

Semi-Supervised Temporal Action Detection with Proposal-Free Masking

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

Abstract. Existing temporal action detection (TAD) methods rely on a large number of training data with segment-level annotations. Collecting and annotating such a training set is thus highly expensive and unscalable. Semi-supervised TAD (SS-TAD) alleviates this problem by leveraging unlabeled videos freely available at scale. However, SS-TAD is also a much more challenging problem than supervised TAD, and consequently much under-studied. Prior SS-TAD methods directly combine an existing proposal-based TAD method and a SSL method. Due to their *sequential* localization (*e.g.*, proposal generation) and classification design, they are prone to proposal error propagation. To overcome this limitation, in this work we propose a novel *Semi-supervised Temporal action detection model based on Proposal-free Temporal mask* (SPOT) with a *parallel* localization (mask generation) and classification architecture. Such a novel design effectively eliminates the dependence between localization and classification by cutting off the route for error propagation in-between. We further introduce an interaction mechanism between classification and localization for prediction refinement, and a new pretext task for self-supervised model pre-training. Extensive experiments on two standard benchmarks show that our SPOT outperforms state-of-the-art alternatives, often by a large margin. The PyTorch implementation of SPOT is available at <https://github.com/sauradip/SPOT>

1 Introduction

Temporal action detection (TAD) aims to predict the temporal duration (*i.e.*, the start and end points) and the class label of each action instances in an untrimmed video [22,8]. Most state-of-the-art TAD methods [65,66,7,54,70,38,37] rely on training datasets containing a large number of videos (*e.g.*, hundreds) with exhaustive segment-level annotations. Obtaining such annotations is tedious and costly. This has severely limited the usability of existing TAD methods in low data setting [40,39].

Semi-supervised learning (SSL) offers a solution to the annotation cost problem by exploiting a large amount of unlabeled data along with limited labeled data [51,47]. This has led to an emerging research interest [23,57] in semi-supervised TAD (SS-TAD). Existing methods adopt an intuitive strategy of

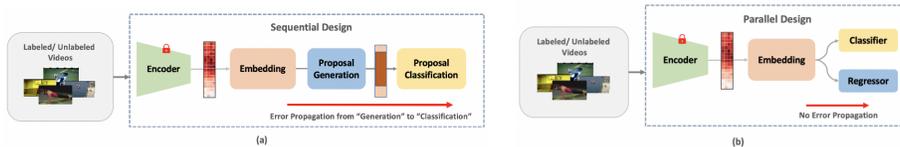


Fig. 1: Illustration of main design difference between existing and our SS-TAD methods. (a) Existing SS-TAD methods suffers from the intrinsic localization error propagation problem, due to the *sequential* localization (*e.g.*, proposal generation) and classification design. (b) We solve this problem by designing a Proposal Free Temporal Mask (SPOT) learning model with a *parallel* localization and classification architecture.

combining an existing TAD models, dominated by proposal-based methods and a SSL method. However, this strategy is intrinsically sub-optimal and prone to an error propagation problem. As illustrated in Fig. 1(a), this is because existing TAD models adopt a *sequential* localization (*e.g.*, proposal generation) and classification design. When extended to SSL setting, the localization errors, inevitable when trained with unlabeled data, can be easily propagated to the classification module leading to accumulated errors in class prediction.

To overcome the above limitation, in this work we propose a novel Semi-supervised PropOsal-free Temporal Masking (SPOT) model with a *parallel* localization (mask generation) and classification architecture (see Fig. 1(b)). Concretely, SPOT consists of a classification stream and a mask based localization stream, established in parallel on a shared feature embedding module. This architecture design has no sequential dependence between localization and classification as in conventional TAD models, therefore eliminating the localization error propagation problem. We further introduce a boundary refinement algorithm and a novel pretext task for self-supervised model pre-training. We integrate the SPOT with pseudo labeling for SS-TAD with new classification and mask loss functions formulated specifically for our parallel design. Moreover, by virtue of being proposal free, our model is $30\times/2\times$ faster in training/inference than existing alternatives.

Contributions. **(1)** To solve the localization error propagation problem suffered by existing SS-TAD methods, we propose a *Proposal-Free Temporal Masking* (SPOT) model with a new parallel classification and localization architecture⁴. **(2)** We further design a novel pretext task for model pre-training and a boundary refinement algorithm. **(3)** Extensive experiments on two standard benchmarks (ActivityNet-V1.3 and THUMOS14) show that the proposed SPOT outperforms significantly alternative state-of-the-art SSL methods.

⁴ Note, instead of contributing a novel generic SSL algorithm, we propose a new TAD architecture designed particularly for facilitating the usage of prior SSL methods (*e.g.*, pseudo labeling) in the sense of minimizing localization error propagation.

2 Related Works

Temporal Action Detection While all existing TAD methods use action proposals, they differ in how the proposals are produced.

Anchor-based proposal learning methods generate proposal with a pre-determined set of anchors. Inspired by object detection in static images [44], R-C3D [62] proposes to use anchor boxes. It follows the structure of proposal generation and classification in design. With a similar model design, TURN [17] aggregates local features to represent snippet-level features, which are then used for temporal boundary regression and classification. Later, GTAN [30] improves the proposal feature pooling procedure with a learnable Gaussian kernel for weighted averaging. G-TAD [66] learns semantic and temporal context via graph convolutional networks for better proposal generation. Recently, VSGN [68] improves short-action localization with 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* directly learn to predict temporal proposals (*i.e.*, start and end times) [70,28,27]. For example, SSN [70] decomposes an action instance into three stages (starting, course, and ending) and employs structured temporal pyramid pooling to generate proposals. BSN [28] predicts the start, end and actionness at each temporal location and generates proposals with high start and end probabilities. Later, BMN [27] additionally generates a boundary-matching confidence map to improve proposal generation. BSN++ [49] further extends BMN with a complementary boundary generator to capture rich context. CSA [48] enriches the proposal temporal context via attention transfer. 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. Critically, both groups of TAD models in essence adopt a *sequential* localization (mask generation) and classification architecture. This will cause the localization error propagation problem for SS-TAD. Our SPOT is designed to address this limitation by removing the dependence between localization and classification thus cutting off the path for error propagation.

Semi-supervised learning (SSL) [71,10] has been widely adopted in computer vision for image classification [5,47,13], object detection [50,69], semantic segmentation [41,21], and pose estimation [15,35]. Two dominant learning paradigms in SSL are pseudo-labeling [47,67,26,24] and consistency regularization [51,25,60,36]. Key to pseudo-labeling is to reliably estimate the labels of unlabeled data which in turn are used to further train the model. Instead, consistency regularization enforces the output of a model to be consistent at the presence of variations in the input space and/or the model space. The variations can be implemented by adding noises, perturbations or forming multiple variants of the same data sample or the model. In this work, we focus on designing a TAD model particularly suitable for SSL, while following the pseudo-labelling paradigm for exploiting unlabeled data for training.

Semi-supervised Temporal Action Detection (SS-TAD). SSL has only been studied very recently in the context of TAD. Existing SS-TAD works

[23,57,46] naively combine existing semi-supervised learning and TAD methods. They are thus particularly prone to the aforementioned localization error propagation problem when trained with unlabeled data. We solve this problem for the first time by introducing a novel proposal-free temporal mask learning model.

Self-supervised learning aims to learn generic feature representations from a large amount of unlabeled data [11,19,18]. It is typically designed to provide a pre-trained model for a downstream task to further fine-tune with a specific labeled training data. We have seen a recent surge of self-supervised learning studies with a focus on both object recognition in images [59,11,19,18,33,53,12] and action classification in videos [1,4,31,34,58]. The most related works to our self-supervised learning based pre-training are very recently introduced in [65,64]. They aim to improve the video encoder for the fully supervised TAD problem. In contrast, we focus on pre-training a superior TAD head in the context of semi-supervised learning.

3 Proposal-Free Temporal Mask Learning

Approach overview In semi-supervised temporal action detection (SS-TAD), we have access to a small set of N_l labeled videos $\mathcal{D}^l = \{V_i, \Psi_i\}_{i=1}^{N_l}$ and a large set of N_u unlabeled videos $\mathcal{D}^u = \{U_i\}_{i=1}^{N_u}$. Each labeled video’s annotation $\Psi_i = \{(\psi_j, \xi_j, y_j)\}_{j=1}^{M_i}$ contains the start time ψ_j , the end time ξ_j , and the class label $y_j \in \mathcal{Y}$ for each of M_i action instances. We denote the label space as $\mathcal{Y} = [1, \dots, K + 1]$ with K action and one background classes. For more effective SS-TAD, we propose a *Proposal-Free Temporal Mask* (SPOT) learning method (see Fig. 2). It has two components: video snippet embedding (Sec. 3.1), and TAD head (Sec. 3.2). The latter is our core contribution.

3.1 Video Snippet Embedding

Given a varying length untrimmed video V , following the standard practice [66,27] we first sample T equidistantly distributed temporal snippets (points) over the entire length and use a fine-tuned two-stream video encoder similar to [27] 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 (16 in this work) consecutive frames. While F contains local spatio-temporal information, it lacks a global context critical for TAD. We hence leverage the self-attention mechanism [52] to learn the global context. Formally, we set the input $\{query, key, value\}$ of a multi-head Transformer encoder $\mathcal{T}(\cdot)$ as the features $\{F, F, F\}$ (Fig. 2(a)). Positional encoding is not applied as it is found to be detrimental (see Appendix B in Supplementary) The final snippet embedding is then obtained as $E = \mathcal{T}(F) \in \mathbb{R}^{C \times T}$ with C being the embedding dimension.

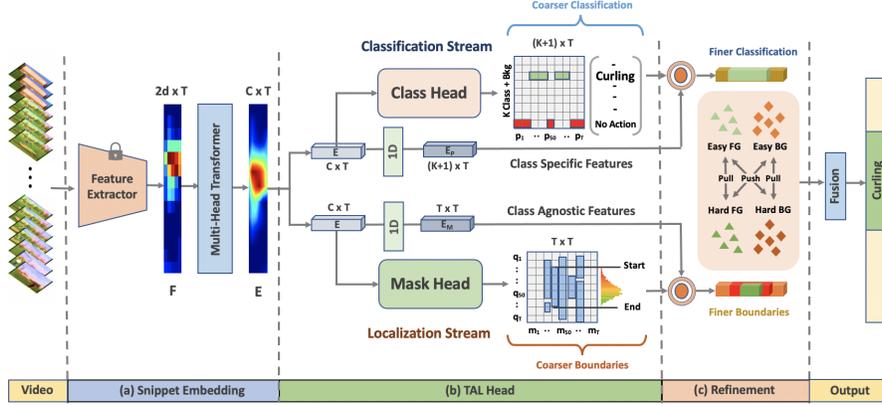


Fig. 2: **Overview of the proposed *Semi-supervised Temporal action detection model based on Proposal-free Temporal mask (SPOT)*.** Given an untrimmed video V , (a) we first extract a sequence of T snippet features with a pre-trained video encoder and conduct self-attention learning to obtain the snippet embedding E with a global context. (b) For each snippet embedding, we then predict a classification score P with the classification stream and a foreground mask M with the mask stream in parallel, (c) both of which are further used for boundary refinement. It is based on mining hard&easy foreground (FG) and background (BG) snippets. For SS-TAD, we alternately predict and leverage pseudo class and mask labels of unlabeled training videos, along with labeled videos.

3.2 Temporal Action Detection Head

To realize a proposal-free design, we introduce a temporal mask learning based TAD head. It consists of two parallel streams (Fig. 2(b)): one for snippet classification and the other for temporal mask inference. This design breaks the sequential dependence between localization and classification which causes unwanted error propagation in the existing TAD models.

Snippet classification stream Given the t -th snippet $E(t) \in \mathbb{R}^c$ (*i.e.*, the t -th column of E), our classification branch predicts a probability distribution $\mathbf{p}_t \in \mathbb{R}^{(K+1) \times 1}$ over \mathcal{Y} . This is realized by a 1-D convolution layer H_c followed by a softmax normalization. For a video with T snippets, the output of the classification branch can be expressed as:

$$\mathbf{P} := \text{softmax}(H_c(E)) \in \mathbb{R}^{(K+1) \times T}. \quad (1)$$

Temporal mask stream In parallel to the classification stream, this stream predicts temporal masks of action instances across the whole temporal span of the video. Given the t -th snippet $E(t)$, it outputs a mask vector $\mathbf{m}_t = [q_1, \dots, q_T] \in \mathbb{R}^{T \times 1}$ with each element $q_i \in [0, 1]$ ($i \in [1, T]$) indicating the

foreground probability of the i -th snippet (see Fig. 2(b) for illustration). This is implemented by a stack of three 1-D convolution layers H_m as:

$$\mathbf{M} := \text{sigmoid}(H_m(E)) \in \mathbb{R}^{T \times T}, \quad (2)$$

where the t -th column of \mathbf{M} is the temporal mask prediction by the t -th snippet. With the proposed mask signal as the model output supervision, proposals are no longer required for facilitating SS-TAD learning.

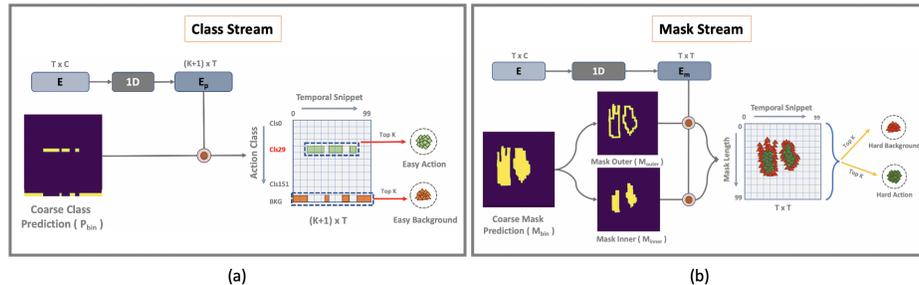


Fig. 3: Snippet selection for inter-stream interaction.

Boundary Refinement TAD methods typically struggle at estimating accurately the boundary between foreground and background segments. This problem is more pronounced for a SS-TAD model by learning from unlabeled data. To mitigate this problem, we design an inter-stream interaction mechanism for TAD refinement at the end of the TAD head (Fig. 2(c)). More specifically, we focus on ambiguous snippets in the transition between foreground and background (*i.e.*, temporal boundary). They are considered as *hard snippets*, in contrast to those easy ones located inside mask or background intervals far away from temporal borders with more trustworthy interference. We detect hard snippets by examining the structure of temporal mask \mathbf{M} . First, we threshold the M to obtain a binary mask as $M_{bin} := \eta(M - \theta_m)$, where $\eta(\cdot)$ is the Heavyside step function and θ_m is the threshold. As shown in Fig. 3(b), we consider the snippets spanned by the eroded mask boundary as *hard background* whereas those by the non-boundary mask as *hard foreground*. We use a differentiable morphological erosion [45] to obtain the eroded mask M_{outer} . We denote the complement of M_{outer} as M_{inner} . They are formally defined as:

$$M_{outer} = \mathbb{E}(M_{bin}, e), \quad M_{inner} = M_{bin} - M_{outer},$$

where $\mathbb{E}(\cdot)$ is the differentiable erosion operation and e is the erosion kernel size. We represent the top k scoring hard snippets by multiplying the downsampled embedding E_m (obtained after applying 1-D convolution on the embedding E). They are obtained as:

$$X^{fg} = \text{topk}(M_{inner} * E_m), \quad X^{bg} = \text{topk}(M_{outer} * E_m),$$

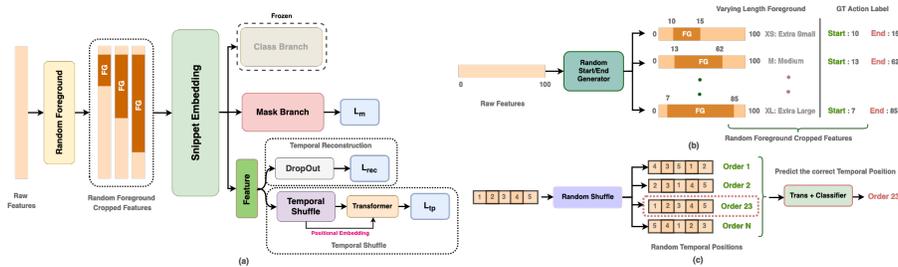


Fig. 4: Illustration of our self-supervised pre-training.

As the snippets from the same instance often predict different length 1-D masks, applying a binarization process on top generates masks of arbitrary shapes, as illustrated in Fig. 3(a). We calculate the boundary features for all of these mask clusters and select the top snippets from all of the clusters jointly. For inter-stream interaction, we further use high-confidence foreground and background snippets from the classification stream. We thus consider them as *easy foreground* and *easy background* due to their *easy to predict* property. As seen in Fig. 3(a), we similarly select top scoring foreground and background snippets from the thresholded classification output $P_{bin} := \eta(P - \theta_c)$ as:

$$Y^{fg} = \text{topk}(\text{argmax}((P_{bin} * E_p)[: K, :])), \quad Y^{bg} = \text{topk}((P_{bin} * E_p)[K + 1, :]).$$

where E_p is obtained by passing the embedding E into a 1-D conv layer for matching the dimension of P . We adopt infoNCE loss [20] for TAD refinement:

$$L_{ref} = \text{tri}(x^{fg}, y^{fg}, y^{bg}) + \text{tri}(y^{bg}, x^{bg}, y^{fg}), \quad (3)$$

where $x^* \in X^*$ and $y^* \in Y^*$ and $\text{tri}(\cdot, \cdot, \cdot)$ defines the foreground and background triplet training samples. In this way, we maximize the correlation between easy and hard snippets of the same category (foreground or background), refining the inter-stream feature representation for better perceiving temporal boundaries.

3.3 Model Training

To better exploit unlabeled data, we formulate a two-staged training pipeline, including self-supervised pre-training and semi-supervised fine-tuning.

Stage I: Self-supervised pre-training We introduce a pretext task based on a novel notion of *random foreground* specially designed for TAD model pre-training. Given a video feature sequence $F \in \mathbb{R}^{T \times d}$, we sample a random token segment (s, e) of varying proportions as foreground and the remaining tokens as background (Fig. 4). Cropping out foreground at feature level has shown to be useful in learning discriminative representation [42]. Motivated by this, we zero out the background snippet features whilst keeping for the pseudo foreground.

This shares a similar spirit with masked region modeling [14,3] with a different objective of detecting the location of masked segment.

With such a masked feature sequence, our pretext task aims to predict jointly (1) the temporal mask with the start s and end e (Fig. 4(b)), (2) the temporal position of each snippet after temporal shuffling (Fig. 4(c)), and (3) the reconstruction of snippet feature (Fig. 4(a)). We treat each of T temporal positions as a specific positional class, and apply a small Transformer with learnable positional embedding on the shuffled snippet sequences. All zeroed snippet features will become *non-zero* after the transformer’s processing, preventing the model from learning a trivial solution. We use the cross-entropy loss L_{tp} for temporal position prediction. This supervision aims to learn the intrinsic temporal structure information in videos. The feature reconstruction loss L_{rec} is measured by L_2 normalized mean squared error between the backbone feature F and the embedding E from the transformer. The motivation is to preserve the video encoder’s discriminative information in learning global context. By dropping random snippets, the Transformer is forced to aggregate and utilize information from the context to predict the dropped snippets. As a result, the model can learn temporal semantic relations and discriminative features useful for TAD. The pretext task loss for pre-training is formulated as:

$$L_{pre} = L_m + \lambda_1 L_{rec} + \lambda_2 L_{tp}, \quad (4)$$

where L_m is the mask learning loss as defined in Eq. (11) and λ_1 and λ_2 are two hyper-parameters set as 0.8 and 0.4 respectively. We use both labeled and unlabeled data for SPOT pre-training without using any ground-truth labels.

Stage II: Semi-supervised fine-tuning We implement temporal mask semi-supervised learning following the pseudo label paradigm [47]. Concretely, we alternate between predicting and applying pseudo labels, starting by using the labeled samples alone.

Pseudo label prediction Given an unlabeled video embedding E , the pseudo class label is obtained by:

$$\hat{y} = \max(\text{softmax}(H_c(E)/\tau_c)), \quad (5)$$

where the sharpening operator $\tau_c = \tau - (\tau - 1)\hat{y}'_{max}$ with \hat{y}'_{max} the maximum probability over the K classes, and τ a hyper-parameter controlling the largest sharpening strength. Similarly, we obtain the pseudo mask label \hat{g} as:

$$\hat{g} = \text{sig}(H_m(E)/\tau_m), \quad (6)$$

where τ_m is the mask sharpening operator, and $\text{sig}()$ is the sigmoid function [29]. Then, we threshold the pseudo class and mask labels at θ_c/θ_m to be binary.

Loss functions For SS-TAD, we use both pseudo and ground-truth labels to minimize the objective loss function as formulated below. For the *classification* stream, we devise a class-balanced loss to tackle the intrinsic class imbalance

challenge with TAD. Inspired by [56,16], we adopt the sigmoid regression based binary cross-entropy loss function. Given a snippet $E(t)$ with class label y , the classification loss function is designed as:

$$L_{bce} = -\log(p_y) - \sum_{k \neq y, k \in \mathcal{Y}} \left(\log(1 - p_k) \right), \quad (7)$$

where p_y denotes the the logits from class stream P . Critically, this binary cross-entropy loss gives us the flexibility to regulate each class individually per training sample. This is useful because untrimmed videos often encompass a dominant proportion of background content, which would overwhelm the under-represented tail action classes \mathcal{Y}_t with small training data. To mitigate this problem, we further improve this loss by encouraging the activation of tail action classes. Concretely, given a *background* snippet, we still allow tail action classes to be activated under a certain degree ε (set as θ_c). This is realized by introducing a weighting mechanism as:

$$L_{wbce} = -\log(p_y) - \sum_{k \neq y, k \in \mathcal{Y}} \left(w_k \log(1 - p_k) \right) \quad (8)$$

$$\text{where } \omega_k = \begin{cases} 0 & \text{if } k \in \mathcal{Y}_t \text{ and } p(k) < \varepsilon \\ 1 & \text{otherwise} \end{cases} \quad (9)$$

Given a video with the background T_{bg} and foreground T_{fg} snippets, our final classification loss is expressed as:

$$L_c = \frac{1}{T} \left(\sum_{t \in T_{fg}} L_{bce}(t) + \sum_{t \in T_{bg}} L_{wbce}(t) \right). \quad (10)$$

For the *mask* stream, we exploit a weighted binary cross-entropy loss L_{mce} for balancing foreground and background classes, along with a binary dice loss L_{dice} [32]. For a snippet location t , we denote $\mathbf{m}(t) \in \mathbb{R}^{T \times 1}$ and $\mathbf{g}(t) \in \mathbb{R}^{T \times 1}$ as the predicted and ground-truth mask. The mask learning loss is designed as:

$$L_m = \beta_{fg} \sum_{t=1}^T \mathbf{g}(t) \log(\mathbf{m}(t)) + \beta_{bg} \sum_{t=1}^T (1 - \mathbf{g}(t)) \log(1 - \mathbf{m}(t)) \\ + \lambda_d \left(1 - \frac{\mathbf{m}^\top \mathbf{g}}{\sum_{t=1}^T (\mathbf{m}(t)^2 + \mathbf{g}(t)^2)} \right), \quad (11)$$

where β_{fg} and β_{bg} are the inverse proportion of foreground and background used for class balance, and λ_d is the dice loss weight which is empirically set as 0.6.

Overall learning objective The overall objective loss is designed as $L = L_c + L_m + L_{ref} + L_{rec}$ where L_{ref} is the refinement loss (Eq. (3)) and L_{rec} is the feature reconstruction loss as described in the pre-training stage. This loss is applied on both ground-truth and pseudo labels for fine-tuning SPOT. Note that the temporal ordering loss term L_{tp} is not used during fine-tuning as it gives performance drop (see Appendix B in Supplementary). A plausible cause is its incompatibility with the fine-tuning loss L_c .

Labels	Methods	ActivityNet				THUMOS					
		0.5	0.75	0.95	Avg	0.3	0.4	0.5	0.6	0.7	Avg
60%	BMN* [27]	47.6	31.7	7.5	31.5	50.8	45.9	34.8	23.7	16.3	34.3
	Mean Teacher [51]+BMN	48.0	32.1	7.4	31.9	53.5	45.0	36.9	27.4	19.0	35.8
	FixMatch [47]+BMN	48.7	32.9	7.7	32.8	53.8	46.2	37.8	28.7	19.5	36.9
	SSP [23]	49.8	34.5	7.0	33.5	53.2	46.8	39.3	29.7	19.8	37.8
	SSTAP [57]	50.1	34.9	7.4	34.0	56.4	49.5	41.0	30.9	21.6	39.9
	SPOT (Ours)	52.8	35.0	8.1	35.2	58.9	50.1	42.3	33.5	22.9	41.5
10%	BMN* [27]	35.4	26.4	8.0	25.8	38.3	28.3	18.8	11.4	5.6	20.5
	Mean Teacher [51]+BMN	36.0	27.2	7.4	26.6	41.2	32.1	23.1	15.0	7.0	23.7
	FixMatch [47]+BMN	36.8	27.9	8.0	26.9	42.0	32.8	23.0	15.9	8.5	24.3
	SSP [23]	38.9	28.7	8.4	27.6	44.2	34.1	24.6	16.9	9.3	25.8
	SSTAP [57]	40.7	29.6	9.0	28.2	45.6	35.2	26.3	17.5	10.7	27.0
	SPOT (Ours)	49.9	31.1	8.3	32.1	49.4	40.4	31.5	22.9	12.4	31.3

Table 1: **SS-TAD results on the validation set of ActivityNet v1.3 and the test set of THUMOS14.** Note: *All methods except SPOT* use the UNet [54] trained with *100% class labels* for proposal classification; These methods hence benefit from *extra classification supervision* in comparison to SPOT. *: Using only labeled training set.

3.4 Model Inference

At test time, we generate action instance predictions for each test video by the classification \mathbf{P} and mask \mathbf{M} predictions. For \mathbf{P} , we only consider the snippets whose class probabilities are greater than θ_c and select top scoring snippets. For each such top scoring action snippet, we then obtain the temporal mask by thresholding the t_i -th column of \mathbf{M} using the localization threshold θ . To produce sufficient candidates, we use a set of thresholds $\Theta = \{\theta_i\}$. For each candidate, we compute a confidence score s by multiplying the classification and max-mask scores. SoftNMS [6] is finally applied to obtain top scoring results.

4 Experiments

Datasets We use two standard TAD datasets in our evaluation. **(1)** *ActivityNet v1.3* is a large-scale benchmark containing 19,994 untrimmed videos with 200 classes. We adopt the standard 2:1:1 training/validation/testing video data split. **(2)** *THUMOS14* provides 200 validation videos and 213 testing videos labeled with temporal annotations for action understanding. We train our model on the validation set and evaluate on the test set.

Implementation details For fair comparisons, we use both TSN [55] and I3D [9] features in our evaluations. For ActivityNet, we use fine-tuned TSN features for fair comparison with [57,27]. For THUMOS, we use TSN and I3D

features both pre-trained on Kinetics [61]. The temporal dimension T is fixed at 100/256 for ActivityNet/THUMOS respectively (see the suppl. for more analysis). For ActivityNet/THUMOS, we first pre-train SPOT (except the classification stream) on the training set including the unlabeled samples for 12 epochs and then we fine-tune SPOT for 15 epochs with a learning rate of $10^{-4}/10^{-5}$, a weight decay of $10^{-3}/10^{-5}$. For boundary refinement, we set top- $k = 40$ snippets, θ_c/θ_m is set as 0.3/0.7 and e is set as 7. In semi-supervised setting, the label sharpening operator τ is set as 1.1 and τ_m is set as 0.7. In testing, we set the threshold set for mask $\theta = \{0.1 \sim 0.9\}$ with a step 0.1. The SoftNMS is performed on ActivityNet/Thumos with a threshold of 0.6/0.4.

4.1 Comparative Results

Setting We introduce two SS-TAD settings with different label sizes. For each dataset, we randomly select 10% or 60% training videos as the labeled set and the remaining as the unlabeled set. Both labeled and unlabeled sets are accessible for SS-TAD model training.

Competitors We compared with the following methods. **(1)** Two state-of-the-art supervised TAD methods: BMN [27]+UNet [54] and GTAD [63]+UNet [54] **(2)** Two SSL+TAD methods: As SS-TAD is a new problem, we need to implement the competitors by extending existing SSL methods to TAD by ourselves. We select two top SSL methods (Mean Teacher [51] and FixMatch [47]), and a state-of-the-art TAD model based on a popular proposal generation method BMN [27] (using TSN features) and GTAD [63] (using I3D features). Both of these models use a common untrimmed video classification model UNet [54]. For FixMatch [47], we use temporal flip (*i.e.*, playing the video backwards) as strong augmentation and temporal feature shift as weak augmentation. For UNet, due to the lack of Caffe based training environment, we can only apply the official weights trained with 100% supervision **(3)** Two recent semi-supervised temporal proposal generation methods [23,57]. Note that the concurrent SS-TAD work [46] does not give reproducible experiment settings and prevents an exact comparison. Besides, [46] achieves only 12.27% avg mAP on ActivityNet v1.2 *vs.* 32.1% by SPOT on v1.3 (v1.2 has only half classes of v1.3 hence simpler), thus significantly inferior to the two compared methods [23,57].

Results The SS-TAD results are reported in Table 1 and Table 5. We make the following observations: **(1)** Without leveraging unlabeled video data, the state-of-the-art fully supervised TAD method BMN [27] (with UNet [54]) achieves the poorest result among all the methods. This clearly suggests the usefulness of unlabeled data and the consistent benefits of adopting semi-supervised learning for TAD – a largely ignored research topic. **(2)** Combining existing TAD methods (*e.g.*, BMN [27]) with previous SSL methods (*e.g.*, Mean Teacher [51] and FixMatch [47]) is indeed effective in improving model generalization. Similar observations can be made on SSP [23] and SS-TAD [57], our two main competitors. Despite such performance improvements, these methods still suffer from the proposal error propagation problem that would limit their ability in exploiting unlabeled videos. **(3)** By solving this problem with a new parallel design, our

Table 2: Analysis of localization error propagation on ActivityNet with 10% supervision. GT: Ground-Truth.

Metric	mAP	
	0.5	Avg
BMN [27] + MLP		
GT proposals	55.7	45.3
Pseudo proposals	32.4	23.6
SPOT		
GT masks	59.2	47.0
Pseudo masks	49.9	32.1

Table 3: Analysis of SPOT model pre-training on ActivityNet with 10% label supervision.

Pre-training loss			mAP	
L_m	L_{rec}	L_{tp}	0.5	Avg
Random initialization				
\times	\times	\times	46.2	30.5
With our pre-training				
\checkmark	\times	\times	47.6	31.4
\checkmark	\checkmark	\times	48.5	31.7
\checkmark	\times	\checkmark	47.9	31.5
\checkmark	\checkmark	\checkmark	49.9	32.1

SPOT achieves the new state of the art on both datasets. It is worth pointing out the larger margin achieved by SPOT over all the competitors at the lower supervision case (10% label). For instance, with 10% label, SPOT surpasses the second best SSTAP by 3.9%/4.3% on ActivityNet/THUMOS. A plausible reason is that the proposal error propagation will become more severe when the labeled set is smaller, causing more harm to existing proposal based TAD methods. This validates the overall efficacy and capability of our model formulation in leveraging unlabeled video data for SS-TAD. (4) From Table 5 we observe similar findings as that of BMN [27] using TSN features (Table 1). Our SPOT outperforms the existing baselines and SS-TAD models by almost similar margin, confirming that the superiority of our method is feature agnostic.

4.2 Further Analysis

Localization error propagation analysis To examine the effect of localization error propagation with previous TAD models, we design a proof-of-concept experiment by measuring the performance drop between ground-truth proposals and pseudo proposals. Due to the lack of access to the now obsolete training environment for UNet [54], we adopt a MLP classifier with BMN [27] as the baseline TAD model. For our SPOT, we contrast ground-truth and pseudo masks. This experiment is tested on ActivityNet with 10% supervision. Table 2 shows that the proposal based TAD baseline suffers almost double performance degradation from localization (*i.e.*, proposal) error due to its sequential localization and classification design. This verifies the advantage of SPOT’s parallel design.

Effectiveness of model pre-training We propose a two-staged training method characterized by model pre-training based on a novel pretext task (Sec. 4.2). We now examine how pre-training helps for model accuracy. Here, we take the 10% label supervision case on ActivityNet, with a comparison to the random initialization baseline. Concretely, we examine three loss terms: the mask prediction

Fig. 5: Effect of our model pre-training on a random ActivityNet video.

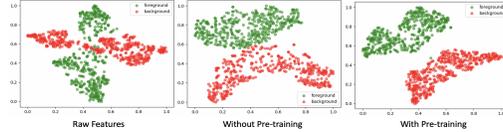


Table 4: Effect of pre-training on ActivityNet with 10% labels.

Method	mAP	
	0.5	Avg
BMN + [54]	35.4	25.8
BMN^\dagger + [54]	36.2	26.3

Table 5: **SS-TAD results** on **THUMOS14 test-set** using I3D features; All method except **SPOT** uses UNet [54]; *:Using only labeled set.

Labels	Method	mAP		
		0.3	0.5	0.7
60%	<i>GTAD*</i> [63]	50.9	35.4	16.5
	FixMatch [47]+GTAD	53.4	38.9	19.1
	SSP [23]	53.5	39.7	20.4
	SSTAP[57]+GTAD	55.9	41.6	22.0
	SPOT (Ours)	58.7	42.4	23.1
10%	<i>GTAD*</i> [63]	36.9	20.1	6.6
	FixMatch [47]+GTAD	42.1	23.8	9.7
	SSP [23]	43.1	25.5	9.6
	SSTAP [57]+GTAD	45.3	27.5	11.0
	SPOT (Ours)	49.1	31.7	12.6

Table 6: Performance of our SPOT model w/ and w/o unlabeled data on ActivityNet.

Labels	SSL Modules		mAP	
	Pre-train	L_c	0.5	Avg
10%	✓	✓	49.9	32.1
	✓	✗	45.3	29.8
	✗	✓	46.2	30.5
	✗	✗	44.5	28.3
60%	✓	✓	52.8	35.2
	✓	✗	51.7	34.4
	✗	✓	52.1	34.9
	✗	✗	51.2	34.0

loss (L_m), the feature reconstruction loss (L_{rec}), and the temporal position loss (L_{tp}). The results in Table 3 reveal that: (1) Compared to random initialization, our pre-training boosts the overall mAP by 1.6%. This gain is also supported by the more distinct separation of foreground and background features as shown in Fig. 5 This validates the importance of TAD model initialization and the efficacy of our pretext task. (2) All the loss terms are effective either in isolation or in the cohorts, suggesting they are complementary with good compatibility. Additionally, we evaluate the generic impact of our pre-training on BMN [27]. We observe from Table 4 that it again gives a good increase of 0.7% in Avg mAP, justifying its generic usefulness.

Effectiveness of using unlabeled data We evaluate the impact of using unlabeled data imposed by the pre-training and the loss term L_c (in Eq. (10)). Table 6 shows that without these components, the model will degrade in performance particularly in case of less labels, as expected. In particular, both SSL components (pretraining and L_c) are clearly effective in model performance. This verifies their designs. Besides, even without the SSL modules our model is better than SSTAP [57] in 10% label case and comparable in 60% label case, further indicating the advantage of our model design.

Table 7: Ablating SPOT objective loss terms on ActivityNet with 10% label. L_c : classification loss; L_{mce} : cross-entropy mask loss; L_{dice} : dice mask loss; L_{ref} : refinement loss; L_{rec} : feature reconstruction loss.

Loss	mAP	
	0.5	Avg
All	49.9	32.1
W/O L_c	45.3	29.8
W/O L_{mce}	47.8	31.5
W/O L_{dice}	47.0	31.2
W/O L_{rec}	45.9	30.2
W/O L_{ref}	46.8	30.7

Table 8: **Fully supervised TAD results on the validation set of ActivityNet.** All the compared methods use TSN features used in [57]

Method	mAP			
	0.5	0.75	0.95	Avg
BSN [28]	46.4	29.9	8.0	30.0
GTAD [66]	50.3	34.6	9.0	34.0
BC-GNN [2]	50.6	34.8	9.4	34.3
BMN [27]	50.0	34.8	8.3	33.8
BSN++ [49]	51.2	35.7	8.3	34.8
TCANet [43]	52.2	36.7	6.8	35.5
GTAD+CSA [48]	51.8	36.8	8.7	35.6
SPOT	53.9	35.9	9.4	35.8

Contributions of different losses We ablate the effect of SPOT’s training loss functions for classification and mask (Sec. 3.3). We test the benefit of the classification loss (L_c), the weighted binary cross-entropy mask loss (L_{mce}), the dice mask loss (L_{dice}), the refinement loss (L_{ref}), and the feature reconstruction loss (L_{rec}). When testing L_c , we replace it with the standard cross-entropy loss. Table 7 shows that each loss term is beneficial for improving model accuracy, verifying their individual and collective efficacy.

Comparisons to fully supervised TAD methods Besides SS-TAD, our SPOT can be also applied for fully supervised TAD with the pseudo-labels replaced by ground-truth labels while keeping the remaining the same. This test is conducted on ActivityNet. Table 8 shows that when trained with fully-labeled data, our SPOT can also outperform state-of-the-art TAD methods in the overall result albeit by a smaller margin as compared to the semi-supervised case. This is as expected because in fully supervised setting there would be less proposal error and hence less harm by its propagation.

5 Conclusions

In this work, we have proposed a novel *Proposal-Free Temporal Mask* (SPOT) learning model for the under-studied yet practically useful semi-supervised temporal action detection (SS-TAD). It is characterized by a *parallel* localization (mask generation) and classification architecture designed to solve the localization error propagation problem with conventional TAD models. For superior optimization, we further introduced a novel pretext task for pre-training and a boundary refinement algorithm. Extensive experiments on ActivityNet and THUMOS demonstrate that our SPOT yields state-of-the-art performance under both semi-supervised and supervised learning settings.

References

1. Alwassel, H., Mahajan, D., Torresani, L., Ghanem, B., Tran, D.: Self-supervised learning by cross-modal audio-video clustering. In: *NeurIPS (2020)* 4
2. 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) 14
3. Bao, H., Dong, L., Wei, F.: Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254* (2021) 8
4. Benaim, S., Ephrat, A., Lang, O., Mosseri, I., Freeman, W.T., Rubinstein, M., Irani, M., Dekel, T.: SpeedNet: Learning the speediness in videos. In: *CVPR (2020)* 4
5. Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.: Mixmatch: A holistic approach to semi-supervised learning. In: *NeurIPS (2019)* 3
6. 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) 10
7. Buch, S., Escorcia, V., Shen, C., Ghanem, B., Carlos Niebles, J.: Sst: Single-stream temporal action proposals. In: *CVPR (2017)* 1
8. 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. 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) 10
10. Chapelle, O., Scholkopf, B., Zien, A.: Semi-supervised learning. *TNNLS* (2009) 3
11. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *ICML (2020)* 4
12. Chen, X., He, K.: Exploring simple siamese representation learning. In: *CVPR (2021)* 4
13. Chen, Y., Zhu, X., Gong, S.: Semi-supervised deep learning with memory. In: *ECCV*. pp. 268–283 (2018) 3
14. Chen, Y.C., Li, L., Yu, L., El Kholy, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: *European conference on computer vision*. pp. 104–120. Springer (2020) 8
15. Chen, Y., Tu, Z., Ge, L., Zhang, D., Chen, R., Yuan, J.: So-handnet: Self-organizing network for 3d hand pose estimation with semi-supervised learning. In: *ICCV (2019)* 3
16. Dong, Q., Zhu, X., Gong, S.: Single-label multi-class image classification by deep logistic regression. In: *AAAI*. vol. 33, pp. 3486–3493 (2019) 9
17. Gao, J., Yang, Z., Chen, K., Sun, C., Nevatia, R.: Turn tap: Temporal unit regression network for temporal action proposals. In: *ICCV (2017)* 3
18. Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., et al.: Bootstrap your own latent: A new approach to self-supervised learning. In: *NeurIPS (2020)* 4
19. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *CVPR (2020)* 4
20. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9729–9738 (2020) 7

21. Ibrahim, M.S., Vahdat, A., Ranjbar, M., Macready, W.G.: Semi-supervised semantic image segmentation with self-correcting networks. In: CVPR (2020) [3](#)
22. 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](#)
23. Ji, J., Cao, K., Niebles, J.C.: Learning temporal action proposals with fewer labels. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7073–7082 (2019) [1](#), [4](#), [10](#), [11](#), [13](#)
24. Kim, J., Hur, Y., Park, S., Yang, E., Hwang, S.J., Shin, J.: Distribution aligning refinery of pseudo-label for imbalanced semi-supervised learning. In: NeurIPS (2020) [3](#)
25. Laine, S., Aila, T.: Temporal ensembling for semi-supervised learning. In: ICLR (2017) [3](#)
26. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: ICML Workshop (2013) [3](#)
27. 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) [3](#), [4](#), [10](#), [11](#), [12](#), [13](#), [14](#)
28. Lin, T., Zhao, X., Su, H., Wang, C., Yang, M.: BSN: Boundary sensitive network for temporal action proposal generation. In: ECCV (2018) [3](#), [14](#)
29. Little, W.A.: The existence of persistent states in the brain. In: From High-Temperature Superconductivity to Microminiature Refrigeration (1974) [8](#)
30. Long, F., Yao, T., Qiu, Z., Tian, X., Luo, J., Mei, T.: Gaussian temporal awareness networks for action localization. In: CVPR (2019) [3](#)
31. Miech, A., Alayrac, J.B., Smaira, L., Laptev, I., Sivic, J., Zisserman, A.: End-to-end learning of visual representations from uncurated instructional videos. In: CVPR (2020) [4](#)
32. 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) [9](#)
33. Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: CVPR (2020) [4](#)
34. Misra, I., Zitnick, C.L., Hebert, M.: Shuffle and learn: Unsupervised learning using temporal order verification. In: ECCV (2016) [4](#)
35. Mitra, R., Gundavarapu, N.B., Sharma, A., Jain, A.: Multiview-consistent semi-supervised learning for 3d human pose estimation. In: CVPR (2020) [3](#)
36. Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. *IEEE TPAMI* (2018) [3](#)
37. Nag, S., Zhu, X., Song, Y.Z., Xiang, T.: Temporal action localization with global segmentation mask transformers (2021) [1](#)
38. Nag, S., Zhu, X., Song, Y.z., Xiang, T.: Proposal-free temporal action detection via global segmentation mask learning. In: ECCV (2022) [1](#)
39. Nag, S., Zhu, X., Song, Y.z., Xiang, T.: Zero-shot temporal action detection via vision-language prompting. In: ECCV (2022) [1](#)
40. Nag, S., Zhu, X., Xiang, T.: Few-shot temporal action localization with query adaptive transformer. *arXiv preprint arXiv:2110.10552* (2021) [1](#)
41. Ouali, Y., Hudelot, C., Tami, M.: Semi-supervised semantic segmentation with cross-consistency training. In: CVPR (2020) [3](#)

42. Patrick, M., Asano, Y.M., Huang, B., Misra, I., Metze, F., Henriques, J., Vedaldi, A.: Space-time crop & attend: Improving cross-modal video representation learning. arXiv preprint arXiv:2103.10211 (2021) [7](#)
43. 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: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 485–494 (June 2021) [14](#)
44. 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](#)
45. 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) [6](#)
46. Shi, B., Dai, Q., Hoffman, J., Saenko, K., Darrell, T., Xu, H.: Temporal action detection with multi-level supervision. In: CVPR. pp. 8022–8032 (2021) [4](#), [11](#)
47. Sohn, K., Berthelot, D., Li, C.L., Zhang, Z., Carlini, N., Cubuk, E.D., Kurakin, A., Zhang, H., Raffel, C.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. arXiv preprint arXiv:2001.07685 (2020) [1](#), [3](#), [8](#), [10](#), [11](#), [13](#)
48. Sridhar, D., Quader, N., Muralidharan, S., Li, Y., Dai, P., Lu, J.: Class semantics-based attention for action detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13739–13748 (2021) [3](#), [14](#)
49. 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](#), [14](#)
50. Tang, Y.S., Lee, G.H.: Transferable semi-supervised 3d object detection from rgb-d data. In: ICCV (2019) [3](#)
51. Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. arXiv preprint arXiv:1703.01780 (2017) [1](#), [3](#), [10](#), [11](#)
52. 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) [4](#)
53. Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K.: Tracking emerges by colorizing videos. In: ECCV (2018) [4](#)
54. Wang, L., Xiong, Y., Lin, D., Van Gool, L.: Untrimmednets for weakly supervised action recognition and detection. In: CVPR. pp. 4325–4334 (2017) [1](#), [10](#), [11](#), [12](#), [13](#)
55. 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) [10](#)
56. Wang, T., Zhu, Y., Zhao, C., Zeng, W., Wang, J., Tang, M.: Adaptive class suppression loss for long-tail object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3103–3112 (2021) [9](#)
57. Wang, X., Zhang, S., Qing, Z., Shao, Y., Gao, C., Sang, N.: Self-supervised learning for semi-supervised temporal action proposal. In: CVPR. pp. 1905–1914 (2021) [1](#), [4](#), [10](#), [11](#), [13](#), [14](#)
58. Wei, D., Lim, J., Zisserman, A., Freeman, W.T.: Learning and using the arrow of time. In: CVPR (2018) [4](#)
59. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR (2018) [4](#)

60. Xie, Q., Dai, Z., Hovy, E., Luong, T., Le, Q.: Unsupervised data augmentation for consistency training. In: *NeurIPS* (2020) [3](#)
61. Xiong, Y., Wang, L., Wang, Z., Zhang, B., Song, H., Li, W., Lin, D., Qiao, Y., Van Gool, L., Tang, X.: Cuhk & ethz & siat submission to activitynet challenge 2016. arXiv preprint arXiv:1608.00797 (2016) [11](#)
62. Xu, H., Das, A., Saenko, K.: R-c3d: Region convolutional 3d network for temporal activity detection. In: *ICCV* (2017) [3](#)
63. 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 (2020) [11](#), [13](#)
64. 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) [4](#)
65. Xu, M., Perez-Rua, J.M., Zhu, X., Ghanem, B., Martinez, B.: Low-fidelity end-to-end video encoder pre-training for temporal action localization. In: *NeurIPS* (2021) [1](#), [4](#)
66. 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](#), [4](#), [14](#)
67. Yan, P., Li, G., Xie, Y., Li, Z., Wang, C., Chen, T., Lin, L.: Semi-supervised video salient object detection using pseudo-labels. In: *ICCV* (2019) [3](#)
68. 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](#)
69. Zhao, N., Chua, T.S., Lee, G.H.: Sess: Self-ensembling semi-supervised 3d object detection. In: *CVPR* (2020) [3](#)
70. Zhao, Y., Xiong, Y., Wang, L., Wu, Z., Tang, X., Lin, D.: Temporal action detection with structured segment networks. In: *ICCV* (2017) [1](#), [3](#)
71. Zhu, X.J.: Semi-supervised learning literature survey. Tech. rep., University of Wisconsin-Madison Department of Computer Sciences (2005) [3](#)