

Action Sensitivity Learning for Temporal Action Localization

Jiayi Shao^{1,*}, Xiaohan Wang¹, Ruijie Quan¹, Junjun Zheng², Jiang Yang², Yi Yang^{1,†}
¹ReLER Lab, CCAI, Zhejiang University, ²Alibaba Group

shaojiayil@zju.edu.cn, wxh1996111@gmail.com, quanruij@hotmail.com,
 {fangcheng.zjj,yangjiang.yj}@alibaba-inc.com, yangyics@zju.edu.cn

Abstract

Temporal action localization (TAL), which involves recognizing and locating action instances, is a challenging task in video understanding. Most existing approaches directly predict action classes and regress offsets to boundaries, while overlooking the discrepant importance of each frame. In this paper, we propose an Action Sensitivity Learning framework (ASL) to tackle this task, which aims to assess the value of each frame and then leverage the generated action sensitivity to recalibrate the training procedure. We first introduce a lightweight Action Sensitivity Evaluator to learn the action sensitivity at the class level and instance level, respectively. The outputs of the two branches are combined to reweight the gradient of the two sub-tasks. Moreover, based on the action sensitivity of each frame, we design an Action Sensitive Contrastive Loss to enhance features, where the action-aware frames are sampled as positive pairs to push away the action-irrelevant frames. The extensive studies on various action localization benchmarks (i.e., MultiThumos, Charades, Ego4D-Moment Queries v1.0, Epic-Kitchens 100, Thumos14 and ActivityNet1.3) show that ASL surpasses the state-of-the-art in terms of average-mAP under multiple types of scenarios, e.g., single-labeled, densely-labeled and egocentric.

1. Introduction

With an increasing number of videos appearing online, video understanding has become a prominent research topic in computer vision. Temporal action localization (TAL), which aims to temporally locate and recognize human actions with a set of categories in a video clip, is a challenging yet fundamental task in this area, owing to its various applications such as sports highlighting, human action analysis and security monitoring [25, 63, 46, 17, 14].

We have recently witnessed significant progress in TAL,

* This work was done during the first author’s internship in Alibaba.

† Corresponding Author.

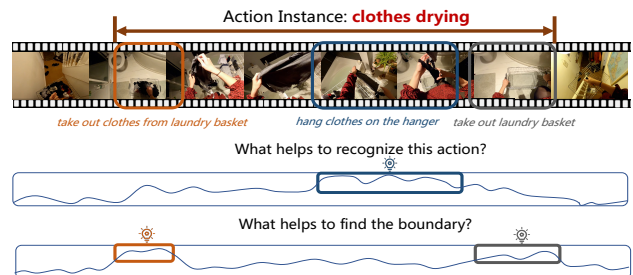


Figure 1. The motivation of our method. We show the action instance of **clothes drying** and depict the possible importance of each frame to recognizing the action category and locating action boundaries. Each frame’s importance is different.

where most methods can be mainly divided into two parts: 1) Two-stage approaches [75, 85] tackle this task accompanied by the generation of class-agnostic action proposals and then perform classification and proposal boundaries refinement in proposal-level; 2) One-stage approaches [79, 72, 32] simultaneously recognize and localize action instances in a single-shot manner. Typical methods [76, 29] of this type predict categories as well as locate corresponding temporal boundaries in frame-level, achieving stronger TAL results currently. In training, they classify every frame as one action category or background and regress the boundaries of frames inside ground-truth action segments. However, these works treat each frame within action segments equally in training, leading to sub-optimal performance.

When humans intend to locate action instances, the discrepant information of each frame is referred to. For the instance of action: **clothes drying**, as depicted in Fig 1, frames in the purple box promote recognizing **clothes drying** most as they describe the intrinsic sub-action: *hang clothes on the hanger*. Analogously, frames in red and gray boxes depict *take out clothes from laundry basket* and *lift laundry basket*, which are more informative to locate precise start and end time respectively. In a word, each frame’s contribution is quite different, due to intrinsic patterns of actions, as well as existing transitional or blurred frames.

Can we discover informative frames for classifying and

localizing respectively? To this end, we first introduce a concept — **Action Sensitivity**, to measure the frame’s importance. It is disentangled into two parts: **action sensitivity to classification** sub-task and **action sensitivity to localization** sub-task. For one sub-task, the higher action sensitivity each frame has, the more important it will be for this sub-task. With this concept, intuitively, more attention should be paid to action sensitive frames in training.

Therefore in this paper, we propose a lightweight Action Sensitivity Evaluator (ASE) for each sub-task to better exploit frame-level information. Essentially, for a specific sub-task, ASE learns the action sensitivity of each frame from two perspectives: class-level and instance-level. The class-level perspective is to model the coarse action sensitivity distribution of each action category and is achieved by incorporating gaussian weights. The instance-level perspective is complementary to class-level modeling and is supervised in a prediction-aware manner. Then the training weights of each frame are dynamically adjusted depending on their action sensitivity, making it more reasonable and effective for model training.

With the proposed ASE, we build our novel Action Sensitivity Learning framework dubbed **ASL** to tackle temporal action localization task (TAL) effectively. Moreover, to furthermore enhance the features and improve the discrimination between actions and backgrounds, we design a novel Action Sensitive Contrastive Loss (ASCL) based on ASE. It is implemented by elaborately generating various types of action-related and action-irrelevant features and performing contrasting between them, which brings multiple merits for TAL.

By conducting extensive experiments on 6 datasets and detailed ablation studies, we demonstrate ASL is able to classify and localize action instances better. In a nutshell, our main contributions can be summarized as follows:

- We propose a novel framework with an Action Sensitivity Evaluator component to boost training, by discovering action sensitive frames to specific sub-tasks, which is modeled from class level and instance level.
- We design an Action Sensitive Contrastive Loss to do feature enhancement and to increase the discrimination between actions and backgrounds.
- We verify ASL on various action localization datasets of multiple types: i) densely-labeled (i.e., Multi-Thumos [74] and Charades [53]). ii) egocentric (Ego4d-Moment Queries v1.0 [19] and Epic-Kitchens 100 [11]). iii) nearly single-labeled (Thumos14 [57] and ActivityNet1.3 [2]), and achieve superior results.

2. Related Works

Temporal Action Localization. Temporal action localization is a long-standing research topic. Contemporary ap-

proaches mostly fall into two categories, i.e. two-stage and one-stage paradigms. Previous two-stage methods usually focused on action proposal generation [31, 33, 56, 58, 65]. Others have integrated action proposal, calibrated backbone, classification and boundary regression or refinement modules into one single model [51, 69, 49, 81]. Recent efforts have investigated the proposal relations [75, 85, 66], utilized graph modeling [72, 75], or designed fine-grained temporal representation [44, 55]. One-stage approaches usually perform frame-level or segment-level classification and directly localization or merging segments [49, 80, 32]. [79, 39] process the video with the assistance of pre-defined anchors or learned proposals, while others utilize existing information and are totally anchor-free [29, 76, 78]. Currently, some works introduce pretrain-finetune to TAL task [70, 71] or attempt to train the model in an efficient end-to-end manner [38, 7, 37]. Others focused on densely-labeled setting [61, 10, 9, 24, 59, 8]. With the success of DETR [3] in object detection, query-based methods have also been proposed [48, 58, 59, 38]. Our method falls into the one-stage TAL paradigm and performs frame-level classification and localization. Notably, [43, 39] incorporate Gaussian kernels to improve receptive fields and optimize the temporal scale of action proposals, [24] use fixed gaussian-like weights to fuse the coarse and fine stage. We also utilize gaussian weights as one part of ASE, but it differs in that: i) Our gaussian-like weights in ASE serve as modeling class-level action sensitivity and to boost effective training, while [24, 43, 39] use it only to better encode the videos. ii) Our learned gaussian weights describe frames’ contributions to each sub-task and can be easily visualized, whereas the semantic meaning of gaussian weights in [24, 43, 39] is unclear. iii) Our gaussian-like weights are totally learnable, category-aware and disentangled to different sub-tasks.

One-stage Object Detection. Analogous to TAL task, the object detection task shares a few similarities. As a counterpart in object detection, the one-stage paradigm has surged recently. Some works remain anchor-based [35], while others are anchor-free, utilizing a feature pyramid network [34, 60] and improved label-assign strategies [77, 83, 84, 52]. Moreover, some works define key points in different ways (e.g. corner [26], center [13, 60] or learned points [73]). These methods bring some inspiration to design a better TAL framework. Some methods [16, 28, 27] aim to tackle the misalignment between classification and localization. But i) we mainly focus on the discrepant information of frames. ii) Misalignment of two sub-tasks (i.e., classification and localization) is only the second issue and we alleviate it by a novel contrastive loss which differs from these works.

Contrastive Learning. Contrastive learning [6, 20, 22] is an unsupervised learning objective that aims to bring sim-

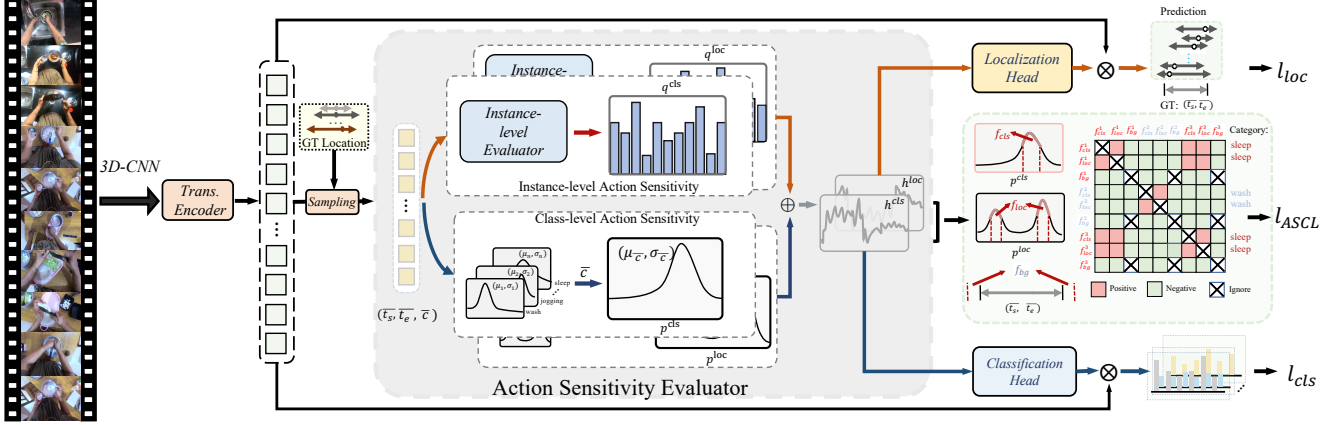


Figure 2. **The overview of ASL.** Given a video clip, we first leverage a pre-trained 3D-CNN to extract the video feature and then utilize a Transformer Encoder to encode feature. We then use ground-truth location sampling to sample all ground-truth segments and feed these into Action Sensitivity Evaluator. In this module, we model sub-task-specific action sensitivity of each frame from class level and instance-level. The former is learned by incorporating learnable gaussian-like weights and the latter is learned with an instance-level evaluator. Then each frame’s weight in training is adjusted based on action sensitivity. Moreover, we propose an Action Sensitive Contrastive Loss to better enhance the feature and alleviate misalignment problems.

ilar examples closer together in feature space while pushing dissimilar examples apart. NCE [21] and Info-NCE [41] are two typical methods that mine data features by distinguishing between data and noise or negative samples. InfoNCE-based contrastive learning has been used in methods of different tasks, such as [67, 36] in cross-modality retrieval and [23, 42] in unsupervised learning. In TAL, [29] leverages ranking loss to boost discrimination between foreground and background while [48] contrasts different actions with a global representation of action segments. But we design a new contrastive loss both across different types of actions and between actions and backgrounds. Moreover, compared to [50] which also contrasts between actions and backgrounds, our proposed contrastive loss contrasts more between i) same and different action classes, ii) sensitive frames of localization and classification to mitigate the misalignment of sub-tasks. Details will be discussed in 3.3.

3. Method

Problem Formulation. The task of temporal action localization (TAL) is to predict a set of action instances $\{(t_m^s, t_m^e, c_m)\}_{m=1}^M$, given a video clip, where M is the number of predicted action instances, t_m^s, t_m^e, c_m are the start, end timestamp and action category of the m -th predicted action instance. ASL is built on an anchor-free representation that classifies each frame as one action category or background, as well as regresses the distances from this frame to the start time and end time.

Overview. The overall architecture of ASL is shown in Fig 2. ASL is composed of four parts: video feature extractor, feature encoder, action sensitivity evaluator, and two

sub-task heads. Concretely, given a video clip, we first extract the video feature using a pre-trained 3D-CNN model. Then we exert a feature encoder involving a pyramid network to better represent the temporal features at multiple levels. We propose an action sensitivity evaluator module to access the action sensitivity of frames to a specific sub-task. The pyramid features combined with frames’ action sensitivity are further processed by sub-task heads to generate predictions. We now describe the details of ASL.

3.1. Feature Encoder

With the success of [76, 29], ASL utilizes a Transformer encoder and feature pyramid network to encode feature sequences into a multiscale representation. To enhance features, in Transformer encoder we design a new attention mechanism that operates temporal attention and channel attention parallelly and then fuses these two outputs.

For normal temporal attention that is performed in the temporal dimension, input features generate query, key and value tensors $(Q_t, K_t, V_t) \in \mathbb{R}^{T \times D}$, where T is the number of frames, D is the embedding dimension, then the output is calculated:

$$f'_{ta} = \text{softmax}\left(\frac{Q_t K_t^T}{\sqrt{D}}\right) V_t \quad (1)$$

For channel attention that is conducted in the channel dimension, input features generate query, key and value tensors $(Q_d, K_d, V_d) \in \mathbb{R}^{D \times T}$, where D is the number of channels. Then the output is calculated:

$$f'_{ca} = \text{softmax}\left(\frac{Q_d K_d^T}{\sqrt{T}}\right) V_d \quad (2)$$

Above two outputs are then added with a coefficient θ : $f' = (1 - \theta)f'_{ta} + \theta f'_{ca}$. Then it is processed by layer nor-

malization and feedforward network to obtain the encoded video representation $f \in \mathbb{R}^{T \times D}$.

3.2. Action Sensitivity Evaluator

As discussed in 1, not all frames inside ground-truth segments contribute equally to the sub-task (i.e., localization or classification). Thus we designed an Action Sensitivity Evaluator (ASE) module, the core idea of which is to determine the sub-task-specific action sensitivity of each frame and help the model pay more attention to those valuable frames. Besides, this module is lightweight, leading to efficient and effective training.

Decoupling to two levels. Digging into action instances, a key observation is that actions of a particular category often share a similar pattern, but they appear slightly different in diverse scenarios or under different behavior agents. For example, action instances of category: **wash vegetables** inherently contain sub-actions: *turn the tap on*, *take vegetables*, *wash*, *turn the tap off*, where frames depicting *washing* are more sensitive to classification, frames depicting *turning the tap on* and *turning the tap off* are more sensitive to localization. But the respective duration or proportion of these sub-actions are dependent on the scenes and context of each action instance, thus making sensitive frames a little different. This motivates us that the action sensitivity of every frame should be decoupled into class-level and instance-level modeling and then recombined from these two parts.

Disentangling to two sub-tasks. Here sub-tasks mean classification and localization. Intuitively action sensitivity for classification needs to be modeled as sensitive frames for classification is not easily determined. Actually, action sensitivity modeling for localization is also necessary. Though the boundaries of action segments are defined already, sensitive frames are not necessarily at the start or the end of an action since i) action boundaries are often unclear, ii) each frame of sub-actions around boundaries also has different semantics. Therefore, action sensitivity modeling should be disentangled for two sub-tasks respectively (i.e., classification and localization).

Formally, for a given ground-truth $\mathcal{G} = \{\bar{t}^s, \bar{t}^e, \bar{c}\}$, each indicating the start time, end time and category of one action, we denote N_f as the number of frames within this action, N_c as the number of all pre-defined action categories. Our goal is to model the class-level action sensitivity p (disentangled into p^{cls} , p^{loc} to classification and localization respectively), instance-level action sensitivity q (disentangled to q^{cls} , q^{loc}). Then we delve into details of action sensitivity learning.

Class-level Modeling. Class-level sensitivity poses a fundamental prior for action sensitivity learning. Two key observations are that: i) video frames are often consecutive. ii) there often exist keyframes that have a peak value of sensitivity among all frames. In this case, we in-

corporate gaussian-like weights with learnable parameters $\mu, \sigma \in \mathbb{R}^{N_c}$ to model class-level action sensitivity p .

For classification sub-task, we model corresponding action sensitivity p_i^{cls} for the i -th frame:

$$p_i^{cls} = \exp\left\{-\frac{(d(i) - \mu_c)^2}{2\sigma_c^2}\right\} \quad (3)$$

where $d(i)$ is the distance from the current i -th frame to the central frame of the ground-truth segment which is normalized by N_f . In this case, $d(i) \in [-0.5, 0.5]$, when $i=1$ (i.e., start frame), $d(i) = -0.5$, when $i = N_f$ (i.e., end frame), $d(i) = 0.5$. Learnable parameters μ_c, σ_c denote mean and variance of each category c 's action sensitivity distribution.

For localization sub-task, different frames are sensitive to locating start time and end time. Therefore action sensitivity p^{loc} is the combination of two parts. We explicitly allocate one gaussian-like weights p^{sot} to model the start time locating sensitivity and another p^{eot} to model the end time locating sensitivity. p^{loc} is calculated:

$$p_i^{loc} = \underbrace{\exp\left\{-\frac{(d(i) - \mu_{c,1})^2}{2\sigma_{c,1}}\right\}}_{p_i^{sot}} + \underbrace{\exp\left\{-\frac{(d(i) - \mu_{c,2})^2}{2\sigma_{c,2}}\right\}}_{p_i^{eot}} \quad (4)$$

In this way, class-level action sensitivity $p^{cls}, p^{loc} \in \mathbb{R}^{N_f \times N_c}$ of all categories are learned with the optimization of model training. In addition, the initialization of μ_c and σ_c counts as there exists prior knowledge [76, 60] according to different sub-tasks. For classification sub-task, near-center frames are more sensitive. Thus we initialize μ_c as 0. For localization sub-task, near-start and near-end frames are more sensitive. Thus we initialize μ_1 as -0.5 and μ_2 as 0.5. For all σ , we initialize as 1.

Instance-level Modeling. Intuitively, a Gaussian can only give a single peak, and thus class-level action sensitivity learning may not discover all sensitive frames. To this end, we introduce instance-level modeling which is complementary and aims to capture additional important frames that haven't been discovered by class-level modeling.

In the instance-level modeling, as more information about frame contexts of each instance is referred to, we obtain instance-level action sensitivity $q \in \mathbb{R}^{N_f}$ using an instance-level evaluator operated directly on each frame, composed of 1D temporal convolution network which aims to encode temporal contexts better, a fully connected layer and a Sigmoid activation function. We denote Φ^{cls} and Φ^{loc} as two sub-task specific instance-level evaluator, then q^{cls} and q^{loc} are computed:

$$\begin{cases} q_i^{cls} = \Phi^{cls}(f_i) \\ q_i^{loc} = \Phi^{loc}(f_i) \end{cases} \quad (5)$$

Unlike class-level modeling that contains some prior knowledge, instance-level sensitivity q is hard to learn in an

unsupervised manner. Intuitively, from the instance level a sensitive frame implies that it can result in fine predictions. Hence we utilize the quality $\{\bar{Q}_i\}_{i=1}^{N_f}$ of each frame’s prediction to supervise the learning of q . For localization, The higher tIoU indicates a higher degree of overlap between two segments. Thus tIoU between the predicted segment and the ground-truth segment can measure the quality of prediction. For classification, the probability of the ground-truth category can serve as the quality of prediction. Therefore, quality \bar{Q}^{cls} and \bar{Q}^{loc} are defined as:

$$\begin{cases} \bar{Q}_i^{cls} = \varphi(s_i[\bar{c}]) \\ \bar{Q}_i^{loc} = \text{tIoU}(\Delta_i, \bar{\Delta}) \end{cases} \quad (6)$$

where s denotes the classification logits, Δ_i is the predicted segment (t^s, t^e) of the i -th frame, $\bar{\Delta}$ is the corresponding ground-truth segment, $\varphi(\cdot)$ is Sigmoid function. We use MSE loss to supervise the calculation of q . For q^{cls} , optimization objective is formed as 7. Optimization of q^{loc} is in a similar way.

$$\mathcal{L}_s = \text{MSE}(q^{cls}, \bar{Q}^{cls}) \quad (7)$$

Optimization with Action Sensitivity. In this way, combining class-level and instance-level, we obtain the final action sensitivity $h(\bar{c}) \in \mathbb{R}^{N_f}$ (disentangled to classification and localization sub-task: $h(\bar{c}) \rightarrow \{h^{cls}(\bar{c}), h^{loc}(\bar{c})\}$) for the ground-truth $\mathcal{G} = \{\bar{t}^s, \bar{t}^e, \bar{c}\}$:

$$\begin{cases} h^{cls}(\bar{c}) = p^{cls} \mathbb{1}[\bar{c}] + q^{cls} \\ h^{loc}(\bar{c}) = p^{loc} \mathbb{1}[\bar{c}] + q^{loc} \end{cases} \quad (8)$$

where $\mathbb{1}[\bar{c}] \in \mathbb{R}^{N_c}$ denotes the one-hot vector of \bar{c} . Action sensitivity h is further used in training. For classification sub-task, we use a focal loss [35] to classify each frame, combined with classification action sensitivity h^{cls} :

$$\mathcal{L}_{cls} = \frac{1}{N_{pos}} \sum_i (\mathbb{1}_{in_i} h_i^{cls}(\bar{c}_i) \mathcal{L}_{focal_i} + \mathbb{1}_{bg_i} \mathcal{L}_{focal_i}) \quad (9)$$

where $\mathbb{1}_{in_i}, \mathbb{1}_{bg_i}$ are indicators that denote if the i -th frame is within one ground-truth action or if is background, N_{pos} is the number of frames within action segments, \bar{c}_i denotes the action category of the i -th frame.

For localization sub-task, we use a DIOU loss [82] performed on frames within any ground-truth action instance, to regress offsets from current frames to boundaries, combined with localization action sensitivity h^{loc} :

$$\mathcal{L}_{loc} = \frac{1}{N_{pos}} \sum_i (\mathbb{1}_{in_i} h_i^{loc}(\bar{c}_i) \mathcal{L}_{DIOU_i}) \quad (10)$$

3.3. Action Sensitive Contrastive Loss

Now with ASE, each frame is equipped with action sensitivity and valuable frames to specific sub-tasks will be discovered. We further boost the training from the perspective

of feature enhancement. Delve into feature representation, three shortcomings may hinder the performance: i) classification sensitive and localization sensitive frames are quite different, resulting in the misalignment of these two sub-tasks. ii) features in actions of different categories are not much discriminable. iii) features within action and outside boundaries are not much distinguished yet.

Therefore, on the basis of ASE, we propose an Action Sensitive Contrastive Loss (ASCL) to correspondingly tackle the above issues. Specifically, for a given video feature $\{f_t\}_{t=1}^T$ and a ground-truth action instance $\mathcal{G} = \{\bar{t}^s, \bar{t}^e, \bar{c}\}$, we generate two action-related features and one action-irrelevant feature. First, to generate more valuable action-related features, we aim to find sensitive frames to these sub-tasks. Thinking that ASCL contrasts action instances of different classes, where class-level discrimination is more important, we hence utilize class-level sensitivity p to parse the sensitive frame ranges T_{cls} for classification and T_{loc} for localization. With one ground-truth category \bar{c} , we get the most sensitive frames $a_{cls}, a_{sot}, a_{eot}$ for classification, start time localization, end time localization respectively. Take a_{eot} as an example:

$$a_{eot} = \arg \max_i (p_i^{eot} \mathbb{1}[\bar{c}]) \quad (11)$$

a_{cls} and a_{sot} are obtained in a similar way. Then, centered on a and extending forward and backward with a range of δN_f , where δ is the sampling length ratio, we get sensitive frame ranges T_{cls} for classification and T_{loc} for localization (T_{cls} and T_{loc} are limited inside the action instance). Furthermore, we utilize class-level sensitivity to compute sensitive features f_{cls} for classification, f_{loc} for localization:

$$\begin{cases} f_{cls} = \frac{1}{T} \sum_t p_t^{cls} \mathbb{1}[\bar{c}] f_t, t \in T_{cls} \\ f_{loc} = \frac{1}{T} \sum_t p_t^{loc} \mathbb{1}[\bar{c}] f_t, t \in T_{loc} \end{cases} \quad (12)$$

Secondly, we aim to simultaneously discriminate actions and backgrounds better. Consequently we generate boundary-related background features f_{bg} :

$$f_{bg} = \frac{1}{T} \sum_t f_t, t \in [\bar{t}^s - \delta N_f, \bar{t}^s] \cup [\bar{t}^e, \bar{t}^e + \delta N_f] \quad (13)$$

The learning objective of ASCL is based on a contrastive loss. As figure 2 shows, the positive samples \mathcal{P} are constructed from f_{cls} and f_{loc} in action instances of the same category while the negative samples \mathcal{N} come from: i) f_{cls} and f_{loc} in action instances of different categories. ii) all background features f_{bg} . ASCL is computed for each batch B with N samples:

$$\mathcal{L}_{ASCL} = \frac{1}{N} \sum_B -\log \frac{\sum_{f_x \in \mathcal{P}_{f_*}} \text{sim}(f_*, f_x)}{\sum_{f_x \in \mathcal{P}_{f_*}} \text{sim}(f_*, f_x) + \sum_{f_x \in \mathcal{N}_{f_*}} \text{sim}(f_*, f_x)} \quad (14)$$

Optimizing ASCL will be of benefits to tackle the corresponding issues above : i) alleviate the misalignment of two sub-tasks by pulling features of their respective sensitive frames closer. ii) discriminate actions and backgrounds better by pushing action features of the same category closer and different categories apart, meanwhile pushing actions and backgrounds apart. Thus ASCL can enhance the feature representation and boost training furthermore.

3.4. Training and Inference

Training. In the training process, our final loss function is designed:

$$\mathcal{L} = \mathcal{L}_{cls} + \mathcal{L}_{loc} + \mathcal{L}_s + \lambda \mathcal{L}_{ASCL} \quad (15)$$

where \mathcal{L}_{cls} , \mathcal{L}_{loc} and \mathcal{L}_s are discussed in equation 9, equation 10 and equation 7. λ denotes the weight of Action Sensitive Contrastive loss.

Inference. At inference time, our model outputs predictions (t^s, t^e, c) for every frame across all pyramids levels, where t^s, t^e denote the start and end time of action, c denote the predicted action category. c also serves as the action confidence score. SoftNMS [1] is then applied on these results to suppress redundant predictions.

4. Experiments

4.1. Datasets and Evaluation Metric

Datasets. To validate the efficacy of the proposed ASL, extensive experiments on 6 datasets of 3 types are conducted: i) densely-labeled: **MultiThumos**[74] and **Charades**[53]; ii) densely-labeled and egocentric: **Ego4D-Moment Queries v1.0**[19] and **Epic-Kitchens 100**[11]; iii) single-labeled: **Thumos14**[57] and **ActivityNet1.3**[2].

MultiThumos is a densely labeled dataset including 413 sports videos of 65 classes. Charades is a large multi-label dataset containing 9848 videos of 157 action classes. These two datasets are both densely labeled and hence have multiple action instances in each video clip, where different actions may occur concurrently.

Ego4D-Moment Queries v1.0 (Ego4D-MQ1.0 for short) is a large-scale egocentric benchmark with 2,488 video clips and 22.2K action instances from 110 pre-defined action categories, which is densely labeled and composed of long clips. EPIC-Kitchens 100 is a large egocentric action dataset containing 100 hours of videos from 700 sessions capturing cooking activities in different kitchens. These two datasets are both large, egocentric and densely labeled.

Thumos14 is composed of 200 validation videos and 212 testing videos from 20 action classes while ActivityNet has 19,994 videos with 200 action classes. These two datasets are singly labeled and thus most of video clips in them have one action instance in each video clip.

Evaluation Metric. Since ASL focuses on action detection, we take mean Average Precision (**mAP**) at certain tIoU thresholds as the evaluation metric. For all six datasets, we also report **average mAP** over several tIoU thresholds as the main metric. The tIoU thresholds are set consistent with the official setup or previous methods, which is detailed in the caption of Table 1, 2, 3, 4.

4.2. Implementation Details.

We follow the practice of using off-the-shelf pre-extracted features as input, specifically I3D [4] RGB features for MultiThumos, Charades, Thumos14 and ActivityNet, EgoVLP [30], Slowfast [15] and Omnivore [18] features for Ego4D-MQ1.0, Slowfast features [15, 12] for Epic-Kitchens 100.

We train our model with a batch size of 2, 16, 2, 2 for 60, 30, 15, 25 epochs on MultiThumos, Charades, Ego4D-MQ1.0 and Epic-Kitchens 100 respectively, where the learning rate is set to $2e^{-4}$. On ActivityNet and Thumos, we train our model with the batch size of 16, 2, the learning rate of $1e^{-3}$, $1e^{-4}$ for 15, 30 epochs. We set λ as 0.3 and θ as 0.2.

In post-processing, we apply softNMS [1] to suppress redundant predictions. For fair comparison, we keep 200, 100, and 2000, 2000 predictions on Thumos14, ActivityNet, Ego4D-MQ1.0 and Epic-Kitchens 100 respectively. As on MultiThumos and Charades, considering that PointTAD [59] splits a video clip into more than 4 parts and generates 48 predictions for each part, we keep 200 predictions on these two datasets.

In the training process, we clamp σ with a threshold (set as 5.0) to ensure σ won't be very large and thus prevent very small p^{cls}, p^{loc} , which may cause trivial solution to minimize the loss. Moreover, We tackle the issue of overlapped actions following [76, 60]: i)use multi-scale mechanism [34] to assign actions with different duration to different feature levels. ii)If a frame, even with multi-scale used, is still assigned to more than one ground-truth action, we choose the action with the shortest duration as its ground-truth target and model its action sensitivity based on this ground-truth.

4.3. Main Results

MultiThumos and Charades: We compare ASL with state-of-the-art methods under detection-mAP on these two densely-labeled TAL benchmarks. PDAN[10], coarse-fine[24], MLAD[61], MS-TCT[9] are based on frame-level representation, while PointTAD[59] are query-based. As shown in Table 1, ASL reaches the highest mAP over all tIoU thresholds, outperforming the previous best method(i.e. PointTAD) by 2.0% absolute increase of average mAP on MultiThumos and 3.3% on Charades. Notably, PointTAD is further trained in an end-to-end manner with

Table 1. **Results on MultiThumos and Charades.** We report detection-*mAP* at different tIoU thresholds. Average *mAP* in [0.1:0.1:0.9] is reported on MultiThumos and Charades. Best results are in **bold**. ‡ indicates results trained with stronger image augmentation [59, 38]. I3D denotes using I3D [4] features and E2E indicates results trained in an end-to-end manner.

Model	Modality	Feature	MultiThumos				Charades			
			0.2	0.5	0.7	Avg.	0.2	0.5	0.7	Avg.
PDAN [10]	RGB	I3D	-	-	-	17.3	-	-	-	8.5
Coarse-Fine [24]	RGB	I3D	-	-	-	-	-	-	-	6.1
MLAD [61]	RGB	I3D	-	-	-	14.2	-	-	-	-
MS-TCT [9]	RGB	I3D	-	-	-	16.3	-	-	-	7.9
PointTAD [59]	RGB	I3D-E2E	36.8	23.3	11.0	21.7	15.9	12.6	8.5	11.3
PointTAD‡ [59]	RGB	I3D-E2E	39.7	24.9	12.0	23.5	17.5	13.5	9.1	12.1
ASL	RGB	I3D	42.4	27.8	13.7	25.5	24.5	16.5	9.4	15.4

Table 2. **Results on Ego4D-Moment Queries v1.0.** We report *mAP* at different tIoU thresholds. Average *mAP* in [0.1:0.1:0.5] is reported on Ego4D-Moment Queries. Best results are in **bold**. EgoVLP, SF and OF denote EgoVLP [30], Slowfast [15] and Omnivore [18] features. InterVideo [5] denotes features extracted from VideoMAE-L [62] and fine-tuned on Ego4D-Moment Queries.

Method/Entry	Feature	mAP at IoUs, Val set				mAP at IoUs, Test set
		0.1	0.3	0.5	Avg.	Avg.
VSGN [79]	SF	9.10	5.76	3.41	6.03	5.68
VSGN [30]	EgoVLP	16.63	11.45	6.57	11.39	10.33
RELER [47]	SF+OV	22.75	17.61	13.43	17.94	17.67
Actionformer [40]	EgoVLP	26.84	20.57	14.54	20.60	-
Actionformer [40]	EgoVLP+SF+OV	28.26	21.88	16.28	22.09	21.76
Actionformer [5]	InternVideo	-	-	-	23.29	23.59
ASL	EgoVLP	29.45	23.03	16.08	22.83	22.25
ASL	EgoVLP+SF+OV	30.50	24.39	17.45	24.15	23.97

Table 3. **Results on EPIC-Kitchens 100 val set.** We report *mAP* at different tIoU thresholds and average *mAP* in [0.1:0.1:0.5]. All methods use the same SlowFast [15, 12] features.

Sub-Task	Method	0.1	0.3	0.5	Avg
Verb	BMN [31]	10.8	8.4	5.6	8.4
	G-TAD [72]	12.1	9.4	6.5	9.4
	Actionformer [76]	26.6	24.2	19.1	23.5
	ASL	27.9	25.5	19.8	24.6
Noun	BMN [31]	10.3	6.2	3.4	6.5
	G-TAD [72]	11.0	8.6	5.4	8.4
	Actionformer [76]	25.2	22.7	17.0	21.9
	ASL	26.0	23.4	17.7	22.6

strong image augmentation while ASL is feature-based, indicating that ASL performs more accurate TAL with more efficiency on densely-labeled datasets.

Ego4D-MQ1.0 and Epic-Kitchens 100: These two datasets are both challenging as they are large-scale, egocentric, densely labeled and composed of longer clips. Table 2 reports the results on Ego4D-MQ1.0. The state-of-the-art methods are all based on Actionformer[76] and perform frame-level recognition and localization with strong features. Using the same feature EgoVLP[30], ASL surpasses the current best entry[40]. Using the combined EgoVLP, slowfast[15] and omnivore[18] features, ASL

gains 2.06% improvement of average *mAP* on Val set and 2.21% on Test set. Moreover, ASL performs better than [5] which uses a stronger but not open-sourced InternVideo [5] feature. Meanwhile, on Epic-Kitchens 100 as table 3 shows, ASL outperforms the strong performance of Actionformer[76], BMN[31] and G-TAD[72] with the same Slowfast feature[15, 12]. The above results demonstrate the advantage of ASL on the challenging, egocentric and densely labeled benchmark.

Thumos14 and ActivityNet1.3: These two datasets are popular and nearly single-labeled, with approximately one action instance in each clip. Table 4 compares the results of ASL with various state-of-the-art methods (e.g., two-stage methods: BSN[33], G-TAD[72], P-GCN[75], RTDNet[58], one-stage methods: AFSD[29], SSN[81], Actionformer[76]). On Thumos14, across all tIoU thresholds, ASL achieves the best and gains 1.1% improvement of average *mAP* (67.9% v.s. 66.8%). On ActivityNet, ASL also outperforms previous methods of *mAP*@0.75 and average *mAP*, though the gap is slight. One possible reason is that due to the success of action recognition on ActivityNet, we follow the common practice [76, 79, 85] to fuse external video-level classification scores [68]. In this case, class-level sensitivity will not play an important role in training. Another reason may be that since each video in ActivityNet is nearly single-labeled, our proposed ASL will

Table 4. **Results on Thumos14 and ActivityNet1.3.** We report *mAP* at different tIoU thresholds. Average *mAP* in [0.3:0.1:0.7] is reported on THUMOS14 and [0.5:0.05:0.95] on ActivityNet1.3. The best results are in **bold**.

Model	Feature	Thumos14						ActivityNet1.3			
		0.3	0.4	0.5	0.6	0.7	Avg.	0.5	0.75	0.95	Avg.
BSN [33]	TSN [64]	53.5	45.0	36.9	28.4	20.0	36.8	46.5	30.0	8.0	30.0
BMN [31]	TSN [64]	56.0	47.4	38.8	29.7	20.5	38.5	50.1	34.8	8.3	33.9
G-TAD [72]	TSN [64]	54.5	47.6	40.3	30.8	23.4	39.3	50.4	34.6	9.0	34.1
P-GCN [75]	I3D [4]	63.6	57.8	49.1	-	-	-	48.3	33.2	3.3	31.1
TCANet [44]	TSN [64]	60.6	53.2	44.6	36.8	26.7	44.3	52.3	36.7	6.9	35.5
ContextLoc [85]	I3D [4]	68.3	63.8	54.3	41.8	26.2	50.9	56.0	35.2	3.6	34.2
VSGN [79]	TSN [64]	66.7	60.4	52.4	41.0	30.4	50.2	52.4	36.0	8.4	35.1
RTD-Net [58]	I3D [4]	68.3	62.3	51.9	38.8	23.7	49.0	47.2	30.7	8.6	30.8
SSN [81]	TS [54]	51.0	41.0	29.8	-	-	-	43.2	28.7	5.6	28.3
GTAN [39]	P3D [45]	57.8	47.2	38.8	-	-	-	52.6	34.1	8.9	34.3
AFSD [29]	I3D [4]	67.3	62.4	55.5	43.7	31.1	52.0	52.4	35.3	6.5	34.4
React [48]	I3D [4]	69.2	65.0	57.1	47.8	35.6	55.0	49.6	33.0	8.6	32.6
TadTR [38]	I3D [4]	62.4	57.4	49.2	37.8	26.3	46.6	49.1	32.6	8.5	32.3
Actionformer [76]	I3D [4]	82.1	77.8	71.0	59.4	43.9	66.8	54.2	36.9	7.6	36.0
ASL	I3D [4]	83.1	79.0	71.7	59.7	45.8	67.9	54.1	37.4	8.0	36.2

Table 5. **Ablation studies of components.** ASE: Action Sensitivity Evaluator. class.: class-level modeling. inst.: instance-level modeling. ASCL: Action Sensitive Contrastive Loss.

#	Components			mAP at different tIoUs			
	ASE class.	ASE inst.	ASCL	0.2	0.5	0.7	Avg.
1				39.6	25.9	11.6	23.4
2	✓			41.0	26.5	12.9	24.5
3		✓		40.5	26.2	12.0	23.9
4			✓	40.2	26.1	11.8	23.7
5	✓		✓	41.9	27.0	13.6	25.1
6	✓	✓		41.8	27.2	13.3	25.0
7	✓	✓	✓	42.4	27.8	13.7	25.5

be short of positive and negative samples, leading to a non-significant increase compared to improvements on densely labeled datasets as Table 1, 2.

4.4. Ablation Study

To further verify the efficacy of our contributions, we analyze main components of ASL on MultiThumos.

Action Sensitive Evaluator. Our proposed ASE can be divided into class-level and instance-level modeling. We first investigate the effect of these parts. In Table 5, baseline 1 denotes using our proposed framework without ASE and ASCL. After being equipped with class-level modeling, it boosts the performance by 1.1% of average mAP (baseline 2 v.s. baseline 1). When further adding instance-level bias, it gains 0.5% absolute increase (baseline 6 v.s. baseline 2). And our ASE contributes a total improvement of 1.6% on average mAP (baseline 7 v.s. baseline 1). It is obvious that action sensitivity modeling from both class-level and instance-level is beneficial to TAL task.

Table 6. **Ablation studies of Gaussians weights.** cls and loc denotes classification and localization sub-task. For gaussian weights in class-level action sensitivity learning, **learnable/fixed** denotes parameters learnable/not learnable. **None** denotes not using gaussian weights.

#	cls.	loc.	0.1	0.3	0.5	Avg.
1	None	None	40.9	26.3	12.3	24.2
2		None	40.9	26.5	12.4	24.4
3	fixed	fixed	41.0	26.6	12.7	24.6
4		learnable	41.7	26.8	13.0	24.9
5		None	41.9	27.1	13.0	24.9
6	learnable	fixed	42.0	26.9	13.4	25.1
7		learnable	42.4	27.8	13.7	25.5

Gaussian Weights. Then we analyze the effect of learnable gaussian weights in class-level action sensitivity learning. Table 6 demonstrates that compared to baseline 1 which does not use any gaussian weights to learn action sensitivity, fixed gaussian weights with prior knowledge do bring benefits (baseline 2,3 v.s. baseline 1). Meanwhile, learnable gaussian weights are more favored (baseline 4 v.s. baseline 3, baseline 7 v.s. baseline 6). Moreover, learnable gaussian weights for both two sub-tasks achieve the best results.

We further study the number of Gaussians used in classification and localization sub-task. As shown in Table 7, using two Gaussians for localization and one Gaussian for classification achieves the best results. It is probably because on the one hand, using two Gaussians for localization explicitly allocates one for modeling start time and one for modeling end time. On the other hand, more Gaussian weights may be a burden for training, leading to inferior performance.

Table 7. **Ablation studies of number of Gaussians weights.** #cls and #loc denote the number of Gaussian weights used in classification and localization sub-task. **shared** indicates two sub-tasks share one Gaussian weights.

#cls	#loc	0.1	0.3	0.5	Avg
1(shared)	1(shared)	42.2	27.2	13.7	25.3
0	0	40.9	26.3	12.3	24.2
	1	41.5	26.9	13.0	24.8
	2	41.6	27.1	13.4	25.0
1	0	42.2	27.1	13.2	25.1
	1	42.0	26.7	13.1	24.9
2	2	42.4	27.8	13.7	25.5
	0	42.3	26.9	13.3	25.1
	1	41.8	26.9	13.0	25.0
	2	42.0	27.2	13.6	25.3




Figure 3. Ablation of hyperparameters in ASCL.

Action Sensitive Contrastive Loss. Moreover, we delve into our proposed ASCL. As shown in Table 5, ASCL improves around 0.6% of average mAP on the basis of class-level prior (baseline 5 v.s. baseline 2) and 0.5% on the basis of ASE (baseline 7 v.s. baseline 6). Baseline 4, where using ASCL alone denotes sampling near the center frame to form f_{cls} and f_{loc} directly, also gains an improvement of 0.3% compared to the vanilla framework (baseline 4 v.s. baseline 1). This indicates the effectiveness of contrast between actions and backgrounds. When performing ASCL based on ASE, it will facilitate the final performance more because it can alleviate the misalignment as discussed in 3.3.

Finally we discussed the hyperparameters in ASCL. Fig 3(a) shows the performance curve of average mAP corresponding to ASCL weight λ . Average mAP on MultiThumos generally improves when λ increases and slightly drop as λ reaches 0.4. Fig 3(b) reports the average mAP to different sampling length ratios δ . When δ equals 0.2, our method achieves the best. In this case, we set λ to 0.3 and δ to 0.2.

4.5. Qualitative Experiment

To better illustrate the effectiveness of ASL, we visualize some qualitative results of Ego4D-MQ1.0 benchmark in Fig 4. We show that i) frames depicting action’s main sub-action (i.e., *hang clothes on the hanger*, *water run through hands*) are of higher action sensitivity for classification. ii) Frames depicting near-start and near-end sub-action (i.e., *turn the tap on*, *lift laundry basket*, e.t.c.) are of higher ac-

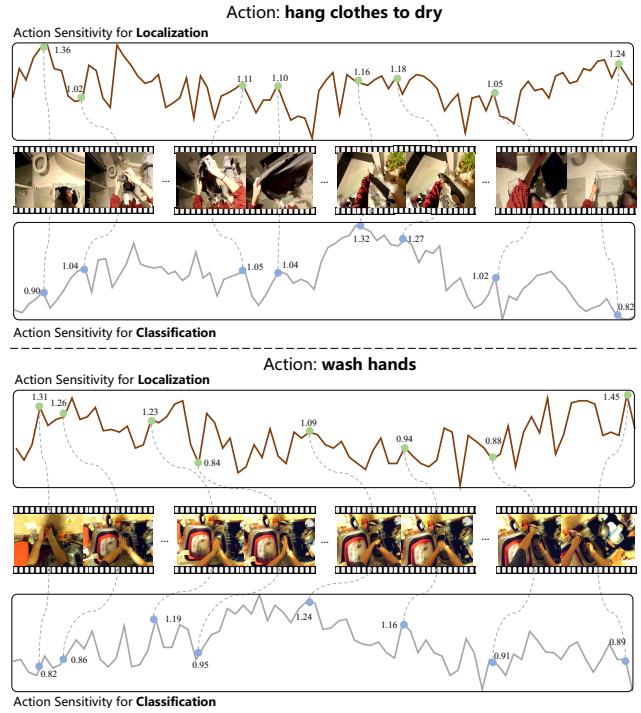


Figure 4. Visualization of (Top) the frame sensitivity to sub-tasks of Action: **hang clothes to dry** and (bottom) Action: **wash hands**. Please zoom in for the best view.

tion sensitivity for localization. Moreover, action sensitivity of frames is not continuous, as our proposed instance-level action sensitivity is discrete partly because blurred or transitional frames exist in video clips.

5. Conclusion

In this paper, we introduce an Action Sensitivity Learning framework (ASL) for temporal action localization (TAL). ASL models action sensitivity of each frame and dynamically change their weights in training. Together with the proposed Action Sensitive Contrastive Loss (ASCL) to further enhance features and alleviate misalignment, ASL is able to recognize and localize action instances effectively. For accurate TAL, fine-grained information should be considered (e.g. frame-level information). We believe that ASL is a step further in this direction. In the future, efforts could be paid to more complicated sensitivity modeling. Besides, ASL could also be redesigned as a plug-and-play component that will be beneficial to various TAL methods.

Acknowledgements: This work is supported by the Fundamental Research Funds for the Central Universities (No.226-2023-00048) and Major Program of the National Natural Science Foundation of China (T2293720/T2293723)

References

- [1] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S Davis. Soft-nms—improving object detection with one line of code. In *Proceedings of the IEEE international conference on computer vision*, pages 5561–5569, 2017. 6
- [2] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–970, 2015. 2, 6
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. 2
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 6, 7, 8
- [5] Guo Chen, Sen Xing, Zhe Chen, Yi Wang, Kunchang Li, Yizhuo Li, Yi Liu, Jiahao Wang, Yin-Dong Zheng, Bingkun Huang, Zhiyu Zhao, Junting Pan, Yifei Huang, Zun Wang, Jiashuo Yu, Yanan He, Hongjie Zhang, Tong Lu, Yali Wang, Limin Wang, and Yu Qiao. Internvideo-ego4d: A pack of champion solutions to ego4d challenges, 2022. 7
- [6] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 2
- [7] Gedas Cheng, Feng & Bertasius. Tallformer: Temporal action localization with a long-memory transformer. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIV*, pages 503–521. Springer, 2022. 2
- [8] Rui Dai, Srijan Das, and Francois Bremond. Ctrn: Class-temporal relational network for action detection. *arXiv preprint arXiv:2110.13473*, 2021. 2
- [9] Rui Dai, Srijan Das, Kumara Kahatapitiya, Michael S Ryoo, and François Brémond. Ms-tct: multi-scale temporal conv-transformer for action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20041–20051, 2022. 2, 6, 7
- [10] Rui Dai, Srijan Das, Luca Minciullo, Lorenzo Garattoni, Gianpiero Francesca, and François Bremond. Pdan: Pyramid dilated attention network for action detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2970–2979, 2021. 2, 6, 7
- [11] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018. 2, 6
- [12] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Jian Ma, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision (IJCV)*, 130:33–55, 2022. 6, 7
- [13] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6569–6578, 2019. 2
- [14] Lijie Fan, Wenbing Huang, Chuang Gan, Stefano Ermon, Boqing Gong, and Junzhou Huang. End-to-end learning of motion representation for video understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6016–6025, 2018. 1
- [15] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 6, 7
- [16] Chengjian Feng, Yujie Zhong, Yu Gao, Matthew R Scott, and Weilin Huang. Tood: Task-aligned one-stage object detection. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3490–3499. IEEE Computer Society, 2021. 2
- [17] Chuang Gan, Naiyan Wang, Yi Yang, Dit-Yan Yeung, and Alex G Hauptmann. Devnet: A deep event network for multimedia event detection and evidence recounting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2568–2577, 2015. 1
- [18] Rohit Girdhar, Mannat Singh, Nikhila Ravi, Laurens van der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A Single Model for Many Visual Modalities. In *CVPR*, 2022. 6, 7
- [19] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022. 2, 6
- [20] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020. 2
- [21] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 297–304. JMLR Workshop and Conference Proceedings, 2010. 3
- [22] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF con-*

- ference on computer vision and pattern recognition*, pages 9729–9738, 2020. 3
- [24] Kumara Kahatapitiya and Michael S Ryoo. Coarse-fine networks for temporal activity detection in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8385–8394, 2021. 2, 6, 7
- [25] Dahun Kim, Donghyeon Cho, and In So Kweon. Self-supervised video representation learning with space-time cubic puzzles. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 8545–8552, 2019. 1
- [26] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018. 2
- [27] Xiang Li, Wenhai Wang, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss v2: Learning reliable localization quality estimation for dense object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11632–11641, 2021. 2
- [28] Xiang Li, Wenhai Wang, Lijun Wu, Shuo Chen, Xiaolin Hu, Jun Li, Jinhui Tang, and Jian Yang. Generalized focal loss: Learning qualified and distributed bounding boxes for dense object detection. *Advances in Neural Information Processing Systems*, 33:21002–21012, 2020. 2
- [29] Chuming Lin, Chengming Xu, Donghao Luo, Yabiao Wang, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang, and Yanwei Fu. Learning salient boundary feature for anchor-free temporal action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3320–3329, June 2021. 1, 2, 3, 7, 8
- [30] Kevin Qinghong Lin, Alex Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Zhongcong Xu, Difei Gao, Rongcheng Tu, Wenzhe Zhao, Weijie Kong, et al. Egocentric video-language pretraining. *arXiv preprint arXiv:2206.01670*, 2022. 6, 7
- [31] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3889–3898, 2019. 2, 7, 8
- [32] Tianwei Lin, Xu Zhao, and Zheng Shou. Single shot temporal action detection. In *Proceedings of the 25th ACM international conference on Multimedia*, pages 988–996, 2017. 1, 2
- [33] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *European Conference on Computer Vision*, 2018. 2, 7, 8
- [34] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 2, 6
- [35] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 2, 5
- [36] Naiyuan Liu, Xiaohan Wang, Xiaobo Li, Yi Yang, and Yueting Zhuang. Reler@zju-alibaba submission to the ego4d natural language queries challenge 2022, 2022. 3
- [37] Shuming Liu, Mengmeng Xu, Chen Zhao, Xu Zhao, and Bernard Ghanem. Etad: Training action detection end to end on a laptop, 2022. 2
- [38] Xiaolong Liu, Qimeng Wang, Yao Hu, Xu Tang, Shiwei Zhang, Song Bai, and Xiang Bai. End-to-end temporal action detection with transformer. *IEEE Transactions on Image Processing*, 31:5427–5441, 2022. 2, 7, 8
- [39] Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. Gaussian temporal awareness networks for action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 344–353, 2019. 2, 8
- [40] Fangzhou Mu, Sicheng Mo, Gillian Wang, and Yin Li. Where a strong backbone meets strong features – action-former for ego4d moment queries challenge, 2022. 7
- [41] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [42] Tian Pan, Yibing Song, Tianyu Yang, Wenhao Jiang, and Wei Liu. Videomoco: Contrastive video representation learning with temporally adversarial examples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11205–11214, 2021. 3
- [43] AJ Piergiovanni and Michael Ryoo. Temporal gaussian mixture layer for videos. In *International Conference on Machine Learning*, pages 5152–5161. PMLR, 2019. 2
- [44] Zhiwu Qing, Haisheng Su, Weihao Gan, Dongliang Wang, Wei Wu, Xiang Wang, Yu Qiao, Junjie Yan, Changxin Gao, and Nong Sang. Temporal context aggregation network for temporal action proposal refinement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 485–494, 2021. 2, 8
- [45] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks, 2017. 8
- [46] Zhaofan Qiu, Ting Yao, Chong-Wah Ngo, Xinmei Tian, and Tao Mei. Learning spatio-temporal representation with local and global diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12056–12065, 2019. 1
- [47] Jiayi Shao, Xiaohan Wang, and Yi Yang. Reler@zju submission to the ego4d moment queries challenge 2022, 2022. 7
- [48] Dingfeng Shi, Yujie Zhong, Qiong Cao, Jing Zhang, Lin Ma, Jia Li, and Dacheng Tao. React: Temporal action detection with relational queries. In *European conference on computer vision*, 2022. 2, 3, 8
- [49] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-deconvolutional networks for precise temporal action localization in untrimmed videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5734–5743, 2017. 2

- [50] Zheng Shou, Hang Gao, Lei Zhang, Kazuyuki Miyazawa, and Shih-Fu Chang. Autoloc: Weakly-supervised temporal action localization in untrimmed videos. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 154–171, 2018. [3](#)
- [51] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1049–1058, 2016. [2](#)
- [52] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 761–769, 2016. [2](#)
- [53] Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 510–526. Springer, 2016. [2](#), [6](#)
- [54] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos, 2014. [8](#)
- [55] Deepak Sridhar, Niamul Quader, Srikanth Muralidharan, Yaoxin Li, Peng Dai, and Juwei Lu. Class semantics-based attention for action detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13739–13748, 2021. [2](#)
- [56] Haisheng Su, Weihao Gan, Wei Wu, Yu Qiao, and Junjie Yan. Bsn++: Complementary boundary regressor with scale-balanced relation modeling for temporal action proposal generation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 2602–2610, 2021. [2](#)
- [57] Yu-Gang Jiang&Jingen Liu&A Roshan Zamir&George Toderici&Ivan Laptev&Mubarak Shah& Rahul Sukthankar. Thumos challenge: Action recognition with a large number of classes. 2014. [2](#), [6](#)
- [58] Jing Tan, Jiaqi Tang, Limin Wang, and Gangshan Wu. Relaxed transformer decoders for direct action proposal generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13526–13535, October 2021. [2](#), [7](#), [8](#)
- [59] Jing Tan, Xiaotong Zhao, Xintian Shi, Bin Kang, and Limin Wang. Pointtad: Multi-label temporal action detection with learnable query points. In *Advances in Neural Information Processing Systems*. [2](#), [6](#), [7](#)
- [60] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9627–9636, 2019. [2](#), [4](#), [6](#)
- [61] Praveen Tirupattur, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. Modeling multi-label action dependencies for temporal action localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1460–1470, 2021. [2](#), [6](#), [7](#)
- [62] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. VideoMAE: Masked autoencoders are data-efficient learners for self-supervised video pre-training. In *Advances in Neural Information Processing Systems*, 2022. [7](#)
- [63] Heng Wang, Dan Oneata, Jakob Verbeek, and Cordelia Schmid. A robust and efficient video representation for action recognition. *International journal of computer vision*, 119:219–238, 2016. [1](#)
- [64] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European conference on computer vision*, pages 20–36. Springer, 2016. [8](#)
- [65] Qiang Wang, Yanhao Zhang, Yun Zheng, and Pan Pan. Rcl: Recurrent continuous localization for temporal action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13566–13575, 2022. [2](#)
- [66] Xiang Wang, Zhiwu Qing, Ziyuan Huang, Yutong Feng, Shiwei Zhang, Jianwen Jiang, Mingqian Tang, Changxin Gao, and Nong Sang. Proposal relation network for temporal action detection. 2021. [2](#)
- [67] Xiaohan Wang, Linchao Zhu, Zhedong Zheng, Mingliang Xu, and Yi Yang. Align and tell: Boosting text-video retrieval with local alignment and fine-grained supervision. *IEEE Transactions on Multimedia*, 2022. [3](#)
- [68] Yuanjun Xiong, Limin Wang, Zhe Wang, Bowen Zhang, Hang Song, Wei Li, Dahua Lin, Yu Qiao, Luc Van Gool, and Xiaoou Tang. Cuhk & ethz & siat submission to activitynet challenge 2016, 2016. [7](#)
- [69] Huijuan Xu, Abir Das, and Kate Saenko. R-c3d: Region convolutional 3d network for temporal activity detection. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017. [2](#)
- [70] Mengmeng Xu, Juan-Manuel Pérez-Rúa, Victor Escorcía, Brais Martínez, Xiatian Zhu, Li Zhang, Bernard Ghanem, and Tao Xiang. Boundary-sensitive pre-training for temporal localization in videos. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7220–7230, 2021. [2](#)
- [71] Mengmeng Xu, Juan Manuel Perez Rua, Xiatian Zhu, Bernard Ghanem, and Brais Martínez. Low-fidelity video encoder optimization for temporal action localization. *Advances in Neural Information Processing Systems*, 34:9923–9935, 2021. [2](#)
- [72] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [1](#), [2](#), [7](#), [8](#)
- [73] Ze Yang, Shaohui Liu, Han Hu, Liwei Wang, and Stephen Lin. Reppoints: Point set representation for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9657–9666, 2019. [2](#)
- [74] Serena Yeung, Olga Russakovsky, Ning Jin, Mykhaylo Andriluka, Greg Mori, and Li Fei-Fei. Every moment counts: Dense detailed labeling of actions in complex videos. *International Journal of Computer Vision*, 126:375–389, 2018. [2](#), [6](#)

- [75] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *ICCV*, 2019. [1](#), [2](#), [7](#), [8](#)
- [76] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *European Conference on Computer Vision*, volume 13664 of *LNCS*, pages 492–510, 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [77] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9759–9768, 2020. [2](#)
- [78] Chen Zhao, Mery Ramazanova, Mengmeng Xu, and Bernard Ghanem. Segtad: Precise temporal action detection via semantic segmentation, 2022. [2](#)
- [79] Chen Zhao, Ali Thabet, and Bernard Ghanem. Video self-stitching graph network for temporal action localization. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 13638–13647, 2021. [1](#), [2](#), [7](#), [8](#)
- [80] Peisen Zhao, Lingxi Xie, Chen Ju, Ya Zhang, Yanfeng Wang, and Qi Tian. Bottom-up temporal action localization with mutual regularization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16*, pages 539–555. Springer, 2020. [2](#)
- [81] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaoou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, 2017. [2](#), [7](#), [8](#)
- [82] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-iou loss: Faster and better learning for bounding box regression. In *The AAAI Conference on Artificial Intelligence (AAAI)*, 2020. [5](#)
- [83] Benjin Zhu, Jianfeng Wang, Zhengkai Jiang, Fuhang Zong, Songtao Liu, Zeming Li, and Jian Sun. Autoassign: Differentiable label assignment for dense object detection. *arXiv preprint arXiv:2007.03496*, 2020. [2](#)
- [84] Chenchen Zhu, Yihui He, and Marios Savvides. Feature selective anchor-free module for single-shot object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 840–849, 2019. [2](#)
- [85] Zixin Zhu, Wei Tang, Le Wang, Nanning Zheng, and G. Hua. Enriching local and global contexts for temporal action localization. In *ICCV*, 2021. [1](#), [2](#), [7](#), [8](#)