal proposals (*i.e.*, start and end times/points) [53,19,18]. For example, SSN [53] decomposes an action instance into three stages (starting, course, and ending) and employs structured temporal pyramid pooling to generate proposals. BSN [19] predicts the start, end and actionness at each temporal location and generates proposals using locations with high start and end probabilities. Later, BMN [18] additionally generates a boundary-matching confidence map to improve proposal generation. BSN++ [34] further extends BMN with a complementary boundary generator to capture rich context. CSA [33] enriches the proposal temporal context via attention transfer. Recently, ContextLoc [57] further pushes the boundaries by adapting global context at proposal-level and handling the context-aware inter-proposal relations. While no pre-defined anchor boxes are required, these methods often have to exhaustively pair all possible locations predicted with high scores. So both anchor-based and anchor-free TAD methods have a large quantity of temporal proposals to evaluate. This results in complex model design, high computational cost and lack of global context modeling. Our TAGS is designed to address all these limitations by being proposal-free.



Fig. 2: Architecture of our proposal-free Temporal Action detection model via Global Segmentation mask (TAGS). Given an untrimmed video, TAGS first extracts a sequence of T snippet features with a pre-trained video encoder (e.q., I3D [7]), and conducts self-attentive learning at multiple temporal scales s to obtain snippet embedding with global context. Subsequently, with each snippet embedding, TAGS classifies different actions (output $\boldsymbol{P}^{s} \in \mathbb{R}^{(K+1) \times T^{s}}$ with K the action class number) and predicts full-video-long foreground mask (output $M^s \in \mathbb{R}^{T^s \times T^s}$) concurrently in a two-branch design. During training, TAGS minimizes the difference of class and mask predictions against the ground-truth. For more accurate localization, an efficient boundary refinement strategy is further introduced, along with mask predictive redundancy and classification-mask consistency regularization. During inference, TAGS selects top scoring snippets from the classification output P, and then thresholds the corresponding foreground masks in M at each scale and then aggregates them to yield action instance candidates. Finally, softNMS is applied to remove redundant candidates.

Self-attention Our snippet representation is learned based on self-attention, which has been firstly introduced in Transformers for natural language processing tasks [36]. In computer vision, non-local neural networks [41] apply the core self-attention block from transformers for context modeling and feature learning. State-of-the-art performance has been achieved in classification [13], self-supervised learning [9], semantic segmentation [50,54], object detection [6,48,56], few-shot action recognition [28,55], and object tracking [10] by using such an attention model. Several recent works [35,40,29,26,25,27,24] also use Transformers for TAD. They focus on either temporal proposal generation [35] or refinement [29]. In this paper, we demonstrate the effectiveness of self-attention in a novel proposal-free TAD architecture.

3 Proposal-Free Global Segmentation Mask

Our global segmentation mask (TAGS) model takes as input an untrimmed video V with a variable number of frames. Video frames are pre-processed by a feature

encoder (e.g., a Kinetics pre-trained I3D network [7]) into a sequence of localized snippets following the standard practice [18]. To train the model, we collect a set of labeled video training set $\mathcal{D}^{train} = \{V_i, \Psi_i\}$. Each video V_i is labeled with temporal segmentation $\Psi_i = \{(\psi_j, \xi_j, y_j)\}_{j=1}^{M_i}$ where ψ_j/ξ_j denote the start/end time, y_j is the action category, and M_i is the action instance number.

Architecture As depicted in Fig. 2, a TAGS model has two key components: (1) a self-attentive snippet embedding module that learns feature representations with global temporal context (Sec. 3.1), and (2) a temporal action detection head with two branches for per-snippet multi-class action classification and binary-class global segmentation mask inference, respectively (Sec. 3.2).

3.1 Self-attentive multi-scale snippet embedding

Given a varying length untrimmed video V, following the standard practice [46,18] we first sample T equidistantly distributed temporal snippets (points) over the entire length and use a Kinetics pre-trained video encoder (*e.g.*, a two-stream model [38]) 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 $F = [X_r; X_o] \in \mathbb{R}^{2d \times T}$. 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 [36] to learn the global context. Formally, we set the Q/K/V of a Transformer encoder as the features F/F/F. To model finer action details efficiently, we consider multiple temporal scales in a hierarchy. We start with the finest temporal resolution (*e.g.*, sampling T = 800 snippets), which is progressively reduced via temporal pooling $P(\theta)$ with kernel size k, stride s and padding p. For efficiency, we first apply temporal pooling: $\hat{Q}^s = P(Q; \theta_Q), \hat{K}^s = P(K; \theta_K)$ and $\hat{V}^s = P(V; \theta_V)$ with the scale $s \in \{1, 2, 4\}$. The self-attention then follows as:

$$A_{i}^{s} = F + softmax(\frac{FW_{\hat{Q}^{s}}(FW_{\hat{K}^{s}})^{\top}}{\sqrt{d}})(FW_{\hat{V}^{s}}),$$
(1)

where $W_{\hat{Q}^s}, W_{\hat{K}^s}, W_{\hat{V}^s}$ are learnable parameters. In multi-head attention (MA) design, for each scale s we combine a set of n_h independent heads A_i to form a richer learning process. The snippet embedding E at scale s is obtained as:

$$E^{s} = MLP(\underbrace{[A_{1}^{s} \cdots A_{n_{h}}^{s}]}_{MA}) \in \mathbb{R}^{T^{s} \times C}.$$
(2)

The Multi-Layer Perceptron (MLP) block has one fully-connected layer with residual skip connection. Layer norm is applied before both the MA and MLP block. We use $n_h = 4$ heads by default.

3.2 Parallel action classification and global segmentation masking

Our TAD head consists of two parallel branches: one for multi-class action classification and the other for binary-class global segmentation mask inference.

Multi-class action classification Given the *t*-th snippet $E^s(t) \in \mathbb{R}^c$ (i.e, the *t*-th column of E^s), our classification branch predicts the probability $p_t \in \mathbb{R}^{(K+1)\times 1}$ that it belongs to one of *K* target action classes or background. This is realized by a 1-D convolution layer H_c followed by a softmax normalization. Since a video has been encoded into T^s temporal snippets, the output of the classification branch can be expressed in column-wise as:

$$\mathbf{P}^{s} := softmax(H_{c}(E^{s})) \in \mathbb{R}^{(K+1) \times T^{s}}.$$
(3)

Global segmentation mask inference In parallel to the classification branch, this branch aims to predict a global segmentation mask for each action instance of a video. Each global mask is action instance specific and class agnostic. For a training video, all temporal snippets of a single action instance are assigned with the same 1D global mask $\in \mathbb{R}^{T \times 1}$ for model optimization (refer to Fig. 3(a)). For each snippet $E^s(t)$, it outputs a mask prediction $\boldsymbol{m}_t = [q_1, \cdots, q_T] \in \mathbb{R}^{T^s \times 1}$ with the k-th element $q_k \in [0, 1]$ indicating the foreground probability of k-th snippet conditioned on t-th snippet. This process is implemented by a stack of three 1-D conv layers as:

$$\boldsymbol{M}^{s} := sigmoid(H_{b}(E^{s})) \in \mathbb{R}^{T^{s} \times T^{s}}, \tag{4}$$

where the t-th column of M is the segmentation mask prediction at the tth snippet. With the proposed mask signal as learning supervision, our TAGS model can facilitate context-aware representation learning, which brings clear benefit on TAD accuracy (see Table 4).

Remarks: Actionness [18,53] is a popular localization method which predicts a single mask in shape of $\in \mathbb{R}^{T \times 1}$. There are several key differences between actionness and TAGS: (1) Our per-snippet mask model TAGS focuses on a single action instance per snippet per mask so that all the foreground parts of a mask are intrinsically related; In contrast, actionness does not. (2) TAGS breaks the single multi-instance 1D actionness problem into a multiple 1D single-instance mask problem (refer to Fig. 3(a)). This takes a divide-and-conquer strategy. By explicitly segmenting foreground instances at different temporal positions, TAGS converts the regression based actionness problem into a position aware classification task. Each mask, associated with a specific time t, focuses on a single action instance. On the other hand, one action instance would be predicted by multiple successive masks. This predictive redundancy, simply removable by NMS, provides rich opportunities for accurate detection. (3) Whilst learning a 2D actionness map, BMN [18] relies on predicting 1D probability sequences which are highly noisy causing many false alarms. Further, its confidence evaluation cannot model the relations between candidates whilst our TAGS can (Eq. (7)). Lastly, our experiments in Table 8 validate the superiority of TAGS over actionness learning.

3.3 Model Training

Ground-truth labels. To train TAGS, the ground-truth needs to be arranged into the designed format. Concretely, given a training video with temporal in-



Fig. 3: Example of label assignment and model inference (see text for details).

tervals and class labels (Fig. 3(a)), we label all the snippets (orange or blue squares) 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, its global mask is defined as the video-length binary mask of that action instance. Each mask is action instance specific. All snippets of a specific action instance share the same mask. For instance, all orange snippets (Fig. 3(a)) are assigned with a *T*-length mask (eg. m_{24} to m_{38}) with one in the interval of [q24, q38].

Learning objectives. The classification branch is trained by a combination of a cross-entropy based focal loss and a class-balanced logistic regression loss [12]. For a training snippet, we denote y the ground-truth class label, p the classification output, and r the per-class regression output obtained by applying sigmoid on top of H_c in Eq. (3) (discarded at inference time). The loss of the classification branch is then written as:

$$L_c = \lambda_1 (1 - \boldsymbol{p}(y))^{\gamma} \log(\boldsymbol{p}_y) + (1 - \lambda_1) \big(\log(\boldsymbol{r}_y) - \frac{\alpha}{|\mathcal{N}|} \sum_{k \in \mathcal{N}} (\log(1 - \boldsymbol{r}(k))) \big), \quad (5)$$

where $\gamma = 2$ is a focal degree parameter, $\alpha = 10$ is a class-balancing weight, and \mathcal{N} specifies a set of hard negative classes at size of K/10 where K is the toTAD action class number. We set the loss trade-off parameter $\lambda_1 = 0.4$.

For training the segmentation mask branch, we combine a novel boundary IOU (bIOU) loss and the dice loss in [23] to model two types of structured consistency respectively: mask boundary consistency and inter-mask consistency. Inspired by the boundary IOU metric [11], bIOU is designed particularly to penalize incorrect temporal boundary prediction w.r.t. the ground-truth segmentation mask. Formally, for a snippet location, we denote $\boldsymbol{m} \in \mathbb{R}^{T \times 1}$ the predicted segmentation mask, and $\boldsymbol{g} \in \mathbb{R}^{T \times 1}$ the ground-truth mask. The overall segmentation mask loss is formulated as:

$$L_{m} = 1 - \left(\frac{\bigcap(m,g)}{\bigcup(m,g)} + \frac{1}{\bigcap(m,g) + \epsilon} \frac{\|\mathbf{m} - \mathbf{g}\|_{2}}{c}\right) + \lambda_{2} \left(1 - \frac{\mathbf{m}^{\top} \mathbf{g}}{\sum_{t=1}^{T} \left(\mathbf{m}(t)^{2} + \mathbf{g}(t)^{2}\right)}\right),$$
(6)



Fig. 4: An example of (a) ground-truth labels and (b) prediction along with an illustration of exploring (c) mask predictive redundancy (Eq. (7)) (d) classification-mask consistency ((Eq. (9)).

where $\cap(m,g) = \Phi(\mathbf{m}) \cap \Phi(\mathbf{g})$ and $\cup(m,g) = \Phi(\mathbf{m}) \cup \Phi(\mathbf{g})$, $\Phi(\cdot)$ represents a kernel of size k (7 in our default setting, see more analysis in Suppl.) used as a differentiable morphological erosion operation [31] on a mask and c specifies the ground-truth mask length. In case of no boundary overlap between the predicted and ground-truth masks, we use the normalized L_2 loss. The constant $\epsilon = e^{-8}$ is introduced for numerical stability. We set the weight $\lambda_2 = 0.4$.

Mask predictive redundancy Although the mask loss Eq. (6) above treats the global mask as a 2D binary mask prediction problem, it cannot always regulate the behaviour of individual 1D mask within an action instance. Specifically, for a predicted mask m_t at time t, thresholding it at a specific threshold $\theta_j \in \Theta$ can result in binarized segments of foreground and background: $\pi[j] = \{(q_s^i, q_e^i, z_i)\}_{i=1}^L$ where q_s^i and q_e^i denotes the start and end of *i*-th segment, and $z_i \in \{0, 1\}$ indicates background or foreground. For a mask corresponding to an action snippet, ideally at least one of these $\{\pi[j]\}$ should be closer to the ground truth. To explore this redundancy, we define a prediction scoring criterion with the outer-inner-contrast [32] as follows:

$$\mathbb{R}(\pi[j]) = \frac{1}{L} \sum_{i=1}^{L} \left(\underbrace{\frac{1}{l_i} \sum_{\substack{r=q_s^i \\ inside}}^{q_e^i} u_i(r)}_{\text{inside}} - \underbrace{\frac{1}{\left[\delta l_i\right] + \left\lceil \delta l_i\right]} \left(\sum_{\substack{r=q_s^i - \left\lceil \delta l_i\right\rceil \\ outside}}^{q_s^i - 1} u_i(r) + \sum_{\substack{r=q_e^i + 1 \\ r=q_e^i + 1}}^{q_e^i + \left\lceil \delta l_i\right\rceil} u_i(r) \right)}_{\text{outside}} \right)$$

$$where \quad u_i(r) = \begin{cases} m_t[r], & \text{if } z_i = 1 \ (i.e., \text{ foreground}) \\ 1 - m_t[r], & \text{otherwise} \end{cases}$$

$$(7)$$

 $l_i = q_e^i - q_s^i + 1$ is the temporal length of *i*-th segment, δ is a weight hyperparameter which is set to 0.25. We obtain the best prediction with the maximal score as $j^* = argmax(\mathbb{R}(\pi[j]))$ (see Fig. 4(c)). Higher $\mathbb{R}(\pi[j^*])$ means a better prediction. To encourage this best prediction, we design a prediction promotion loss function:

$$L_{pp} = \left(1 - \mathbb{R}(\pi[j^*])\right)^{\beta} \|m_t - g_t\|_2,$$
(8)

where we set $\beta = 2$ for penalizing lower-quality prediction stronger. We average this loss across all snippets of each action instance per training video.

Classification-mask consistency In TAGS, there is structural consistency in terms of foreground between class and mask labels by design (Fig. 4(a)). To leverage this consistency, we formulate a feature consistency loss as:

$$L_{fc} = 1 - \operatorname{cosine}\left(\hat{F}_{clf}, \hat{F}_{mask}\right),\tag{9}$$

where $\hat{F}_{clf} = topk(argmax((P_{bin} * E_p)[: K, :]))$ is the features obtained from the top scoring foreground snippets obtained from the thresholded classification output $P_{bin} := \eta(P - \theta_c)$ with θ_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: $\hat{F}_{mask} =$ $topk(\sigma(1DPool(E_m * M_{bin})))$ 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, * is element-wise multiplication, $\eta(.)$ is the binarization function, and σ is sigmoid activation. Our intuition is that the foreground features should be closer and consistent after the classification and masking process (refer to Fig. 4(d)).

Overall objective The overall objective loss function for training TAGS is defined as: $L = L_c + L_m + L_{pp} + L_{fc}$. This loss is calculated for each temporal scale s and finally aggregated over all the scales.

3.4 Model Inference

Our model inference is similar as existing TAD methods [18,46]. Given a test video, at each temporal scale s the action instance predictions are first generated separately based on the classification P^s and mask M^s predictions and combined for the following post-processing. Starting with the top scoring snippets from P (Fig 3(b)), we obtain their segmentation mask predictions (Fig 3(c)) by thresholding the corresponding columns of M (Fig 3(d)). To generate sufficient candidates, we apply multiple thresholds $\Theta = \{\theta_i\}$ to yield action candidates with varying lengths and confidences. For each candidate, we compute its confidence score sc_{final} by multiplying the classification score (obtained from the corresponding top-scoring snippet in P) and the segmentation mask score (*i.e.*, the mean predicted foreground segment in M). Finally, we apply SoftNMS [3] on top scoring candidates to obtain the final predictions.

4 Experiments

Datasets We conduct extensive experiments on two popular TAD benchmarks. (1) *ActivityNet-v1.3* [5] has 19,994 videos from 200 action classes. We follow the

Type	Model 	Bkb	THUMOS14					ActivityNet-v1.3				
1900			0.3	0.4	0.5	0.6	0.7	Avg.	0.5	0.75	0.95	Avg.
	R-C3D [42]	C3D	44.8	35.6	28.9	-	-	-	26.8	-	-	-
	TAD [8]	I3D	53.2	48.5	42.8	33.8	20.8	39.8	38.2	18.3	1.3	20.2
Anchor	GTAN [22]	P3D	57.8	47.2	38.8	-	-	-	52.6	34.1	8.9	34.3
	PBR-Net [20]	I3D	58.5	54.6	51.3	41.8	29.5	-	53.9	34.9	8.9	35.0
	MUSES [21]	I3D	68.9	64.0	56.9	46.3	31.0	53.4	50.0	34.9	6.5	34.0
	VSGN [51]	I3D	66.7	60.4	52.4	41.0	30.4	50.1	52.3	36.0	8.3	35.0
	BMN [18]	TS	56.0	47.4	38.8	29.7	20.5	38.5	50.1	34.8	8.3	33.9
	DBG [17]	TS	57.8	49.4	42.8	33.8	21.7	41.1	-	-	-	-
	[G-TAD [46]	TS	54.5	47.6	40.2	30.8	23.4	39.3	50.4	34.6	9.0	34.1
	BU-TAL [52]	I3D	53.9	50.7	45.4	38.0	28.5	43.3	43.5	33.9	9.2	30.1
	BSN++[34]	TS	59.9	49.5	41.3	31.9	22.8	-	51.2	35.7	8.3	34.8
Actn	GTAD+CSA[33]	TS	58.4	52.8	44.0	33.6	24.2	42.6	51.8	36.8	8.7	35.7
	BC-GNN [2]	TS	57.1	49.1	40.4	31.2	23.1	40.2	50.6	34.8	9.4	34.3
	TCANet [29]	TS	60.6	53.2	44.6	36.8	26.7	-	52.2	36.7	6.8	35.5
	ContextLoc [57]	I3D	68.3	63.8	54.3	41.8	26.2	-	56.0	35.2	3.5	34.2
	[RTD-Net [35]	I3D	68.3	62.3	51.9	38.8	23.7	-	47.2	30.7	8.6	30.8
Mixed	A2Net [47]	I3D	58.6	54.1	45.5	32.5	17.2	41.6	43.6	28.7	3.7	27.8
	GTAD+PGCN[49]	I3D	66.4	60.4	51.6	37.6	22.9	47.8	-	-	-	-
PF	TAGS (Ours) TAGS (Ours)	I3D TS	68.6 61.4	63.8 52.9	$\begin{array}{c} 57.0\\ 46.5\end{array}$	$\begin{array}{c} 46.3\\ 38.1 \end{array}$	$\begin{array}{c} 31.8\\ 27.0 \end{array}$	52.8 44.0	$\begin{array}{c} 56.3 \\ 53.7 \end{array}$	36.8 36.1	9.6 9.5	$\begin{array}{c} 36.5\\ 35.9\end{array}$

Table 1: Performance comparison with state-of-the-art methods on THUMOS14 and ActivityNet-v1.3. The results are measured by mAP at different IoU thresholds, and average mAP in [0.3 : 0.1 : 0.7] on THUMOS14 and [0.5 : 0.05 : 0.95] on ActivityNet-v1.3. Actn = Actioness; PF = Proposal Free; Bkb = Backbone.

standard setting to split all videos into training, validation and testing subsets in ratio of 2:1:1. (2) *THUMOS14* [16] has 200 validation videos and 213 testing videos from 20 categories with labeled temporal boundary and action class.

Implementation details We use two pre-extracted encoders for feature extraction, for fair comparisons with previous methods. One is a fine-tuned twostream model [18], with downsampling ratio 16 and stride 2. Each video's feature sequence F is rescaled to T = 800/1024 snippets for AcitivtyNet/THUMOS using linear interpolation. The other is Kinetics pre-trained I3D model [7] with a downsampling ratio of 5. Our model is trained for 15 epochs using Adam with learning rate of $10^{-4}/10^{-5}$ for AcitivityNet/THUMOS respectively. The batch size is set to 50 for ActivityNet and 25 for THUMOS. For classification-mask consistency, the threshold θ_m/θ_p is set to 0.5/0.3 and in top-k to 40. In testing, we set the threshold set for mask $\Theta = \{0.1 \sim 0.9\}$ with step 0.05. We use the same set of threshold Θ for mask predictive redundancy during training.

4.1 Main Results

Results on ActivityNet From Table 1, we can make the following observations: (1) TAGS with I3D feature achieves the best result in average mAP. Despite the fact that our model is much simpler in architecture design compared to the existing alternatives. This validates our assumption that with proper global context modeling, explicit proposal generation is not only redundant but also less effective. (2) When using the relatively weaker two-stream (TS) features, our model remains competitive and even surpasses I3D based BU-TAL [52], A2Net [47] and the very recent ContextLoc [57] and MUSES [21] by a significant margin. TAGS also surpasses a proposal refinement and strong G-TAD based approach CSA [33] on avg mAP. (3) Compared to RTD-Net which employs an architecture similar to object detection Transformers, our TAGS is significantly superior. This validates our model formulation in exploiting the Transformer for TAD.

Results on THUMOS14 Similar conclusions can be drawn in general on THUMOS from Table 1. When using TS features, TAGS achieves again the best results, beating strong competitors like TCANet [39], CSA [33] by a clear margin. There are some noticeable differences: (1) We find that I3D is now much more effective than two-stream (TS), *e.g.*, 8.8% gain in average mAP over TS with TAGS, compared with 0.6% on ActivityNet. This is mostly likely caused by the distinctive characteristics of the two datasets in terms of action instance duration and video length. (2) Our method achieves the second best result with marginal edge behind MUSES [21]. This is partly due to that MUSES benefits from additionally tackling the scene-changes. (3) Our model achieves the best results in stricter IOU metrics (*e.g.*, IOU@0.5/0.6/0.7) consistently using both TS and I3D features, verifying the effectiveness of solving mask redundancy.

Computational cost comparison One of the key motivations to design a proposal-free TAD model is to reduce the model training and inference cost. For comparative evaluation, we evaluate TAGS against two representative and recent TAD methods (BMN [18] and G-TAD [46]) using their released codes. All the methods are tested on the same machine with one Nvidia 2080 Ti GPU. We measure the convergence time in training and average inference time per video in testing. The two-stream video features are used. It can be seen in Table 2 that our TAGS is drastically faster, *e.g.*, $20/25 \times$ for training and clearly quicker – $1.6/1.8 \times$ for testing in comparison to G-TAD/BMN, respectively. We also notice that our TAGS needs less epochs to converge. Table 3 also shows that our TAGS has the smallest FLOPs and the least parameter number.

4.2 Ablation study and further analysis

Transformers vs. CNNs We compare our multi-scale Transformer with CNN for snippet embedding. We consider two CNN designs: (1) a 1D CNN with 3 dilation rates (1, 3, 5) each with 2 layers, and (2) a multi-scale MS-TCN [14], and (3) a standard single-scale Transformer [36]. Table 4 shows that the Transformers are clearly superior to both 1D-CNN and a relatively stronger MS-TCN. This suggests that our global segmentation mask learning is more compatible with

ana tost cost.							
Model	Epoch	Train	Test				
BMN G-TAD	13 11	6.45 hr 4.91 hr	0.21 sec 0.19 sec				
TAGS	9	0.26 hr	0.12 sec				

Table 2: Analysis of model trainingTableand test cost.# an

Tat	ole 3	3:	Anal	ysis	ot	model	paran	neters
# E	and	FΙ	LOPs					

Model	Params (in M)	FLOPs (in G)
BMN	5.0	91.2
GTAD	9.5	97.2
TAGS	6.2	17.8

self-attention models due to stronger contextual learning capability. Besides, multi-scale learning with Transformer gives 0.4% gain in avg mAP validating the importance of larger snippets. As shown in Table 5, the gain almost saturates from 200 snippets, and finer scale only increases the computational cost.

Table 4: Ablation of Transformer vs. CNN on ActivityNet.

Table 5:	Ablation	on snippet	t embedding d	le-
sign and	multiple	temporal	scales.	

Network	m	4P	Seelo	Spippot	Params	Infer	mAP	
	0.5	\mathbf{Avg}		Supper	$ ^{(in M)} $	(in sec)	0.5	Avg
1D CNN	46.8	26.4	{1}	100	2.9	0.09	55.8	36.1
MS-TCN	53.1	33.8	(1.0)	100 000		0.10	1	
Transformer	55.8	36.1	$\{1,2\}$	100,200	6.2	0.12	56.3	36.5
MS-Transformer	56.3	36.5	$\{1,2,4\}$	100,200,400	9.8	0.16	56.5	36.4

Proposal-based vs. proposal-free We compare our proposal-free TAGS with conventional proposal-based TAD methods BMN [2] (anchor-free) and R-C3D [42] (anchor-based) via false positive analysis [1]. We sort the predictions by the scores and take the top-scoring predictions per video. Two major errors of TAD are considered: (1) *Localization error*, which is defined as when a proposal/mask is predicted as foreground, has a minimum tIoU of 0.1 but does not meet the tIoU threshold. (2) *Background error*, which happens when a proposal/mask is predicted as foreground but its tIoU with ground truth instance is smaller than 0.1. In this test, we use ActivityNet. We observe in Fig. 5 that TAGS has the most true positive samples at every amount of predictions. The proportion of localization error with TAGS is also notably smaller, which is the most critical metric for improving average mAP [1]. This explains the gain of TAGS over BMN and R-C3D.

Direction of improvement analysis Two subtasks are involved in TAD – temporal localization and action classification, each of which would affect the final performance. Given the two-branch design in TAGS, the performance effect of one subtask can be individually examined by simply assigning ground-truth to the other subtask's output at test time. From Table 7, the following observations can be made: (1) There is still a big scope for improvement on both subtasks.



Fig. 5: False positive profile of TAGS, BMN and R-C3D on ActivityNet. We use top up-to 10G predictions per video, where G is the number of ground truth action instances.

(2) Regarding the benefit from the improvement from the other subtask, the classification subtask seems to have the most to gain at mAP@0.5, whilst the localization task can benefit more on the average mAP metric. Overall, this analysis suggests that further improving the efficacy on the classification subtask would be more influential to the final model performance.

Analysis of components We can see in Table 6 that without the proposed segmentation mask branch, the model will degrade significantly, *e.g.*, a drop of 7.6% in average mAP. This is due to its fundamental capability of modeling the global temporal structure of action instances and hence yielding better action temporal intervals. Further, for TAGS we use a pre-trained UntrimmedNet (UNet) [37] as an external classifier instead of using the classification branch, resulted in a 2-stage method. This causes a performance drop of 4.7%, suggesting that both classification and mask branches are critical for model accuracy and efficiency.

Table 6: Analysis of TAGS's two branches on ActivityNet.

Table 7: Improvement	analysis	of TAGS
on ActivityNet.		

Model	mAP	Model		mAP		
110 401	$\left 0.5 \left \mathrm{Avg} \right. \right $		$\left 0.5 \right $	\mathbf{Avg}		
TAGS(Full)	56.3 36.5	TAGS (full)	56.3	36.5		
w/o Mask Branch w/o Class Branch + UNe	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	+ Ground-trut + Ground-trut	$\begin{array}{c c} \text{th class} & 61.0 \\ \text{th mask} & 69.2 \\ \end{array} \\ \begin{array}{c} 48 \end{array}$	$\begin{array}{c} 3.8 (\uparrow 7.3\%) \\ 8.5 (\uparrow 12.0\%) \end{array}$		

Global mask design We compare our global mask with previous 1D actionness mask [18,43]. We integrate actionness with TAGS by reformulating the mask branch to output 1D actionness. From the results in Table 8, we observe a significant performance drop of 11.5% in mAP@0.5 IOU. One reason is that the number of action candidates generated by actionness is drastically limited,

THOS ON MENNING GAUGED.						
Mask Design	$\left \begin{array}{c} \mathrm{m} \mathrm{m} \mathrm{m} \mathrm{m} \mathrm{m} \mathrm{m} \mathrm{m} m$	$\frac{AP}{Avg}$	Avg masks / video			
Actionness	44.8	27.1	30			
Our Global Mask	56.3	36.5	250			





leading to poor recall. Additionally, we visualize the cosine similarity scores of all snippet feature pairs on a random ActivityNet val video. As shown in Fig. 6, our single-instance mask (global mask) design learns more discriminating feature representation with larger separation between background and action, as compared to multi-instance actionness design. This validates the efficacy of our design in terms of jointly learning multiple per-snippet masks each with focus on a single action instance.

5 Limitations

In general, short foreground and background segments with the duration similar as or less than the snippet length would challenge snippet-based TAD methods. For instance, given short background between two foreground instances, our TAGS might wrongly predict it as part of the foreground. Besides, given a snippet with mixed background and foreground, TAGS tends to make a background prediction. In such cases, the ground-truth annotation involves uncertainty which however is less noted and investigated thus far.

6 Conclusion

In this work, we have presented the first proposal-free TAD model by Global Segmentation Mask (TAGS) learning. Instead of generating via predefined anchors, or predicting many start-end pairs (*i.e.*, temporal proposals), our model is designed to estimate the full-video-length segmentation mask of action instances directly. As a result, the TAD model design has been significantly simplified with more efficient training and inference. With our TAGS learning, we further show that learning global temporal context is beneficial for TAD. Extensive experiments validated that the proposed TAGS yields new state-of-the-art performance on two TAD benchmarks, and with clear efficiency advantages on both model training and inference.

15

References

- Alwassel, H., Heilbron, F.C., Escorcia, V., Ghanem, B.: Diagnosing error in temporal action detectors. In: ECCV. pp. 256–272 (2018) 12
- Bai, Y., Wang, Y., Tong, Y., Yang, Y., Liu, Q., Liu, J.: Boundary content graph neural network for temporal action proposal generation. In: ECCV. pp. 121–137. Springer (2020) 10, 12
- 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) 9
- 4. Buch, S., Escorcia, V., Shen, C., Ghanem, B., Carlos Niebles, J.: Sst: Single-stream temporal action proposals. In: CVPR (2017) 1
- 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) 1, 9
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: Endto-end object detection with transformers. In: ECCV (2020) 4
- 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) 4, 5, 10
- Chao, Y.W., Vijayanarasimhan, S., Seybold, B., Ross, D.A., Deng, J., Sukthankar, R.: Rethinking the faster r-cnn architecture for temporal action localization. In: CVPR (2018) 1, 10
- Chen, M., Radford, A., Child, R., Wu, J., Jun, H., Dhariwal, P., Luan, D., Sutskever, I.: Generative pretraining from pixels. In: ICML (2020) 4
- Chen, X., Yan, B., Zhu, J., Wang, D., Yang, X., Lu, H.: Transformer tracking. In: CVPR (2021) 4
- Cheng, B., Girshick, R., Dollár, P., Berg, A.C., Kirillov, A.: Boundary iou: Improving object-centric image segmentation evaluation. In: CVPR. pp. 15334–15342 (2021) 7
- Dong, Q., Zhu, X., Gong, S.: Single-label multi-class image classification by deep logistic regression. In: AAAI. vol. 33, pp. 3486–3493 (2019) 7
- 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) 4
- 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) 1, 3
- 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) 1, 10
- Lin, C., Li, J., Wang, Y., Tai, Y., Luo, D., Cui, Z., Wang, C., Li, J., Huang, F., Ji, R.: Fast learning of temporal action proposal via dense boundary generator. In: AAAI. vol. 34, pp. 11499–11506 (2020) 10
- 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) 1, 3, 5, 6, 9, 10, 11, 13

- 16 Nag et al.
- 19. Lin, T., Zhao, X., Su, H., Wang, C., Yang, M.: BSN: Boundary sensitive network for temporal action proposal generation. In: ECCV (2018) 1, 3
- Liu, Q., Wang, Z.: Progressive boundary refinement network for temporal action detection. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 11612–11619 (2020) 3, 10
- Liu, X., Hu, Y., Bai, S., Ding, F., Bai, X., Torr, P.H.: Multi-shot temporal event localization: a benchmark. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12596–12606 (2021) 3, 10, 11
- 22. Long, F., Yao, T., Qiu, Z., Tian, X., Luo, J., Mei, T.: Gaussian temporal awareness networks for action localization. In: CVPR (2019) 1, 3, 10
- 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) 7
- Nag, S., Zhu, X., Song, Y.Z., Xiang, T.: Temporal action localization with global segmentation mask transformers (2021) 4
- Nag, S., Zhu, X., Song, Y.z., Xiang, T.: Semi-supervised temporal action detection with proposal-free masking. In: ECCV (2022) 4
- Nag, S., Zhu, X., Song, Y.z., Xiang, T.: Zero-shot temporal action detection via vision-language prompting. In: ECCV (2022) 4
- 27. Nag, S., Zhu, X., Xiang, T.: Few-shot temporal action localization with query adaptive transformer. arXiv preprint arXiv:2110.10552 (2021) 1, 4
- Perrett, T., Masullo, A., Burghardt, T., Mirmehdi, M., Damen, D.: Temporalrelational crosstransformers for few-shot action recognition. In: CVPR (2021) 4
- Qing, Z., Su, H., Gan, W., Wang, D., Wu, W., Wang, X., Qiao, Y., Yan, J., Gao, C., Sang, N.: Temporal context aggregation network for temporal action proposal refinement. In: CVPR. pp. 485–494 (2021) 4, 10
- 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) 3
- Riba, E., Mishkin, D., Ponsa, D., Rublee, E., Bradski, G.: Kornia: an open source differentiable computer vision library for pytorch. In: WACV. pp. 3674–3683 (2020)
 8
- Shou, Z., Gao, H., Zhang, L., Miyazawa, K., Chang, S.F.: Autoloc: Weaklysupervised temporal action localization in untrimmed videos. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 154–171 (2018) 8
- Sridhar, D., Quader, N., Muralidharan, S., Li, Y., Dai, P., Lu, J.: Class semanticsbased attention for action detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13739–13748 (2021) 3, 10, 11
- 34. 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) 3, 10
- Tan, J., Tang, J., Wang, L., Wu, G.: Relaxed transformer decoders for direct action proposal generation. In: ICCV (2021) 4, 10
- 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) 2, 4, 5, 11
- Wang, L., Xiong, Y., Lin, D., Van Gool, L.: Untrimmednets for weakly supervised action recognition and detection. In: CVPR. pp. 4325–4334 (2017) 13
- Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Van Gool, L.: Temporal segment networks: Towards good practices for deep action recognition. In: European conference on computer vision. pp. 20–36. Springer (2016) 5

- Wang, L., Yang, H., Wu, W., Yao, H., Huang, H.: Temporal action proposal generation with transformers. arXiv preprint arXiv:2105.12043 (2021) 11
- 40. Wang, X., Zhang, S., Qing, Z., Shao, Y., Zuo, Z., Gao, C., Sang, N.: Oadtr: Online action detection with transformers. In: ICCV (2021) 4
- Wang, X., Girshick, R., Gupta, A., He, K.: Non-local neural networks. In: CVPR (2018) 4
- Xu, H., Das, A., Saenko, K.: R-c3d: Region convolutional 3d network for temporal activity detection. In: ICCV (2017) 1, 3, 10, 12
- Xu, M., Pérez-Rúa, J.M., Escorcia, V., Martinez, B., Zhu, X., Zhang, L., Ghanem, B., Xiang, T.: Boundary-sensitive pre-training for temporal localization in videos. arXiv (2020) 13
- Xu, M., Pérez-Rúa, J.M., Escorcia, V., Martinez, B., Zhu, X., Zhang, L., Ghanem, B., Xiang, T.: Boundary-sensitive pre-training for temporal localization in videos. In: ICCV. pp. 7220–7230 (2021) 1
- Xu, M., Perez-Rua, J.M., Zhu, X., Ghanem, B., Martinez, B.: Low-fidelity end-toend video encoder pre-training for temporal action localization. In: NeurIPS (2021) 1
- 46. Xu, M., Zhao, C., Rojas, D.S., Thabet, A., Ghanem, B.: G-tad: Sub-graph localization for temporal action detection. In: CVPR (2020) 1, 3, 5, 9, 10, 11
- Yang, L., Peng, H., Zhang, D., Fu, J., Han, J.: Revisiting anchor mechanisms for temporal action localization. IEEE Transactions on Image Processing 29, 8535– 8548 (2020) 10, 11
- 48. Yin, M., Yao, Z., Cao, Y., Li, X., Zhang, Z., Lin, S., Hu, H.: Disentangled non-local neural networks. In: ECCV (2020) 4
- Zeng, R., Huang, W., Tan, M., Rong, Y., Zhao, P., Huang, J., Gan, C.: Graph convolutional networks for temporal action localization. In: ICCV (2019) 10
- Zhang, L., Xu, D., Arnab, A., Torr, P.H.: Dynamic graph message passing networks. In: CVPR (2020) 4
- 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) 3, 10
- Zhao, P., Xie, L., Ju, C., Zhang, Y., Wang, Y., Tian, Q.: Bottom-up temporal action localization with mutual regularization. In: ECCV. pp. 539–555. Springer (2020) 10, 11
- 53. Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., Lin, D.: Temporal action detection with structured segment networks. In: ICCV (2017) 3, 6
- 54. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: CVPR (2021) 4
- Zhu, X., Toisoul, A., Perez-Rua, J.M., Zhang, L., Martinez, B., Xiang, T.: Fewshot action recognition with prototype-centered attentive learning. arXiv preprint arXiv:2101.08085 (2021) 4
- 56. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint (2020) 4
- Zhu, Z., Tang, W., Wang, L., Zheng, N., Hua, G.: Enriching local and global contexts for temporal action localization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13516–13525 (2021) 3, 10, 11