ActionSwitch: Class-agnostic Detection of Simultaneous Actions in Streaming Videos

Hyolim Kang[®], Jeongseok Hyun[®], Joungbin An[®], Youngjae Yu[®], and Seon Joo Kim[®]

Yonsei University

Abstract. Online Temporal Action Localization (On-TAL) is a critical task that aims to instantaneously identify action instances in untrimmed streaming videos as soon as an action concludes—a major leap from frame-based Online Action Detection (OAD). Yet, the challenge of detecting overlapping actions is often overlooked even though it is a common scenario in streaming videos. Current methods that can address concurrent actions depend heavily on class information, limiting their flexibility. This paper introduces ActionSwitch, the first class-agnostic On-TAL framework capable of detecting overlapping actions. By obviating the reliance on class information, ActionSwitch provides wider applicability to various situations, including overlapping actions of the same class or scenarios where class information is unavailable. This approach is complemented by the proposed "conservativeness loss", which directly embeds a conservative decision-making principle into the loss function for On-TAL. Our ActionSwitch achieves state-of-the-art performance in complex datasets, including Epic-Kitchens 100 targeting the challenging egocentric view and FineAction consisting of fine-grained actions.

Keywords: Online Video Understanding \cdot Class-agnostic Detection

1 Introduction

As the demand for autonomous driving systems [22], robotics applications [27], and ego-centric perception [9] continues to grow, the importance of online video understanding tasks has become more pronounced. Online Action Detection (OAD) [10] is a pivotal online video understanding task that has gained considerable attention [1,13,15,24,45,48,49], as it aims to detect the action class of each frame from a streaming video. However, merely assigning class scores to individual frames does not capture the essence of actions in the video, which inherently extend over multiple frames and form *action instances*. In particular, this instance-level understanding is a de facto standard in conventional video understanding tasks, including Temporal Action Localization (TAL) [28,31,37,47,53].

To provide instance-level understanding in streaming videos, Online Temporal Action Localization (On-TAL) [20, 23, 25, 41] has recently been proposed. Its goal is to instantaneously identify action instances in untrimmed streaming video, accurately determining the start and end times, and classifying each instance as soon as it concludes.



Fig. 1: Overview of the ActionSwitch Framework: State label is derived from the *sum* of the ids of activated switches. For example, the state is labeled as '3' between t2 and t3 when switches '1' and '2' are simultaneously active, whereas it registers as '2' from t3 to t4 when only switch '2' is active. State changes signify action instance boundaries, and our 'conservativeness loss' minimizes state fluctuations to improve detection accuracy.

While the use of (soft-) non-maximum suppression (NMS) [3] is indispensable in conventional TAL, the online constraints of On-TAL preclude any retrospective modifications after the initial generation of action instances. This requisite, coupled with the need for instant detection, poses challenges in directly applying TAL methods to On-TAL.

Thus, a straightforward approach is extending the OAD framework; this involves thresholding the per-frame output from the OAD model and aggregating these outputs to generate action instances. Nonetheless, a simple threshold-based grouping of either a class-agnostic or class-aware OAD model gives rise to two significant challenges: i) the difficulty in detecting actions occurring simultaneously and ii) the generation of fragmented and noisy action instances [20]. The main objective of this work is to explore these challenges and offer effective solutions.

Detecting overlapping action instances in streaming videos, despite their frequent occurrence, remains relatively underexplored. This aspect holds considerable importance as it harmonizes with the intuition that actions are intricately connected [51]. In particular, an extension of a class-agnostic OAD model [20] is entirely incapable of recognizing concurrent actions. While a classaware model [41] may attempt to do so, it faces considerable limitations. This is largely because class-aware models rely on independently grouping per-class scores to generate action instances. As a result, their ability to detect overlapping instances is intricately tied to the predefined action classes in a dataset. The Epic-Kitchens dataset [9], with its many overlapping instances within the same class, highlights this limitation that class-aware OAD models struggle to address. Additionally, these models are highly sensitive to predetermined thresholds, a challenge that intensifies with an increase in the number of classes. Moreover, the dependence on class information to distinguish action instances presents difficulties in general On-TAL applications, especially when predefining all possible action classes is not practical. The unique advantage of disentangling class information from action proposal generation is further elaborated in Sec. 2.3.

To this end, we propose the first class-agnostic On-TAL framework capable of detecting overlapping action instances. We begin by considering a machine with multiple switches (Fig. 4), namely **ActionSwitch**. Each switch performs class-agnostic yet mutually exclusive action instance detection where the activated switch indicates detecting ongoing action. By incorporating multiple switches, our system adeptly detects overlapping instances all without necessitating class information. We devise a finite state machine corresponding to our ActionSwitch framework to instantiate this concept. Leveraging a conventional OAD framework that outputs state labels for each frame, we create what we refer to as a state-emitting OAD. We can effectively generate action instances online by appropriately grouping the framewise output of the state-emitting OAD model, even if there are overlaps among them.

Another remaining challenge is noisy and fragmented action proposal generation, a byproduct of the discrepancy between the frame-centric nature of OAD and the instance-level demands of On-TAL. From an instance-level perspective, even a single erroneous decision can shatter the continuity of an action instance, which is particularly problematic when identifying long action instances. Consequently, the agent must be conservative in altering its decisions. Past approaches [20, 41] learned this principle from data by modeling decision context, yet we advocate for a more direct method: directly infuse conservatism into the loss function. We propose a **Conservativeness loss**, an auxiliary loss term that encourages the agent to depend on its previous decision. This method integrates conservatism into the standard loss function without necessitating complex architectural modifications.

We conduct comprehensive experiments on three primary action localization datasets [9, 19, 32] to validate our proposed method's efficacy and establish robust baselines for future research of On-TAL. It is also worth noting that while the majority of action localization datasets (including the above three) require the ActionSwitch framework with just two switches, three or more switch configurations are also tested in MultiTHUMOS [51] dataset, providing valuable insight for further research.

Our contributions can be summarized as follows:

- We introduce ActionSwitch, the first class-agnostic On-TAL framework that is capable of detecting overlapping action instances by incorporating a finite state machine concept.
- We introduce *Conservativeness loss* to address the challenge of noisy and fragmented action proposals, effectively incorporating conservatism directly into the loss function with minimal modification.

- 4 Kang. et al.
- We demonstrate the effectiveness of ActionSwitch through experimental results from multiple action localization datasets [9,19,30,32,51] and extensive ablation studies.

2 Related Work

2.1 Streaming Video Understanding

Online Action Detection (OAD) [10] is a well-established task in online video understanding that requires identifying the action class of the current input frame in a streaming setting. A plethora of work [1, 6, 11, 45, 48-50, 55] has been introduced, mainly focusing on temporal modeling of the past visual context. Despite the progress, the frame-centric approach of OAD and its evaluation with per-frame mean Average Precision (mAP) falls short of addressing the needs of real-world streaming video analysis that necessitates instance-level recognition.

Online Detection of Action Start (ODAS) [14, 36], on the other hand, is a task that pinpoints the initial timestep of each action within a streaming video. By retrieving one unique timestep for each action instance, ODAS implicitly conveys the concept of action instances. However, its concentration solely on the initiation point of action instance restricts its broader applicability.

2.2 Online Temporal Action Localization

Online Temporal Action Localization (On-TAL) targets real-time identification of action instances in streaming video. A direct method involves extending the OAD framework [20, 41] for On-TAL, which necessitates incorporating the agent's decision history to produce accurate action instances. Several approaches have been devised to tackle this challenge, including the incorporation of a unique grouping module complemented by a distinct training strategy [20], as well as the implementation of a decision context token [41]. The exact processes for generating action instances in these methods are detailed further in our supplementary materials to make the paper self-contained.

On the other hand, the TAL extension has been adapted for On-TAL using data-driven online filtering [23], which mimics the online version of NMS. Yet, all these methods rely on class-specific information to separate concurrent actions. For example, SimOn [41] groups the same class decisions to generate action instances whereas OAT [23] uses a handcrafted threshold value to explicitly suppress overlapping action instances of the same class, indicating that neither can handle overlapping action instances that have the same action class. In contrast, our ActionSwitch framework is the first On-TAL approach to identify overlapping actions irrespective of class information.

2.3 Class-agnostic Detection

Class-agnostic proposal generation is already a widely accepted standard in both Object Detection (OD) [8, 16, 40] and TAL [28, 47, 52] literature, and recent

 $\mathbf{5}$

works highlight the unique benefits of disentangling the proposal generation and classification. OLN [21] suggests that focusing on localization without class constraints enhances generalization, and Maaz et al. [33] show benefits of classagnostic detector for open-world detection and self-supervised learning [2]. In videos, recognizing class-agnostic event boundaries, as discussed in [35], aligns with human perception, which does not depend on predefined action classes. These recent advances [18, 44] and flourishing development of video-language models [34] validate our approach to separate action proposal generation from classification in On-TAL, indicating the potential for a more flexible framework suited to open-world and open-vocabulary contexts.

3 Methodology

3.1 Problem Setting

On-TAL Let us consider an untrimmed video $V = \{x_{\tau}\}_{\tau=1}^{T}$ with M action instances $\Psi = \{\psi_m\}_{m=1}^{M} = \{(t_m^s, t_m^e, c_m)\}_{m=1}^{M}$ is given in a streaming format. x_{τ} indicates τ th frame, t_m^s , t_m^e represent the start and the end timestep, and c_m is the class label of the mth action instance ψ_m . Following the previous convention [20, 28, 45, 47], consecutive k frames are converted to a D dimensional visual feature sequence $f \in \mathbb{R}^{\lfloor \frac{T}{k} \rfloor \times D}$ with a pretrained snippet encoder. Subsequent operations, including online action instance generation, are performed on this feature sequence. The goal of On-TAL task is to generate and accumulate action proposals ψ as soon as their completion is detected, aiming to reconstruct Ψ without retrospective modification to each ψ .

OAD-extension vs TAL-extension While On-TAL requires immediate identification of action endpoints, it allows flexible timing in identifying action startpoints. This flexibility is exploited by the TAL-based approach [23]; it identifies action starts at the time of action instance generation, thereby benefiting from rich context. In contrast, OAD-based methods [20, 41] are inherently designed to recognize the start of actions in real-time, often working within a highly constrained initial context. This distinction emphasizes the unique role of OAD-based methods in On-TAL. They are able to adeptly handle early action detection, a task TAL-based methods cannot perform. In particular, promptly identifying an action's start can extend the On-TAL model's utility to scenarios where immediate detection of an action's initiation is critical, an aspect tackled in Online Detection of Action Start (ODAS) [36].

3.2 State-emitting OAD Model

In order to address the class-agnostic On-TAL problem, we begin by extending class-agnostic OAD models, following the baseline approach described in [20]. In this setting, the OAD model produces a framewise actionness decision; 0 indicates the absence of action, and 1 the presence of action. Action start and end can be instantaneously determined by detecting changes from 0 to 1 and 1 to





Fig. 2: (a) State diagram of ActionSwitch framework. Some connections are omitted for simplicity. (b) Overall architecture of state-emitting OAD model.

0 respectively. Additionally, the progress status of an action, which is helpful for many practical applications, can be easily calculated by measuring the distance between the action's start timestep and the current timestep.

The primary challenge with this approach is its inability to detect overlapping action instances, as decisions are based solely on the presence of the action in the current frame. Aggregation of these binary action decisions tends to merge multiple overlapping instances into a single instance. A possible solution is to use multiple OAD models for mutually exclusive action instance detection. For instance, in a setup with two OAD models and a single action, one model should signal "no action" when the other detects the action. However, this interdependence of the models' decisions makes the implementation complicated.

Instead, we abstract the concept into a *single* machine with multiple switches (Fig. 2 (a)). A two-switch configuration is exemplified for simplicity, although configurations with three or more switches are also feasible (See Section 4.5). Switch 1 is activated when the first action is detected, and Switch 2 is activated when the other action is detected while Switch 1 is still active. The finite state machine corresponding to this machine has four states: i) no switch activated, ii) switch 1 activated, iii) switch 2 activated, and iv) both switches activated. These states are illustrated in the State Table in Fig. 2 (a). State transitions here indicate the commencement or termination of a certain action instance, enabling the model to determine the start, end, or progress status of multiple ongoing actions in real-time by analyzing the stored state history.

In order to realize this concept using neural networks, we need to make an important design choice regarding the interpretation of the network's output. For simplicity, we directly interpret the output of the network as the state label, yielding a state-emitting OAD model (Fig. 2 (b)). At each timestep $t, f_t \in \mathbb{R}^D$ is fed into the uni-directional sequential encoder, producing hidden state $g_t \in \mathbb{R}^D$. Subsequently, the categorical probability distribution p_t is generated by $p_t = softmax(SC(g_t))$, where SC denotes a State Classifier SC : $\mathbb{R}^D \to \mathbb{R}^S$ that can be any arbitrary neural network, with S representing the number of



Fig. 3: Training process in ActionSwitch. CE and \mathcal{L}_c denote the terms in Eq. 2. GT state and conservative pseudo-state are used for training. GT states are encoded from GT action instances while the pseudo-states come from the model's own predictions. At t6, action 1 is temporarily lost and results in the fragmentation of the action instance. However, with our conservativeness loss, the output of action instances becomes robust against such fragmentation (t6) and noisy output (t11).

```
Algorithm 1 Conservativeness loss in a PyTorch style.
```

```
# logits: tensor in the shape of (B, L, n_state)
# return loss penalized by context change (cc)
pred_state = torch.argmax(logits, dim=2) # (B, L)
cc_mask = pred_state[:, 1:] != pred_state[:, :-1]
cc_targets = pred_state[:, :-1][cc_mask]
cc_logits = logits[:, 1:][cc_mask]
return F.cross_entropy(cc_logits, cc_targets)
```

states. The state corresponding to the current timestep t is obtained by simply calculating $s_t = \operatorname{argmax}(p_t)$. Note that argmax operation here eliminates the need for handcrafted threshold selection, which was the primary problem of previous works [23, 41].

3.3 Conservativeness Loss

Action boundaries are typically much less frequent than non-boundary frames, which is a widely observed phenomenon in most videos. It is evident that leveraging this prior is advantageous in grouping framewise OAD decisions in generating action instances. Previous methods [20,41] attempted to learn this prior by modeling decision history. Our approach, on the other hand, involves encoding it directly into the loss function, thereby leading to the novel conservativeness loss term \mathcal{L}_c :

$$\mathcal{L}_c(p_t, s_{t-1}) = \begin{cases} -log(p_t[s_{t-1}]), & \text{if } \operatorname{argmax}(p_t) \neq s_{t-1}; \\ 0, & \text{otherwise.} \end{cases}$$
(1)

This loss term penalizes state change by applying standard cross-entropy loss with a pseudo-label s_{t-1} , an output from the model's own prediction in the

previous step. Note that this loss term is only imposed at the context-changing timestep, as shown in Fig. 3. A significant advantage of this loss term is its simplicity; it integrates seamlessly into the standard OAD framework for On-TAL extension without necessitating any additional modules or architectural modifications. Furthermore, the loss term can be easily implemented with just five lines of code (Algorithm 1), making it a practical and efficient solution.

With a ground truth state label y_t , final loss term \mathcal{L} for each timestep t is defined as follows:

$$\mathcal{L}(p_t, y_t, s_{t-1}) = CE(p_t, y_t) + \alpha \mathcal{L}_c(p_t, s_{t-1}), \tag{2}$$

where CE refers to the typical cross entropy term and α denotes a weight to balance both losses. For training, we generate a framewise ground truth state label from the ground truth action instances (Encode Action in Fig. 3), assuming that switch 1 is first activated and then switch 2 to avoid ambiguity.

During the inference stage, we store s_t for each timestep t in a history queue. By comparing s_t and s_{t-1} , we can instantly infer the start and end timesteps. The history queue can then be decoded into action instances using a straightforward algorithm (Decode State in Fig. 2 (b) and 3), which we elaborate on in our supplementary material. This approach ensures "no room for boundary mismatches", as the predicted state label sequence corresponds to one clear action scenario, and hence does not need a separate boundary-matching module. It significantly differentiates our approach from other boundary-matching based algorithms [28, 29], which involve exhaustive matching between boundaries and filtering processes.

4 Experiments

In this section, we present the experimental results and ablation studies of our proposed method, mainly focusing on the On-TAL performance. Moreover, we include results from the ODAS benchmark for a comprehensive evaluation.

4.1 Datasets and Features

Datasets To validate the ActionSwitch framework's efficacy and general applicability, we evaluate our method on the variety of popular TAL datasets, including standard THUMOS14 [19], large-scale FineAction [32], and large egocentric Epic-Kitchens 100 [9]. Additional experimental results on the MUSES [30] dataset and a detailed analysis of each dataset in terms of appropriately evaluating On-TAL are provided in the supplementary materials.

Features For the THUMOS14 dataset, we adopt the two-stream TSN model [43] trained on Kinetics [5] to extract features, following [20, 23, 41, 47]. For the MUSES dataset, we use the officially available features for our experiments. However, in the FineAction dataset, the provided features are too coarse-grained (16 frames per one feature vector), which results in an insuperable bottleneck

9

for detecting fine-grained action instances. To address this issue, we follow [1] and increase the temporal resolution of the features by a factor of four. For the Epic-Kitchens 100 dataset, we used publicly available¹ Kinetics400 pretrained Slowfast [12] network weight with stride four to extract features. It is important to note that all experiments in the same table are conducted with the *same features*, whether for offline or online TAL methods to ensure a fair comparison.

4.2 Evaluation Metric

F1 score for class-agnostic evaluation In various TAL literatures [28, 29, 37, 39, 54], Average Recall (AR) under varying Intersection over Union (IoU) thresholds is used to evaluate class-agnostic action proposal generation. This metric is suited to standard TAL, where the output from proposal generators undergoes NMS-like processing to eliminate redundant detection. However in the On-TAL context, as discussed in Section 3.1, such retrospective modifications to generated proposals are not allowed, making it essential to assess both the precision and recall.

Therefore, we use the **F1 score** as the main evaluation metric in the classagnostic On-TAL to capture the balance between precision and recall. We first run the Hungarian algorithm [26] to provide optimal bipartite matching between the ground truth action instances and predictions based on their temporal IoU (tIoU). Then, a prediction is considered a true positive based on whether it surpasses a certain tIoU threshold. Note that using Hungarian matching here is in the same context as using it in popular query-based object detection methods [4]. For the sake of concise demonstration, F1 and recall are mainly presented at a setting of tIOU=0.5. F1 and recall in other tIOU thresholds, straightforward pseudocode of the metric calculation, and interesting discussions about the F1 score and its relationship with standard classwise mAP metric can be found in the supplementary materials.

mAP metric for class-dependent evaluation To provide an apple-to-apple comparison among previous works, we also include the standard mAP metric, which has been widely used in both offline-TAL [47,53] and online-TAL [20,23] literature. We report mAPs with varying tIOUs in a set $\{0.3, 0.4, 0.5, 0.6, 0.7\}$ for THUMOS14 and $\{0.5, 0.75, 0.95\}$ for FineAction. An average value of those mAPs with multiple tIOUs is also reported for succinct comparison. For Epic-Kitchens 100 dataset, we evaluate mAP@0.5 using "noun" annotations, as they present a greater challenge compared to its verb annotations [53]. To evaluate the performance of ODAS, we measure the point-level average precision (p-AP) and calculate p-mAP by averaging p-AP over all action classes, following previous works [14, 20, 36]. Given that the ODAS task aims to achieve fast detection of action starts, we restrict the offset value to 3 seconds, as late detection beyond this threshold is considered not useful.

¹ https://pytorchvideo.org/

4.3 Implementation Details

Architecture Inspired by the latest research [1], a stacked GRU [7] and MLPs with residual connection are used for the unidirectional sequential encoder and a state classifier respectively. This efficient recurrent design enables our model to process over 500 fps, underscoring its suitability for online applications. Exploring effective architectural design choices is a promising direction, but we stick to the minimalist design in this paper as it is not our primary focus.

Class label While the ActionSwitch framework focuses on class-agnostic action instances and our main F1 metric does not require class and confidence scores for evaluation, the class labels and corresponding confidence scores are necessary to measure per class mAP. Hence, we train an extra classifier that takes feature sequences as its input and predicts the class labels for the given sequences. We take a vanilla transformer [42] as its architecture, and the max class logit is directly treated as the confidence score. To ensure a fair comparison, we apply this same classifier to previous class-agnostic methods, resulting in slightly better performance than reported in the original paper.

4.4 Main Results

Tab. 1, Tab. 2 and Tab. 3 present the main experimental results of our Action-Switch framework. For clarity, the performance of several state-of-the-art offline TAL methods is also reported. Here, the standard mAP metric is observed to have a stronger correlation with recall than precision (see supplementary materials for further discussion). Therefore in offline TAL methods, while using all action proposals without filtering boosts mAP scores, it becomes crucial to eliminate low-confidence proposals to achieve a reasonable F1 score.

While there is still a performance gap compared to the state-of-the-art offline TAL method [53], ActionSwitch sets a new standard for both mAP and F1 scores among OAD-extended On-TAL methods. Unlike CAG-QIL [20] which requires two-stage training and SimOn [41] which depends on an On-TAL specific architecture, ActionSwitch offers a straightforward extension from the OAD

Table 1: Comparison to other TAL methods on FineAction dataset [32]. We report mAP at different tIoU thresholds and average mAP in [0.5:0.05:0.95].[†] indicates the results with the optimal threshold that drops the low-confidence proposals and achieves the best F1 score. For each metric, the best is bolded.

	Setting	Method	F1	Recall	mAP@0.5	0.75	0.95	Avg.
	Offline TAL	ActionFormer [53] ActionFormer [†] [53]	3.62 30.14	62.83 25.25	21.21 14.13	11.02 9.05	1.68 1.57	11.74 9.16
Online TAL	TAL-Extension	OAT (downsampled by 4) [23] OAT [23]	18.64 7.21	38.10 8.14	8.88 1.19	1.56 0.12	0.02 0.00	3.04 0.34
	OAD-Extension	CAG-QIL [20] SimOn [41] ActionSwitch (Ours)	15.67 10.53 19.44	12.73 18.06 21.76	8.00 7.95 10.58	4.07 3.53 4.71	1.05 0.98 0.64	4.45 4.12 5.36

Table 2: Comparison on Epic-Kitchens 100 dataset [9]. # proposal denotes the number of generated proposals of each method, and # ground truth refers to the number of ground truth proposals. We report class-specific mAP at tiou = 0.5 using "noun" annotations. SimOn [41] and OAT [23] exploit class information for action instance generation, in contrast to OAD-Grouping [20], CAG-QIL [20], and our method, which do not. For each metric, the best is bolded.

Setting	Method	F1	precision	recall	mAP@0.5	# proposal	# ground truth
TAL-extension	OAT [23]	27.583	17.595	63.798	3.296	35054	
	[CAG-QIL [20]	$2\bar{3}.\bar{1}1\bar{7}$	21.347	25.206	2.442	11416	
OAD automaion	SimOn [41]	4.395	2.351	33.481	1.846	137685	9668
OAD-extension	OAD-Grouping [20]	21.416	25.533	18.442	2.267	6983	
	ActionSwitch (Ours)	32.444	29.858	35.519	3.597	11501	

Table 3: Comparison to other TAL methods on THUMOS14 dataset [19]. We report mAP at different tIoU thresholds and average mAP in [0.3:0.1:0.7]. In the offline TAL results, [†] indicates the results with the optimal threshold that drops the low-confidence proposals and achieves the best F1 score. OAT [23] performs On-TAL within a more flexible constraint than other works [20, 41]. (Section 3.1) For each metric, the best is bolded.

	Setting	Method	F1	Recall	mAP@0.3	0.4	0.5	0.6	0.7	Avg.
		G-TAD [47]	6.4	83.4	58.8	52.2	43.6	33.3	22.9	42.2
	Offline TAL	TadTR [31]	3.5	88.1	68.8	62.7	55.9	45.3	32.3	53.0
		ActionFormer [53]	13.8	94.0	77.5	73.5	66.0	55.4	40.6	62.6
		$G-TAD^{\dagger}$ [47]	51.1	49.9	45.6	40.7	34.4	26.7	18.7	33.2
		$TadTR^{\dagger}$ [31]	68.4	63.0	55.4	51.4	46.5	38.4	28.0	43.9
		ActionFormer [†] [53]	75.5	73.6	66.9	63.2	56.3	47.4	35.1	53.8
	TAL-Extension	OAT [23]	62.9	70.3	64.1	57.4	47.8	36.7	20.3	45.3
Online		CAG-QIL [20]	45.8	50.4	48.8	40.9	33.6	24.9	17.3	33.1
TAL	OAD-Extension	SimOn [41]	28.7	52.0	52.2	43.6	32.3	22.5	14.2	33.0
		ActionSwitch (Ours)	53.2	60.1	57.2	50.8	41.7	30.5	21.3	40.3

framework to On-TAL with a minor modification of the output layer and incorporation of additional loss term. Notably as the number of classes grows, SimOn's class-aware grouping strategy significantly falters with the excessive number of proposals and low precision (Tab. 1 and 2), showing the limitations of classwise threshold-based grouping. In contrast, our method shows balanced precision and recall across all the datasets, and the use of the **argmax** operation streamlines the process by eliminating the need to predefine the threshold.

Furthermore, the TAL extension method, OAT [23], achieves impressive performance on relatively small datasets (*e.g.*, THUMOS14) but experiences training instability when applied to large-scale datasets that contain temporally varying action instances. For example, in the FineAction dataset (Tab. 1), OAT [23] performs reasonably well only when the temporal resolution is downsampled, and even in that case, our ActionSwitch performs better.

Also on Epic-Kitchens 100 (Tab. 2), ActionSwitch surpasses OAT in both F1 and mAP scores, as OAT shows a stark imbalance between precision and recall.

Table 5: Ablation study of the number of states in ActionSwitch and the conservativeness loss whose weight is α as shown in Eq. 2.

# switch	Thumos14 [19]					FineAction [32]					
<i>,,</i>	α	F1	Recall	Precision $\#$	proposal	α	F1	Recall	Precision $\#$	proposal	
1	0.000	42.67	59.49	33.26	6006	0.000	14.53	14.71	14.35	24846	
	0.010	49.25	57.77	42.92	4519	0.005	15.78	14.45	17.37	20159	
	0.025	50.41	56.04	45.81	4108	0.010	16.53	13.28	21.90	14694	
	0.050	49.91	49.61	50.21	3318	0.025	9.34	6.28	18.17	8380	
2	0.000	45.87	63.96	35.75	6007	0.000	17.57	22.36	14.46	37469	
	0.010	49.28	61.25	41.23	4989	0.005	18.32	22.00	15.70	33949	
	0.025	53.20	60.10	47.73	4228	0.010	19.44	21.76	17.56	29972	
	0.050	50.20	52.76	47.88	3701	0.025	18.71	16.76	21.15	19205	

This is significant as our ActionSwitch also has the inherent ability to handle ODAS tasks, a functionality not supported by OAT. (Section 3.1)

We present the ODAS performance in Tab. 4, which shows that our method achieves a new stateof-the-art performance. Given that early detection is crucial in ODAS, our method's superior performance at lower offsets highlights its strength in scenarios where immediate action detection is vital. Note that reported mAP is much higher than those in the original papers [20,41] as we utilized Kinetics pre-trained features in all three models.

Table 4: Comparison of other SOTA ODASmethods on THUMOS14 dataset [19].

ODAS Method	Offsets					
ODA5 Method	1	2	3			
CAG-QIL [20]	28.30	42.86	50.12			
SimOn [41]	31.45	46.22	54.11			
ActionSwitch (Ours)	33.06	47.06	54.44			

4.5 Ablation Studies

Number of switches and conservativeness loss Tab. 5 presents comprehensive ablation studies on our ActionSwitch setting. The 1-switch setting with $\alpha = 0$ is considered the baseline for our approach which is a naive OAD grouping extension. By adopting one additional switch, the model can capture overlapping action instances. The recall of the 2-switch setting is much higher than its 1-switch counterpart, supporting the aforementioned claim. This tendency is even more apparent in the FineAction dataset, which contains more overlapping action instances than THUMOS14.

However, the precision remains low due to the grouping error that is prone in OAD extension On-TAL approaches. Here, conservativeness loss significantly alleviates the grouping error. We observe a tradeoff relationship between precision and recall when applying varying alpha values. A higher α value indicates a stronger conservativeness prior, resulting in better precision but lower recall, and vice versa. An adequate α value leads to balanced precision and recall, leading to the best F1 value.

Method	# switch	\mathcal{L}_{c}	F1	Precision	Recall	# proposal	$ \# \text{ GT_proposal} $
	1	0	16.51	42.85	10.22	4816	
	2	0	30.82	34.37	27.94	16405	
Action Switch (Ours)	3	0	32.76	30.04	36.02	24211	
Actionswitch (Ours)	4	0	29.25	22.67	41.21	36682	
	3	x	26.63	20.53	37.91	37287	
	4	х	23.16	16.02	41.86	52758	20186
CAG-QIL $[20]$	-		22.10	48.7	14.29	5926	
SimOn [41]	-		20.59	12.42	60.17	97759	
OAT [23]	-		29.63	18.48	74.74	81626	

Table 6: Multiple switch configurations tested in Multithumos [51] dataset. Ablation study on conservativeness loss term \mathcal{L}_c ($\alpha = 0.025$) is also conducted in multi-switch settings. A good F1 score is observed when the number of generated predictions (# proposal) is close to the number of ground truth instances (# GT_proposal).

Can we include 3 or more switches? Expanding our framework to incorporate three or more action switches is straightforward; it only requires adding additional states corresponding to the combination of switches. Yet, in most of the established TAL benchmark datasets [9, 19, 32], over 99% of the frames have no more than two simultaneous actions, providing almost no ground truth data for scenarios where a third switch would be activated.

The Multithumos dataset [51], however, is characterized by significant overlap among multiple action instances, despite its smaller scale. Therefore, we chose this dataset to perform an ablation study on configurations with three or more switches, with the results presented in Tab. 6. The result not only shows that our method established a new state-of-the-art performance in the Multithumos dataset, but also highlights critical insights. Similar patterns to those in Tab. 5 emerge here, but due to the dataset's dense overlaps, the optimal F1 score occurs



Fig. 4: Qualitative results of On-TAL models. If the action instances are overlapped, they are placed in other lines. We show ground-truth (GT) and the output of 2-state (2S), 4-state (4S), conservative loss (Cons), SimOn [41] and CAG-QIL [20]. Refer to Sec. 4.6 for an analysis of four cases (C1 \sim C4) which are annotated by the red boxes.

with a **3** switch setup. More switches enhance the model's ability to detect additional overlapping actions, as indicated by the improved recall. However, while introducing the conservativeness loss as an auxiliary term can mitigate the issues, it also leads to a precision trade-off. This becomes evident when the model overestimates action instances, as shown by comparing the number of predicted proposals with the actual ground truth counts.

4.6 Qualitative Results

In Fig. 4, we demonstrate the effectiveness of our proposed components. Specifically, C1 illustrates how the conservativeness loss effectively addresses the action fragmentation issue. Additionally, C2 and C3 demonstrate the ability of our ActionSwitch framework to detect overlapping instances. C3 shows that SimOn [41] can detect multiple overlapping actions, but its detection is restricted by the predefined action classes, thus preventing it from detecting overlapping instances of the same action class. In C4, the conservativeness loss leads the model to keep predicting state 1 instead of 3, preventing the excessive proposal generation.

5 Discussion and Conclusion

Towards an Open-Vocabulary Framework Existing On-TAL approaches are evaluated in a closed-vocabulary setting with predefined action classes. However, there is a growing interest in open-vocabulary research [17, 46, 56] as realworld scenarios need deeper comprehension beyond the identification of predefined actions and events in datasets. Class-aware On-TAL methods [23, 41] struggle in this context due to their reliance on preset classes. ActionSwitch, on the other hand, presents a versatile, class-agnostic framework that facilitates integration with video-language models [34]. It paves the path for expansive vocabulary classification, a promising direction for future research.

Is different instantiation possible? Instead of the state-emitting OAD approach, one can treat the output of the OAD model as a state transition, where 0 refers to maintaining the current state, while 1 and 2 indicate state transitions made by manipulating switch 1 and 2, respectively. While this formulation can reduce the required state number, it brings up severe class imbalance due to the sparsity of action boundaries. (Sec. 3.3) It will be promising future research to develop reinforcement learning algorithms [38] with this formulation.

6 Conclusion

We introduce ActionSwitch, the first On-TAL method capable of detecting overlapping action instances without relying on class information. Coupled with novel conservativeness loss, experiments on multiple datasets demonstrate its effectiveness despite its simplicity. We hope that our proposed ActionSwitch can serve as a strong baseline and inspire further On-TAL research. Acknowledgement This work was supported by Institute of Information & communications Technology Planning & evaluation (IITP) grant funded by the Korea government (MSIT) (No.RS-2022-II220113, Developing a Sustainable Collaborative Multi-modal Lifelong Learning Framework), the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF- 2022R1A2C2004509), and Artificial Intelligence Graduate School Program, Yonsei University, under Grant 2020-0-01361

References

- An, J., Kang, H., Han, S.H., Yang, M.H., Kim, S.J.: Miniroad: Minimal rnn framework for online action detection. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2023) 1, 4, 9, 10
- Bar, A., Wang, X., Kantorov, V., Reed, C.J., Herzig, R., Chechik, G., Rohrbach, A., Darrell, T., Globerson, A.: Detreg: Unsupervised pretraining with region priors for object detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 5
- Bodla, N., Singh, B., Chellappa, R., Davis, L.S.: Soft-nms improving object detection with one line of code. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2017) 2
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: Endto-end object detection with transformers. In: European Conference on Computer Vision (ECCV) (2020) 9
- Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 8
- Chen, J., Mittal, G., Yu, Y., Kong, Y., Chen, M.: Gatehub: Gated history unit with background suppression for online action detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 4
- Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. In: NIPS 2014 Workshop on Deep Learning, December 2014 (2014) 10
- Dai, J., Li, Y., He, K., Sun, J.: R-fcn: Object detection via region-based fully convolutional networks. Advances in Neural Information Processing Systems (NeurIPS) (2016) 4
- Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Kazakos, E., Ma, J., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Rescaling egocentric vision. arXiv preprint arXiv:2006.13256 (2020) 1, 2, 3, 4, 8, 11, 13
- De Geest, R., Gavves, E., Ghodrati, A., Li, Z., Snoek, C., Tuytelaars, T.: Online action detection. In: European Conference on Computer Vision (ECCV) (2016) 1, 4
- Eun, H., Moon, J., Park, J., Jung, C., Kim, C.: Learning to discriminate information for online action detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) 4
- Feichtenhofer, C., Fan, H., Malik, J., He, K.: Slowfast networks for video recognition. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
 9
- 13. Gao, J., Yang, Z., Nevatia, R.: Red: Reinforced encoder-decoder networks for action anticipation. The British Machine Vision Conference (BMVC) (2017) 1

- 16 Kang. et al.
- Gao, M., Xu, M., Davis, L.S., Socher, R., Xiong, C.: Startnet: Online detection of action start in untrimmed videos. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2019) 4, 9
- Gao, M., Zhou, Y., Xu, R., Socher, R., Xiong, C.: Woad: Weakly supervised online action detection in untrimmed videos. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 1
- Girshick, R.: Fast r-cnn. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2015) 4
- Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. In: ICLR (2021) 14
- Jaiswal, A., Wu, Y., Natarajan, P., Natarajan, P.: Class-agnostic object detection. In: IEEE Winter Conference on Applications of Computer Vision (WACV) (2021)
 5
- Jiang, Y.G., Liu, J., Roshan Zamir, A., Toderici, G., Laptev, I., Shah, M., Sukthankar, R.: THUMOS challenge: Action recognition with a large number of classes. http://crcv.ucf.edu/THUMOS14/ (2014) 3, 4, 8, 11, 12, 13
- Kang, H., Kim, K., Ko, Y., Kim, S.J.: Cag-qil: Context-aware actionness grouping via q imitation learning for online temporal action localization. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2021) 1, 2, 3, 4, 5, 7, 8, 9, 10, 11, 12, 13
- Kim, D., Lin, T.Y., Angelova, A., Kweon, I.S., Kuo, W.: Learning open-world object proposals without learning to classify. IEEE Robotics and Automation Letters (2022) 5
- Kim, J., Misu, T., Chen, Y.T., Tawari, A., Canny, J.: Grounding human-to-vehicle advice for self-driving vehicles. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 1
- Kim, Y.H., Kang, H., Kim, S.J.: A sliding window scheme for online temporal action localization. In: European Conference on Computer Vision (ECCV) (2022) 1, 4, 5, 7, 8, 9, 10, 11, 13, 14
- Kim, Y.H., Nam, S., Kim, S.J.: Temporally smooth online action detection using cycle-consistent future anticipation. Pattern Recognition (2021) 1
- 25. Kim, Y.H., Nam, S., Kim, S.J.: 2pesnet: Towards online processing of temporal action localization. Pattern Recognition (2022) 1
- Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly (1955) 9
- 27. Lee, Y.: Scaling Robot Learning with Skills. Ph.D. thesis, University of Southern California, Viterbi School of Engineering (2022) 1
- Lin, T., Liu, X., Li, X., Ding, E., Wen, S.: Bmn: Boundary-matching network for temporal action proposal generation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 1, 4, 5, 8, 9
- Lin, T., Zhao, X., Su, H., Wang, C., Yang, M.: Bsn: Boundary sensitive network for temporal action proposal generation. In: European Conference on Computer Vision (ECCV) (2018) 8, 9
- Liu, X., Hu, Y., Bai, S., Ding, F., Bai, X., Torr, P.H.: Multi-shot temporal event localization: a benchmark. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021) 4, 8
- Liu, X., Wang, Q., Hu, Y., Tang, X., Zhang, S., Bai, S., Bai, X.: End-to-end temporal action detection with transformer. IEEE Transactions on Image Processing 31, 5427–5441 (2022) 1, 11

- Liu, Y., Wang, L., Wang, Y., Ma, X., Qiao, Y.: Fineaction: A fine-grained video dataset for temporal action localization. IEEE Transactions on Image Processing 31, 6937–6950 (2022) 3, 4, 8, 10, 12, 13
- Maaz, M., Rasheed, H., Khan, S., Khan, F.S., Anwer, R.M., Yang, M.H.: Classagnostic object detection with multi-modal transformer. In: European Conference on Computer Vision (ECCV) (2022) 5
- Rasheed, H., Khattak, M.U., Maaz, M., Khan, S., Khan, F.S.: Fine-tuned clip models are efficient video learners. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2023) 5, 14
- Shou, M.Z., Lei, S.W., Wang, W., Ghadiyaram, D., Feiszli, M.: Generic event boundary detection: A benchmark for event segmentation. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2021) 5
- Shou, Z., Pan, J., Chan, J., Miyazawa, K., Mansour, H., Vetro, A., Giro-i Nieto, X., Chang, S.F.: Online detection of action start in untrimmed, streaming videos. In: European Conference on Computer Vision (ECCV) (2018) 4, 5, 9
- 37. Su, H., Gan, W., Wu, W., Qiao, Y., Yan, J.: Bsn++: Complementary boundary regressor with scale-balanced relation modeling for temporal action proposal generation. In: Association for the Advancement of Artificial Intelligence (AAAI) (2021) 1, 9
- 38. Sutton, R.S., Barto, A.G.: Reinforcement Learning: An Introduction (2018) 14
- Tan, J., Tang, J., Wang, L., Wu, G.: Relaxed transformer decoders for direct action proposal generation. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2021) 9
- Tan, M., Pang, R., Le, Q.V.: Efficientdet: Scalable and efficient object detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) 4
- Tang, T.N., Park, J., Kim, K., Sohn, K.: Simon: A simple framework for online temporal action localization. arXiv preprint arXiv:2211.04905 (2022) 1, 2, 3, 4, 5, 7, 8, 10, 11, 12, 13, 14
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems (NeurIPS) (2017) 10
- 43. 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 (ECCV) (2016) 8
- 44. Wang, W., Feiszli, M., Wang, H., Malik, J., Tran, D.: Open-world instance segmentation: Exploiting pseudo ground truth from learned pairwise affinity. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 5
- Wang, X., Zhang, S., Qing, Z., Shao, Y., Zuo, Z., Gao, C., Sang, N.: Oadtr: Online action detection with transformers. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2021) 1, 4, 5
- 46. Wu, X., Zhu, F., Zhao, R., Li, H.: Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7031–7040 (2023) 14
- 47. Xu, M., Zhao, C., Rojas, D.S., Thabet, A., Ghanem, B.: G-tad: Sub-graph localization for temporal action detection. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020) 1, 4, 5, 8, 9, 11
- Xu, M., Gao, M., Chen, Y.T., Davis, L.S., Crandall, D.J.: Temporal recurrent networks for online action detection. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2019) 1, 4

- 18 Kang. et al.
- Xu, M., Xiong, Y., Chen, H., Li, X., Xia, W., Tu, Z., Soatto, S.: Long short-term transformer for online action detection. Advances in Neural Information Processing Systems (NeurIPS) (2021) 1, 4
- Yang, L., Han, J., Zhang, D.: Colar: Effective and efficient online action detection by consulting exemplars. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2022) 4
- Yeung, S., Russakovsky, O., Jin, N., Andriluka, M., Mori, G., Fei-Fei, L.: Every moment counts: Dense detailed labeling of actions in complex videos. International Journal of Computer Vision (IJCV) (2018) 2, 3, 4, 13
- Zeng, R., Huang, W., Tan, M., Rong, Y., Zhao, P., Huang, J., Gan, C.: Graph convolutional networks for temporal action localization. In: IEEE/CVF International Conference on Computer Vision (ICCV) (2019) 4
- Zhang, C.L., Wu, J., Li, Y.: Actionformer: Localizing moments of actions with transformers. In: European Conference on Computer Vision (ECCV) (2022) 1, 9, 10, 11
- Zhao, P., Xie, L., Ju, C., Zhang, Y., Wang, Y., Tian, Q.: Bottom-up temporal action localization with mutual regularization. In: European Conference on Computer Vision (ECCV) 9
- Zhao, Y., Krähenbühl, P.: Real-time online video detection with temporal smoothing transformers. In: European Conference on Computer Vision (ECCV) (2022) 4
- Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twentythousand classes using image-level supervision. In: European Conference on Computer Vision (ECCV). pp. 350–368. Springer (2022) 14