

# Decomposed Cross-modal Distillation for RGB-based Temporal Action Detection

Pilhyeon Lee<sup>1</sup> Taeoh Kim<sup>2</sup> Minhoo Shim<sup>2</sup> Dongyoon Wee<sup>2</sup> Hyeran Byun<sup>1\*</sup>

<sup>1</sup>Yonsei University <sup>2</sup>Naver Cloud, AI Tech.

✉ lph1114@yonsei.ac.kr

## Abstract

Temporal action detection aims to predict the time intervals and the classes of action instances in the video. Despite the promising performance, existing two-stream models exhibit slow inference speed due to their reliance on computationally expensive optical flow. In this paper, we introduce a decomposed cross-modal distillation framework to build a strong RGB-based detector by transferring knowledge of the motion modality. Specifically, instead of direct distillation, we propose to separately learn RGB and motion representations, which are in turn combined to perform action localization. The dual-branch design and the asymmetric training objectives enable effective motion knowledge transfer while preserving RGB information intact. In addition, we introduce a local attentive fusion to better exploit the multimodal complementarity. It is designed to preserve the local discriminability of the features that is important for action localization. Extensive experiments on the benchmarks verify the effectiveness of the proposed method in enhancing RGB-based action detectors. Notably, our framework is agnostic to backbones and detection heads, bringing consistent gains across different model combinations.

## 1. Introduction

With the popularization of mobile devices, a significant number of videos are generated, uploaded, and shared every single day through various online platforms such as YouTube and TikTok. Accordingly, there arises the importance of automatically analyzing untrimmed videos. As one of the major tasks, temporal action detection (or localization) [56] has attracted much attention, whose goal is to find the time intervals of action instances in the given video. In recent years, a lot of efforts have been devoted to improving the action detection performance [28–30, 36, 37, 74, 79, 84].

Most existing action detectors take as input two-stream data consisting of RGB frames and motion cues, *e.g.*, optical flow [21, 66, 78]. Indeed, it is widely known that differ-

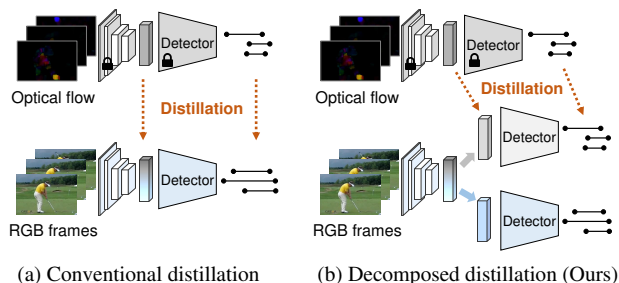


Figure 1. Comparison between conventional distillation and ours.

Framework	Method	Average mAP (%)		
		RGB+OF	RGB	$\Delta$
Anchor-based	G-TAD [74]	41.5	26.9	-14.6
	AFSD [34]	52.4	43.3	-9.1
Anchor-free	Actionformer [80]	62.2	55.5	-6.7
DETR-like	TadTR [42]	56.7	46.0	-10.7
Proposal-free	TAGS [47]	52.8	47.9	-4.9

Table 1. Impact of motion modality. We measure the average mAP under the IoU thresholds of [0.3:0.7:0.1] on THUMOS’14.

ent modalities provide complementary information [6, 24, 58, 69]. To examine how much two-stream action detectors rely on the motion modality, we conduct an ablative study using a set of representative models<sup>1</sup>. As shown in Table 1, regardless of the framework types, all the models experience sharp performance drops when the motion modality is absent, probably due to the static bias of video models [11, 27, 31, 32]. This indicates that explicit motion cues are essential for accurate action detection.

However, two-stream action detectors impose a cycle of dilemmas for real-world applications due to the heavy computational cost of motion modality. For instance, the most popular form of motion cues for action detection, TV- $L^1$  optical flow [66], is not real-time, taking 1.8 minutes to process a 1-min  $224 \times 224$  video of 30 fps on a single GPU [58]. Although cheaper motion clues such as temporal gradient [63, 70, 85] can be alternatives, two-stream models still exhibit inefficiency at inference by doubling the net-

\*Corresponding author

<sup>1</sup>Each model is reproduced by its official codebase.

work forwarding process. Therefore, it would be desirable to build strong RGB-based action detectors that can bridge the performance gap with two-stream methods.

To this end, we focus on cross-modal knowledge distillation [12, 16], where the helpful knowledge of motion modality is transferred to an RGB-based action detector during training in order to improve its performance. In contrast to conventional knowledge distillation [19, 20, 46, 52] where the superior teacher guides the weak student, cross-modal distillation requires exploiting the complementarity of the teacher and student. However, existing cross-modal distillation approaches [12, 13] fail to consider the difference and directly transfer the motion knowledge to the RGB model (Fig. 1a), as conventional distillation does. By design, the RGB and motion information are entangled with each other, making it difficult to balance between them. As a result, they often achieve limited gains without careful tuning.

To tackle the issue, we introduce a novel framework, named *decomposed cross-modal distillation* (Fig. 1b). In detail, our model adopts the split-and-merge paradigm, where the high-level features are decomposed into appearance and motion components within a dual-branch design. Then only the motion branch receives the distillation signal, while the other branch remains intact to learn appearance information. For explicit decomposition, we adopt the shared detection head and the asymmetric objective functions for the branches. Moreover, we design a novel attentive fusion to effectively combine the multimodal information provided by the two branches. In contrast to existing attention methods, the proposed fusion preserves local sensitivity which is important for accurate action detection. With these key components, we build a strong action detector that produces precise action predictions given only RGB frames.

We conduct extensive experiments on the popular benchmarks, THUMOS'14 [22] and ActivityNet1.3 [4]. Experimental results show that the proposed framework enables effective cross-modal distillation by separating the RGB and motion features. Consequently, our model largely improves the performance of RGB-based action detectors, exhibiting its superiority over conventional distillation. The resulting RGB-based action detectors effectively bridge the gap with two-stream models. Moreover, we validate our approach by utilizing another motion clue, *i.e.*, temporal gradient, which has been underexplored for action detection.

To summarize, our contributions are three-fold: 1) We propose a decomposed cross-modal distillation framework, where motion knowledge is transferred in a separate way such that appearance information is not harmed. 2) We design a novel attentive fusion method that is able to exploit the complementarity of two modalities while sustaining the local discriminability of features. 3) Our method is generalizable to various backbones and detection heads, showing consistent improvements.

## 2. Related Works

### 2.1. Temporal Action Detection

Temporal action detection requires predicting temporal intervals as well as action categories for all action instances occurring in the video. Conventional methods adopt the two-stage pipeline (*i.e.*, proposal-and-classification) and generate proposals by either sliding windows [7, 55, 72, 76, 77, 86] or predicting the per-frame starting and ending probabilities [33, 36, 37, 84]. Besides, several methods focus on proposal refinement to improve the detection performance [49, 79, 88]. Meanwhile, analogous to one-stage object detection [39, 50], anchor-free models are proposed for efficient action detection [34, 80]. Inspired by the recent success of DETR [5], query-based action detectors are also designed to streamline the complicated detection pipeline [42, 53, 60]. In an orthogonal direction, some works showcase the benefit of end-to-end training of video backbones for temporal action detection [10, 40, 73].

Most of the current action detectors leverage two-stream inputs for accurate action localization. However, the optical flow takes heavy computations, suggesting the necessity of RGB-based action detection models. In this paper, we propose a novel distillation framework to build a strong RGB-based action detector by transferring motion knowledge.

### 2.2. Cross-modal Knowledge Distillation

Knowledge distillation [20] is originally devised to transfer the knowledge of large-scale models (teachers) to smaller ones (students). Existing approaches can be categorized into three groups based on what types of knowledge are distilled: responses, features, and relations. Response-based methods encourage the student to produce similar predictions to those of the teacher [3, 9, 20, 81]. Meanwhile, feature-based distillation methods pursue the matching of intermediate features between students and teachers [46, 52, 57, 65]. Lastly, relation-based models focus on aligning the inter-sample relationship between the teacher and the student [43, 48, 61]. In practice, the three types of distillation are often used in a combinational form [14, 81, 87].

As a subset, cross-modal knowledge distillation aims to transfer knowledge of one modality to another modality. Generally, the model taking the superior modality on the target task is selected as the teacher and guides the inferior modality model [51, 83]. On the other hand, in the regime of action analysis, there is no dominant modality, and it is crucial to grasp the complementarity of different modalities [6, 58]. However, existing works [12, 13, 16, 45] ignore this property and directly inject motion knowledge into the RGB model, leading to entangled representations. In contrast, we propose a decomposed distillation framework, where different modality information is separately learned and recomposed for effective action detection.

### 3. Method

**Problem formulation.** The input of temporal action detection models is a video  $V \in \mathbb{R}^{T \times 3 \times H \times W}$  consisting of a total of  $T$  frames with the size of  $H \times W$ . Here the video length  $T$  can vary across input videos. Since most of the action detectors are based on two-stream inputs, the optical flow maps are additionally extracted from each pair of consecutive RGB frames. We denote the optical flow maps by  $F \in \mathbb{R}^{T \times 2 \times H \times W}$ , where each map  $F_t$  is composed of two channels that respectively estimate the displacements of  $x$  and  $y$  axes. During training, the input video is labeled with its corresponding annotation  $\Psi = \{(\varphi_m, \mathbf{y}_m)\}_{m=1}^M$ , where  $\mathbf{y}_m$  is the category label of the  $m$ -th action instance and  $\varphi_m = (t_{s_m}, t_{e_m})$  indicates its starting and ending timestamps. Typically, the class label  $\mathbf{y}_m \in \mathbb{R}^C$  is a one-hot vector, where  $C$  indicates the number of categories. In the test phase, a two-stream action detector localizes the action instances based on both  $V$  and  $F$ . Differently, to bypass time-consuming optical flow extraction and inefficient inference, we aim to build an RGB-based action detector that takes only  $V$  as input. To this end, we leverage a pre-trained motion model as the teacher and transfer its knowledge to the RGB-based action detector in a decomposed way.

**Motion modality.** Conventionally, optical flow is utilized as the motion modality for action detection [7, 37]. Meanwhile, another form of the motion modality, *i.e.*, temporal gradient, has recently been explored for action recognition [70, 71]. Inspired by this, we explore temporal gradient for temporal action detection for the first time. The temporal gradient maps  $G \in \mathbb{R}^{T \times 3 \times H \times W}$  is defined by the residual difference between two consecutive frames, *i.e.*,  $G_t = V_t - V_{t-1}$ , which implicitly captures motion information such as camera moving, object moving, *etc.* Fig. 2 exemplifies an RGB frame and the corresponding optical flow and temporal gradient. Although vulnerable to noises like environmental changes, the temporal gradient can be obtained on-the-fly and thus have the potential to be utilized as a weak form of motion modality. Indeed, it is shown in the experiments that its motion knowledge helps to improve the action detection performance. In the following, we will elaborate on our framework using optical flow maps  $F$  as the motion modality by default, but they can be replaced by temporal gradient maps  $G$  without loss of generality.

#### 3.1. Motion Teacher Training

In this section, we describe the training of the motion teacher. Given the input optical flow  $F$ , the video backbone extracts 1D spatially pooled features  $z^{\text{mot}} \in \mathbb{R}^{T/r_T \times C}$ , where  $r_T$  indicates the temporal downsampling rate and  $C$  is the channel dimension. Here  $r_T$  can vary depending on the backbone choices. The extracted features go through the action detection head, resulting in a set of predicted action boundaries and scores  $\hat{\Psi}^{\text{mot}} = \{(\hat{\varphi}_n^{\text{mot}}, \hat{\mathbf{y}}_n^{\text{mot}})\}_{n=1}^N$ ,




Figure 2. An example of an RGB frame and the corresponding optical flow map and temporal gradient map.

where  $\hat{\varphi}_n^{\text{mot}} = (\hat{t}_{s_n}^{\text{mot}}, \hat{t}_{e_n}^{\text{mot}})$  is the predicted action proposals,  $N$  is the number of predictions, and  $\hat{\mathbf{y}}_n$  is the class probability with Sigmoid activation. We leave the model architecture and the detailed prediction process to be abstract since our framework can be applied to any framework type, *e.g.*, anchor-based [44, 74], anchor-free [34, 80], query-based [42, 60]. We show in experiments that our approach brings consistent improvements for different types of heads. The action predictions are grouped by positive and negative sets based on pre-defined principles; some works perform thresholding the intersection-over-union (IoU) with ground truths [7, 55], while other query-based methods rely on the Hungarian matching [40]. We denote the positive and negative sets by  $\hat{\Psi}_P^{\text{mot}}$  and  $\hat{\Psi}_N^{\text{mot}}$ , respectively. The predictions in the positive set are labeled by the best-matched ground-truth intervals and classes  $\{(\varphi_n, \mathbf{y}_n)\}_{n=1}^{|\hat{\Psi}_P^{\text{mot}}|}$ . On the other hand, the predictions in the negative set are annotated as the background class. Despite minor differences, the training objective of action detectors generally consists of a classification loss, a regression loss, and a completeness loss.

1) The classification loss is defined by a weighted sum of the focal binary cross-entropy function as follows.

$$\mathcal{L}_{\text{cls}} = \frac{\alpha_P}{|\hat{\Psi}_P^{\text{mot}}|} \sum_{\hat{\mathbf{y}}_n \in \hat{\Psi}_P^{\text{mot}}} \ell_{\text{focal}}(\hat{\mathbf{y}}_n, \mathbf{y}_n) + \frac{\alpha_N}{|\hat{\Psi}_N^{\text{mot}}|} \sum_{\hat{\mathbf{y}}_n \in \hat{\Psi}_N^{\text{mot}}} \ell_{\text{focal}}(\hat{\mathbf{y}}_n, \mathbf{0}), \quad (1)$$

where  $\ell_{\text{focal}}(\hat{\mathbf{y}}, \mathbf{y})$  denotes the focal binary cross-entropy loss [38],  $\alpha_*$  is the term for balancing between positive and negative samples, and  $\mathbf{0} \in \mathbb{R}^C$  is the  $C$ -dimensional zero vector indicating the background class.

2) The regression loss is defined by the  $L_1$ -distance between ground truths and the predicted offsets as follows.

$$\mathcal{L}_{\text{reg}} = \frac{1}{|\hat{\Psi}_P^{\text{mot}}|} \sum_{\hat{\varphi}_n \in \hat{\Psi}_P^{\text{mot}}} \ell_1(\hat{\varphi}_n, \varphi_n), \quad (2)$$

where  $\ell_1(\hat{\varphi}, \varphi)$  denotes the (smooth)  $L_1$  distance between the two offsets. It is worth noting that the regression loss is computed only for the positive set.

3) The completeness loss is utilized to maximize the IoUs between ground truths and the predicted proposals.

$$\mathcal{L}_{\text{comp}} = \frac{1}{|\hat{\Psi}_P^{\text{mot}}|} \sum_{\hat{\varphi}_n \in \hat{\Psi}_P^{\text{mot}}} \left(1 - \frac{|\hat{\varphi}_n \cap \varphi_n|}{|\hat{\varphi}_n \cup \varphi_n|}\right). \quad (3)$$

The completeness loss is defined only for the positive set.

The overall objective function is a weighted sum of the above losses:  $\lambda_{\text{cls}}\mathcal{L}_{\text{cls}} + \lambda_{\text{reg}}\mathcal{L}_{\text{reg}} + \lambda_{\text{comp}}\mathcal{L}_{\text{comp}}$ . After training, the motion model  $\mathcal{M}^{\text{mot}}$  is able to extract discriminative motion features and will serve as the teacher for distillation.

### 3.2. Decomposed Cross-modal Distillation

**Motivation.** Previous cross-modal knowledge distillation frameworks [12, 13] transfer motion knowledge to the RGB model in a direct way, so that the RGB model is encouraged to produce similar predictions with the motion model. Such an approach, however, would result in sub-optimal solutions, as the complementarity of the two modalities is not considered. To handle the challenge, we design a novel decomposed distillation framework, where RGB and motion information are separately learned and later fused to effectively exploit their complementarity. The training process of our distillation framework is illustrated in Fig. 3a.

Similar to the motion model, given RGB frames  $V$ , the video backbone first produces the spatiotemporal features  $z^{\text{RGB}} \in \mathbb{R}^{T/r_T \times C}$ . The representation  $z^{\text{RGB}}$  is projected into two different spaces, *i.e.*, motion and appearance spaces. The projections are implemented by a series of 1D convolutional layers, and we denote this process by  $\phi_{\text{app}}(z^{\text{RGB}}), \phi_{\text{mot}}(z^{\text{RGB}}) \in \mathbb{R}^{T/r_T \times D}$ , respectively. Afterward, the projected features serve as input to the subsequent separate branches to predict the temporal intervals of action instances. In the following, we describe how the separated features learn their own representations.

#### 3.2.1 Motion branch

The motion branch is supposed to learn motion information from the RGB frames. For this purpose, the pre-trained motion model  $\mathcal{M}^{\text{mot}}$  (teacher) explicitly guides the motion branch (student) by distilling its knowledge. A variety of distillation approaches can be applied to our framework, and we adopt two representative methods in the following.

1) The response-based distillation loss [9, 20] utilizes the teacher’s predictions as pseudo-ground truths and encourages the student to mimic the behavior of the teacher.

$$\mathcal{L}_{\text{respon}} = \lambda_{\text{cls}}\mathcal{L}_{\text{respon}}^{\text{cls}} + \lambda_{\text{reg}}\mathcal{L}_{\text{respon}}^{\text{reg}},$$

where  $\mathcal{L}_{\text{respon}}^{\text{cls}} = \frac{1}{|\hat{\Psi}|} \sum_{n=1}^{|\hat{\Psi}|} \ell_{\text{focal}}(\hat{\mathbf{y}}_n, \hat{\mathbf{y}}_n^{\text{mot}})$  (4)

and  $\mathcal{L}_{\text{respon}}^{\text{reg}} = \frac{1}{|\hat{\Psi}|} \sum_{n=1}^{|\hat{\Psi}|} \ell_1(\hat{\varphi}_n, \hat{\varphi}_n^{\text{mot}})$ .

where  $(\hat{\varphi}_n^{\text{mot}}, \hat{\mathbf{y}}_n^{\text{mot}})$  is the corresponding teacher’s predictions of the student’s ones. Here  $\ell_{\text{focal}}$  and  $\ell_1$  are defined as same in Eq. 1 and 2, respectively.

2) The feature-based distillation loss [52] encourages the matching of the intermediate features from the student and

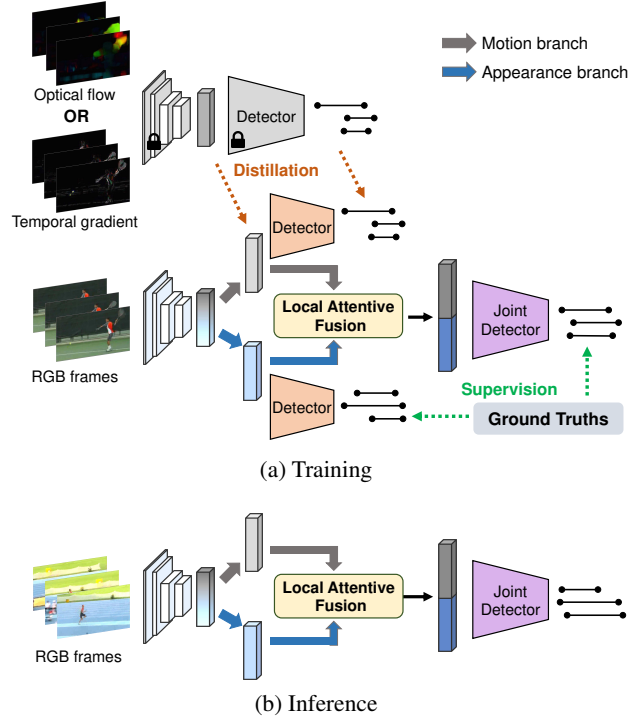


Figure 3. Overall workflow of the proposed framework. (a) During training, our model performs decomposed cross-modal distillation by explicitly separating motion and appearance branches. Here the detection heads of two branches are shared. (b) At inference, our model takes as input only RGB frames and precisely localizes action instances based on multimodal information.

the teacher, which is formulated as:

$$\mathcal{L}_{\text{feat}} = \frac{1}{T} \sum_{t=1}^T \|\phi_{\text{mot}}(z_t^{\text{RGB}}) - z_t^{\text{mot}}\|_2^2, \quad (5)$$

where the projection layer  $\phi_{\text{mot}}(\cdot)$  can be viewed as an adaptation layer [8] that helps to ease the feature matching between different modality features. Note that although this equation represents the matching of backbone features, other features from any intermediate layer can be aligned.

To summarize, the total training objective function of the motion branch is defined by the weighted sum of the distillation losses:  $\mathcal{L}_{\text{distill}} = \mathcal{L}_{\text{respon}} + \lambda_{\text{feat}}\mathcal{L}_{\text{feat}}$ . With the guidance of the pre-trained motion model, the motion branch learns motion representations given RGB frames as inputs. Note that our framework is general and can utilize other improved distillation methods, but it is beyond the scope of this paper.

#### 3.2.2 Appearance branch

On the other side, to prevent the model from losing rather static RGB information during distillation, we train the appearance branch jointly with the motion branch. The training objective is the same with the motion teacher  $\mathcal{M}^{\text{mot}}$ ,

except that the input is RGB frames. Here we define the training objective of the appearance branch as:  $\mathcal{L}_{\text{app}} = \lambda_{\text{cls}}\mathcal{L}_{\text{cls}} + \lambda_{\text{reg}}\mathcal{L}_{\text{reg}} + \lambda_{\text{comp}}\mathcal{L}_{\text{comp}}$ . Being trained with the conventional detection loss, the appearance branch is able to keep the appearance information intact during distillation which is useful for action detection.

### 3.2.3 Discussion

Our decomposed distillation model contains two separate branches with different design purposes. Specifically, the motion branch needs to imitate the behavior of the motion teacher, while the appearance branch is supposed to learn the original appearance information from the same RGB inputs. To achieve this goal, we propose to endow the branches with conflicting training objectives. However, the two branches may reach a degenerate solution, learning similar representations without distinction and relying solely on the powerful detection head to produce different predictions. To tackle the potential issue, we make the branches share the weights of the classifier and the regressor in the detection heads. With these two key designs, we can lead the two branches to learn the different information of RGB and motion modalities, thus achieving their complementarity of theirs in the subsequent fusion stage.

### 3.3. Local Attentive Fusion

Provided that the appearance and motion features have learned the expected information, the next question is how to fuse the different information from them. An intuitive way would be to perform concatenation before feeding them to the joint detection head, also known as early fusion [23]. However, we argue that such a naive way does not help much since erroneous predictions of one modality can propagate to the other. Instead, we utilize an attentive fusion mechanism to maximize the harmonizing effect of the two different modalities.

Our motivation is that one modality can enhance the other by highlighting the agreement between different modalities. One can try applying cross-attention for that purpose, where features of one modality aggregate information from those of the other modality based on feature similarity. Although the cross-attention proves to be effective in various fields [18, 67], we experimentally find that it does not help much to improve the action detection performance. We conjecture that this is because the cross-attention hampers local discriminability of the features by gathering information along the temporal dimension, *i.e.*, feature over-smoothing, which is also observed in the literature [17, 60]. To bypass the issue, we propose a new attentive fusion that sustains locally discriminative features for action detection.

Let us call the modality to be enhanced the target modality and the other the reference modality. At first, we build a

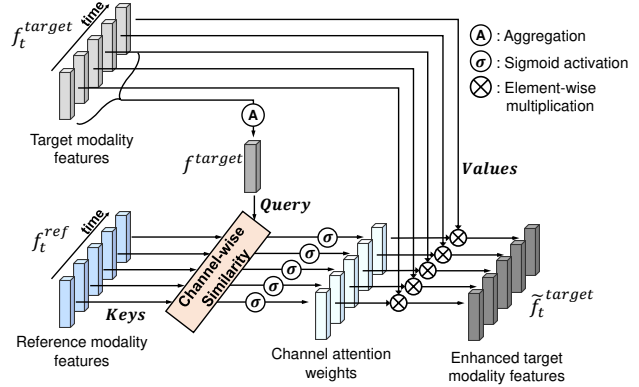


Figure 4. Illustration of the proposed local attentive fusion.

single representative feature for the target modality by aggregating the whole features along the temporal dimension, *i.e.*,  $f^{\text{target}} = \psi(f_1^{\text{target}}, \dots, f_T^{\text{target}}) \in \mathbb{R}^D$ , where the aggregation function can be implemented by various pooling methods. Thereafter, it serves as the query and the individual features from the assistant modality are deemed as the key. Then, the query-key matching process is defined by:

$$\omega_t = \sigma[(W_{\text{query}}^\top f^{\text{target}}) \odot (W_{\text{key}}^\top f_t^{\text{ref}})], \quad (6)$$

where  $W_{\text{query}}, W_{\text{key}} \in \mathbb{R}^{D \times D}$  respectively indicate the projection matrices for the query and keys,  $\odot$  denotes the Hadamard product, and  $\sigma(\cdot)$  is the sigmoid activation. Here the resulting weight  $\omega_t \in [0, 1]^D$  can be viewed as the channel attention weights that are derived from the similarity between the two modalities. Therefore, we can suppress erroneous predictions of the target modality and emphasize the mutually agreed information by applying the weight to the features of the target modality, *i.e.*,  $\tilde{f}_t^{\text{target}} = \omega_t \cdot f_t^{\text{target}}$ . The overall process of our attentive fusion is illustrated in Fig. 4. Importantly, our attentive fusion does not fuse information of different temporal points and thus can preserve the local sensitivity of the features. The advantage of our local attentive fusion over the conventional cross-attention will be verified in Sec. 4. After enhancing the appearance and motion features based on each other, we concatenate and put them into the joint detection head. The detection head is trained to perform the action detection given the enhanced multimodal features. The loss function of the detection head is the same as that of the appearance branch:  $\mathcal{L}_{\text{joint}} = \lambda_{\text{cls}}\mathcal{L}_{\text{cls}} + \lambda_{\text{reg}}\mathcal{L}_{\text{reg}} + \lambda_{\text{comp}}\mathcal{L}_{\text{comp}}$ .

### 3.4. Joint Training and Inference

Our framework is trained in an end-to-end manner. The total training objective of our model is defined as  $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{app}} + \mathcal{L}_{\text{distill}} + \mathcal{L}_{\text{joint}}$ . At inference, we discard the detection heads of the branches and predict action intervals and classes using the joint head, as depicted in Fig. 3b. It is noteworthy that our model takes only RGB frames as input and localizes action instances based on multimodal information.

distillation		local attn.	mAP@IoU (%)					AVG
conven.	decomp.		0.3	0.4	0.5	0.6	0.7	
✗	✗	✗	62.3	55.2	46.2	33.8	20.4	43.6
✓			62.5	55.7	47.3	35.1	21.8	44.5
	✓		63.3	56.2	47.9	36.1	22.9	45.2
	✓	✓	64.4	58.0	49.0	37.5	24.1	46.6

Table 2. Ablation study on the effect of each component. The comparative methods are evaluated on THUMOS’14.

## 4. Experiments

### 4.1. Experimental Setups

**Datasets.** We evaluate our framework on the two most popular benchmarks for temporal action detection: THUMOS’14 [22] and ActivityNet1.3 [4]. THUMOS’14 consists of 200 and 213 videos respectively for training and testing with 20 action categories. It is a challenging dataset since it contains frequent action instances, *e.g.*, 15 instances per video on average. Therefore, we utilize it as the main dataset for experiments. On the other hand, ActivityNet1.3 is a relatively large-scale dataset containing 10,024, 4,926, and 5,044 videos respectively for training, validation, and testing. As the ground truths of the test set are unavailable, we evaluate the comparative models on the validation set.

**Evaluation metrics.** The standard protocol for evaluating action detectors is mean average precisions (mAPs) at different intersection-over-union (IoU) thresholds. Following the convention [7, 41, 74], we set the thresholds to [0.3:0.7] with a step size of 0.1 for THUMOS’14 and [0.5:0.95] with a step size of 0.05 for ActivityNet1.3.

**Implementation details.** Our framework is agnostic to the choices of backbones and action detection heads. To confirm the generalizability of the proposed method, we conduct experiments using different combinations. The video backbones are TSM18 [35], TSM50, I3D [6], and Slowfast50 [15]. For the detection head, we utilize the representative models of three different detector types: anchor-based (GTAD [74]), anchor-free (Actionformer [80]), and query-based (TadTR [42]). For model implementation and hyperparameter settings, we strictly follow the official codebases. All the backbones are pre-trained on Kinetics-400 [6]. Following the previous works [34, 37, 80], we use the TV- $L^1$  algorithm [1] to extract dense optical flow. Since only the I3D model has pre-trained weights on the optical flow, we utilize it as the teacher network in the distillation setting of optical flow. Besides, we also conduct experiments using the temporal gradient as the motion modality, where the identical backbone networks serve as the teacher to guide the RGB model. We follow Liu *et al.* [40] for data processing, where the input is a video with the size of  $96 \times 96$  and the length of 25.6 seconds. Our model is trained in an end-to-end fashion using the Adam optimizer [26] with the learning rate of  $1e-4$  for 20 epochs.

Fusion	mAP@IoU (%)					AVG
	0.3	0.4	0.5	0.6	0.7	
concat.	63.3	56.2	47.9	36.1	22.9	45.2
sum.	62.6	56.1	47.5	36.1	23.0	45.1
self-attn.	63.8	56.3	46.7	34.2	21.9	44.6
cross-attn.	63.1	54.5	46.4	35.4	21.7	44.2
diff.-attn.	61.8	54.8	46.3	32.6	21.0	43.3
local attn. (Ours)	64.4	58.0	49.0	37.5	24.1	46.6

Table 3. Ablation study on the fusion mechanism. The comparative methods are evaluated on THUMOS’14.

### 4.2. Analysis

#### 4.2.1 Effect of each component

To analyze the impact of each component, we conduct an ablation study on THUMOS’14 in Table 2. For this study, we utilize TSM18 [35] and Actionformer [80] as the backbone and the detection head, respectively. We first set our baseline using the pure RGB-based action detector, whose average mAP is 43.6%. On top of it, we apply two different types of distillation with temporal gradient as the motion modality, namely the conventional and decomposed ones. For the decomposed distillation, we simply concatenate features from the two branches and feed them to the joint branch for action detection. As a result, the conventional distillation achieves a limited performance gain, while our decomposed distillation greatly boosts the detection performance (2<sup>nd</sup>-3<sup>rd</sup> rows). This indicates that the way of directly transferring the motion knowledge to the RGB model is prone to a sub-optimal solution due to entangled information. When adopting the proposed local attentive fusion, our model better grasps the multimodal complementarity and shows a further performance gain of 1.4% (4<sup>th</sup> row), indicating the importance of information fusion. The resulting distilled RGB model improves the baseline by 3.0%, which clearly verifies the effectiveness of the proposed methods.

#### 4.2.2 Ablation study on fusion

To verify the effectiveness of the proposed local attentive fusion, we compare it with a variety of fusion methods. In specific, we employ two naïve fusion methods, *i.e.*, concatenation and summation, and three attention-based approaches [62, 75], *i.e.*, self-, cross-, and difference-attention. The results are shown in Table 3, where we observe that the naïve fusion shows decent performance, while the attention-based methods perform even poorer than the naïve ones. We conjecture that this is due to the over-smoothing of the features, which are known to be a side effect of attention-based methods [17, 54, 60], leading to less discriminability in action boundary detection. On the contrary, the proposed local attentive fusion successfully preserves the local sensitivity of features during multimodal information exchange, thereby achieving the best action localization performance.

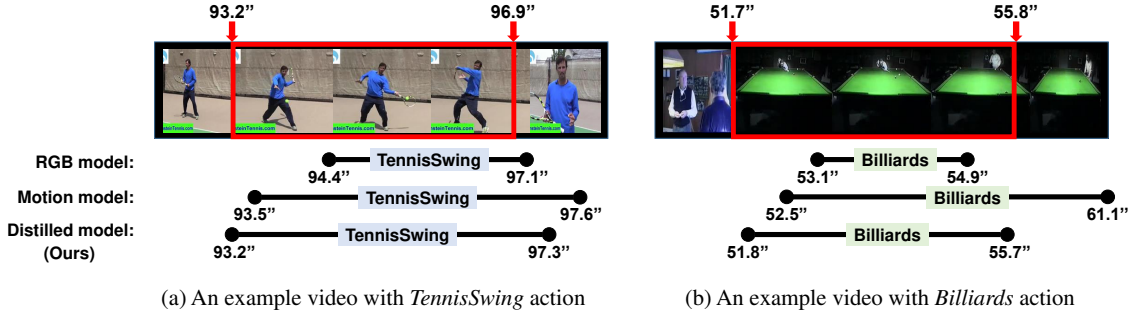


Figure 5. Qualitative results on THUMOS'14.

Backbone	Distill.	mAP@IoU (%)					AVG
		0.3	0.4	0.5	0.6	0.7	
TSM18 [35]	✗	62.3	55.2	46.2	33.8	20.4	43.6
	TG	64.4	58.0	49.0	37.5	24.1	46.6 (+3.0)
	OF	65.3	59.5	50.9	39.6	25.5	48.2 (+4.6)
TSM50 [35]	✗	65.0	59.2	50.0	38.2	25.0	47.5
	TG	68.1	61.8	52.4	41.7	27.5	50.3 (+2.8)
	OF	66.5	62.3	55.3	44.5	32.9	52.3 (+4.8)
I3D [6]	✗	53.8	47.0	38.6	30.0	19.9	37.9
	TG	57.6	51.4	42.5	32.9	22.1	41.3 (+3.4)
	OF	57.7	52.1	44.6	34.9	24.0	42.6 (+4.7)
Slowfast50 [15]	✗	67.4	62.9	56.8	46.8	35.0	53.8
	TG	68.9	64.1	58.1	48.2	35.6	55.0 (+1.2)
	OF	70.5	65.8	59.2	50.1	38.2	56.8 (+3.0)

Table 4. Experiments with different backbones on THUMOS'14. The detection head is fixed to Actionformer [80]. 'TG' and 'OF' denote temporal gradient and optical flow, respectively.

### 4.2.3 Generalizability

It is worth noting that our distillation framework is general and can be applied to any video backbones and detection heads. To analyze the generalizability of our method, we conduct comprehensive experiments on THUMOS'14.

The experimental results on four different video backbones are presented in Table 4. We use Actionformer [80] as the detection head for this experiment. It can be noticed that regardless of the backbone choices, our decomposed distillation consistently improves the performance with large gains. In addition, even the strong backbone, *i.e.*, Slowfast, benefits from the motion knowledge distillation, showing a performance gain of 3.0 % when using optical flow as the teacher modality. In addition, the results using different detection heads are provided in Table 5, where TSM18 [35] is employed as the backbone network. Again, it is noticeable that all the heads benefit from the decomposed distillation to a large extent. In detail, they show significant performance boosts when being distilled the motion knowledge from the optical flow. On the other hand, albeit being efficiently obtained from RGB frames, the temporal gradient also brings nontrivial performance gains, shedding light on its potential to serve as an efficient motion modality. To summarize, these experiments clearly validate the generalizability of the

Head	Distill.	mAP@IoU (%)					AVG
		0.3	0.4	0.5	0.6	0.7	
G-TAD [74]	✗	51.4	44.7	36.0	26.4	16.8	35.1
	TG	54.8	48.9	38.1	28.0	18.1	37.6 (+2.5)
	OF	55.3	49.4	39.2	30.6	19.7	38.8 (+3.6)
TadTR [42]	✗	62.8	56.7	47.5	37.3	25.5	46.0
	TG	63.8	57.4	49.9	39.2	26.9	47.4 (+1.4)
	OF	64.1	58.3	51.2	40.9	28.8	48.7 (+2.7)
Actionformer [80]	✗	62.3	55.2	46.2	33.8	20.4	43.6
	TG	64.4	58.0	49.0	37.5	24.1	46.6 (+3.0)
	OF	65.3	59.5	50.9	39.6	25.5	48.2 (+4.6)

Table 5. Experiments with different detection heads on THUMOS'14. The backbone is fixed to TSM18 [35]. 'TG' and 'OF' indicate temporal gradient and optical flow, respectively.

proposed distillation framework.

### 4.3. Qualitative Results

To analyze where the performance gains come from, we provide several qualitative comparisons in Fig. 5. Specifically, we visualize the detection results from the RGB baseline, the motion teacher, and the resultant RGB-based model of our decomposed distillation. In both examples, it can be observed that the RGB-based baseline shows unsatisfying performance, indicating the importance of motion modality. Meanwhile, the motion teacher suffers from the static video with tiny movements (Fig. 5b), resulting in inaccurate localization results. On the other hand, our model produces precise action proposals for both examples by successfully fusing the multimodal information learned within the decomposed distillation framework.

### 4.4. Comparison with State-of-the-arts

The state-of-the-art comparison of THUMOS'14 and ActivityNet1.3 is shown in Table 6. For the comparison, we utilize Slowfast50 [15] and Actionformer [80] for the backbone and the detection head, respectively. To make the comparison clear, we separate the entries based on whether the models take as input two-stream data or not. The first thing we can observe is that the RGB-based approaches fall largely behind the two-stream approaches, especially on THUMOS'14 compared to ActivityNet1.3. This is because

Method	Venue	OF	THUMOS'14						ActivityNet1.3			
			0.3	0.4	0.5	0.6	0.7	AVG	0.5	0.75	0.95	AVG
TAL-Net [7]	CVPR'18	✓	53.2	48.5	42.8	33.8	20.8	39.8	38.23	18.30	1.30	20.22
BSN [37]	ECCV'18	✓	53.5	45.0	36.9	28.4	20.0	-	46.45	29.96	8.02	30.03
BMN [36]	ICCV'19	✓	56.0	47.4	38.8	29.7	20.5	38.5	50.07	34.70	8.29	33.85
P-GCN [79]	ICCV'19	✓	63.6	57.8	49.1	-	-	-	48.26	33.16	3.27	31.11
G-TAD [74]	CVPR'20	✓	54.5	47.6	40.2	30.8	23.4	39.3	50.36	34.60	9.02	34.09
BC-GNN [2]	ECCV'20	✓	57.1	49.1	40.4	31.2	23.1	40.2	50.56	34.75	9.37	34.26
BU-MR [84]	ECCV'20	✓	53.9	50.7	45.4	38.0	28.5	43.3	43.47	33.91	9.21	30.12
AFSD [34]	CVPR'21	✓	67.3	62.4	55.5	43.7	31.1	52.0	52.38	35.27	6.47	34.39
MUSES [41]	CVPR'21	✓	68.9	64.0	56.9	46.3	31.0	53.4	50.02	34.97	6.57	33.99
RTD-Net [60]	ICCV'21	✓	68.3	62.3	51.9	38.8	23.7	49.0	47.21	30.68	8.61	30.83
VSGN [82]	ICCV'21	✓	66.7	60.4	52.4	41.0	30.4	50.2	52.38	36.01	8.37	35.07
RCL [64]	CVPR'22	✓	70.1	62.3	52.9	42.7	30.7	51.7	55.15	39.02	8.27	37.65
RefactorNet [68]	CVPR'22	✓	70.7	65.4	58.6	47.0	32.1	54.8	56.60	40.70	7.50	38.60
TAGS [47]	ECCV'22	✓	68.6	63.8	57.0	46.3	31.8	52.8	56.30	36.80	9.60	36.50
ReAct [53]	ECCV'22	✓	69.2	65.0	57.1	47.8	35.6	55.0	49.60	33.00	8.60	32.60
Actionformer [80]	ECCV'22	✓	82.1	77.8	71.0	59.4	43.9	66.8	53.50	36.20	8.20	35.60
CDC [55]	CVPR'17	✗	40.1	29.4	23.3	13.1	7.9	22.8	45.30	26.00	0.20	23.80
GTAN [44]	CVPR'19	✗	57.8	47.2	38.8	-	-	-	52.61	34.14	8.91	34.31
G-TAD* [74]	CVPR'20	✗	52.5	45.9	37.6	28.5	19.1	36.7	49.22	34.55	4.74	33.17
AFSD* [34]	CVPR'21	✗	57.7	52.8	45.4	34.9	22.0	43.6	-	-	-	32.90
TadTR* [42]	TIP'22	✗	59.6	54.5	47.0	37.8	26.5	45.1	49.56	35.24	9.93	34.35
E2E-TAD [40]	CVPR'22	✗	69.4	64.3	56.0	46.4	34.9	54.2	50.47	35.99	10.83	35.10
TAGS† [47]	ECCV'22	✗	59.8	57.2	50.7	42.6	29.1	47.9	54.44	34.95	8.71	34.95
Actionformer† [80]	ECCV'22	✗	69.8	66.0	58.7	48.3	34.6	55.5	53.21	35.15	8.03	34.94
Ours	-	✗	70.5	65.8	59.2	50.1	38.2	56.8	53.73	35.87	8.61	35.58

Table 6. Comparison with state-of-the-art methods. The average mAPs under the IoU thresholds 0.3:0.7 and 0.5:0.95 are reported respectively for THUMOS'14 and ActivityNet1.3. Entries are grouped by whether the model relies on optical flow (OF) at inference time. The results of models with asterisk (\*) are taken from Liu *et al.* [40], while those with dagger (†) are the reproduced results by official code.

the THUMOS'14 benchmark contains frequent actions occurring, *e.g.*, 15 instances per video on average, and therefore the motion modality has a large impact on localizing the actions. On the other hand, ActivityNet1.3 have sparse and relatively long action instances, *e.g.*, 1.5 instances per video on average, and it is widely known that classification rather than localization is important for the dataset due to a large number of action classes. In the comparison results on THUMOS'14, our model achieves the state-of-the-art performance among the RGB-based action detectors. In addition, it surpasses many two-stream action detection models even without relying on the motion modality during inference. This signifies the efficacy of our distillation framework that enables simulating two-stream predictions by effectively distilling the motion knowledge and exploiting the multimodal complementarity. On ActivityNet1.3, our model also sets a new state-of-the-art among the RGB-based detectors and shows comparable performance with two-stream approaches.

## 5. Conclusion

In this paper, we have presented a new paradigm for cross-modal distillation. Specifically, we pointed out that

existing distillation approaches inevitably entangle the RGB and motion information during the distilling process. To handle the issue, we propose a decomposed distillation pipeline that enables separate learning of different modalities. Furthermore, we design a local attentive fusion to sustain the local discriminability of features during integrating multimodal information, thereby accomplishing accurate action detection. In the extensive experiments, we verified the effectiveness of our distillation framework and local attentive fusion. Moreover, our model successfully bridges the gap between two-stream and RGB-based action detectors while preserving efficiency at test time. Notably, our framework is generalizable to various combinations of video backbones and action detection heads, demonstrating consistent performance improvements. In the future, it would be interesting to explore our decomposed distillation framework for other multimodal tasks.

## Acknowledgements

This project was partly supported by the NAVER Corporation and the National Research Foundation of Korea grant funded by the Korea government (MSIT) (No. 2022R1A2B5B02001467). Part of the computational work was conducted on NAVER Smart Machine Learning (NSML) platform [25, 59].



## References

- [1] Stuart Andrews, Ioannis Tsochantaridis, and Thomas Hofmann. Support vector machines for multiple-instance learning. In *Neurips*, 2002. 6
- [2] Yueran Bai, Yingying Wang, Yunhai Tong, Yang Yang, Qiyue Liu, and Junhui Liu. Boundary content graph neural network for temporal action proposal generation. In *ECCV*, pages 121–137, 2020. 8
- [3] Lucas Beyer, Xiaohua Zhai, Amélie Royer, Larisa Markeeva, Rohan Anil, and Alexander Kolesnikov. Knowledge distillation: A good teacher is patient and consistent. In *CVPR*, pages 10925–10934, 2022. 2
- [4] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *CVPR*, pages 961–970, 2015. 2, 6
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229, 2020. 2
- [6] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 6299–6308, 2017. 1, 2, 6, 7
- [7] Yu-Wei Chao, Sudheendra Vijayanarasimhan, Bryan Seybold, David A Ross, Jia Deng, and Rahul Sukthankar. Rethinking the faster r-cnn architecture for temporal action localization. In *CVPR*, pages 1130–1139, 2018. 2, 3, 6, 8
- [8] Defang Chen, Jian-Ping Mei, Hailin Zhang, Can Wang, Yan Feng, and Chun Chen. Knowledge distillation with the reused teacher classifier. In *CVPR*, pages 11933–11942, 2022. 4
- [9] Guobin Chen, Wongun Choi, Xiang Yu, Tony Han, and Manmohan Chandraker. Learning efficient object detection models with knowledge distillation. *Neurips*, 30, 2017. 2, 4
- [10] Feng Cheng and Gedas Bertasius. Tallformer: Temporal action localization with long-memory transformer. In *ECCV*, 2022. 2
- [11] Jinwoo Choi, Chen Gao, Joseph CE Messou, and Jia-Bin Huang. Why can't i dance in the mall? learning to mitigate scene bias in action recognition. *Neurips*, 32, 2019. 1
- [12] Nieves Crasto, Philippe Weinzaepfel, Karteek Alahari, and Cordelia Schmid. Mars: Motion-augmented rgb stream for action recognition. In *CVPR*, pages 7882–7891, 2019. 2, 4
- [13] Rui Dai, Srijan Das, and François Bremond. Learning an augmented rgb representation with cross-modal knowledge distillation for action detection. In *ICCV*, pages 13053–13064, 2021. 2, 4
- [14] Xing Dai, Zeren Jiang, Zhao Wu, Yiping Bao, Zhicheng Wang, Si Liu, and Erjin Zhou. General instance distillation for object detection. In *CVPR*, pages 7842–7851, 2021. 2
- [15] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, pages 6202–6211, 2019. 6, 7
- [16] Nuno C Garcia, Pietro Morerio, and Vittorio Murino. Modality distillation with multiple stream networks for action recognition. In *ECCV*, pages 103–118, 2018. 2
- [17] Chengyue Gong, Dilin Wang, Meng Li, Vikas Chandra, and Qiang Liu. Vision transformers with patch diversification. *arXiv preprint arXiv:2104.12753*, 2021. 5, 6
- [18] Satya Krishna Gorti, Noël Vouitsis, Junwei Ma, Keyvan Golestan, Maksims Volkovs, Animesh Garg, and Guangwei Yu. X-pool: Cross-modal language-video attention for text-video retrieval. In *CVPR*, pages 5006–5015, 2022. 5
- [19] Byeongho Heo, Jeesoo Kim, Sangdoon Yun, Hyojin Park, Nojun Kwak, and Jin Young Choi. A comprehensive overhaul of feature distillation. In *ICCV*, pages 1921–1930, 2019. 2
- [20] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *Neurips Deep Learning and Representation Learning Workshop*, 2014. 2, 4
- [21] Berthold K.P. Horn and Brian G. Schunck. Determining optical flow. *Artificial Intelligence*, 17(1):185–203, 1981. 1
- [22] Y.-G. Jiang, J. Liu, A. Roshan Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://crcv.ucf.edu/THUMOS14/>, 2014. 2, 6
- [23] Hamid Reza Vaezi Joze, Amirreza Shaban, Michael L Iuzolino, and Kazuhito Koishida. Mmtm: Multimodal transfer module for cnn fusion. In *CVPR*, pages 13289–13299, 2020. 5
- [24] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *ICCV*, pages 5492–5501, 2019. 1
- [25] Hanjoo Kim, Minkyu Kim, Dongjoo Seo, Jinwoong Kim, Heungseok Park, Soeun Park, Hyunwoo Jo, KyungHyun Kim, Youngil Yang, Youngkwan Kim, et al. Nsm1: Meet the mlaas platform with a real-world case study. *arXiv preprint arXiv:1810.09957*, 2018. 8
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 6
- [27] Matthew Kowal, Mennatullah Siam, Md Amirul Islam, Neil DB Bruce, Richard P Wildes, and Konstantinos G Derpanis. A deeper dive into what deep spatiotemporal networks encode: Quantifying static vs. dynamic information. In *CVPR*, pages 13999–14009, 2022. 1
- [28] Pilhyeon Lee and Hyeran Byun. Learning action completeness from points for weakly-supervised temporal action localization. In *ICCV*, pages 13648–13657, 2021. 1
- [29] Pilhyeon Lee, Youngjung Uh, and Hyeran Byun. Background suppression network for weakly-supervised temporal action localization. In *AAAI*, pages 11320–11327, 2020. 1
- [30] Pilhyeon Lee, Jinglu Wang, Yan Lu, and Hyeran Byun. Weakly-supervised temporal action localization by uncertainty modeling. In *AAAI*, pages 1854–1862, 2021. 1
- [31] Yingwei Li, Yi Li, and Nuno Vasconcelos. Resound: Towards action recognition without representation bias. In *ECCV*, pages 513–528, 2018. 1
- [32] Yi Li and Nuno Vasconcelos. Repair: Removing representation bias by dataset resampling. In *CVPR*, pages 9572–9581, 2019. 1
- [33] Chuming Lin, Jian Li, Yabiao Wang, Ying Tai, Donghao Luo, Zhipeng Cui, Chengjie Wang, Jilin Li, Feiyue Huang, and Rongrong Ji. Fast learning of temporal action proposal

- via dense boundary generator. In *AAAI*, pages 11499–11506, 2020. [2](#)
- [34] 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 *CVPR*, pages 3320–3329, 2021. [1](#), [2](#), [3](#), [6](#), [8](#)
- [35] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *ICCV*, pages 7083–7093, 2019. [6](#), [7](#)
- [36] Tianwei Lin, Xiao Liu, Xin Li, Errui Ding, and Shilei Wen. Bmn: Boundary-matching network for temporal action proposal generation. In *ICCV*, pages 3888–3897, 2019. [1](#), [2](#), [8](#)
- [37] Tianwei Lin, Xu Zhao, Haisheng Su, Chongjing Wang, and Ming Yang. Bsn: Boundary sensitive network for temporal action proposal generation. In *ECCV*, pages 3–19, 2018. [1](#), [2](#), [3](#), [6](#), [8](#)
- [38] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. [3](#)
- [39] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *ECCV*, pages 21–37, 2016. [2](#)
- [40] Xiaolong Liu, Song Bai, and Xiang Bai. An empirical study of end-to-end temporal action detection. In *CVPR*, pages 20010–20019, 2022. [2](#), [3](#), [6](#), [8](#)
- [41] Xiaolong Liu, Yao Hu, Song Bai, Fei Ding, Xiang Bai, and Philip HS Torr. Multi-shot temporal event localization: a benchmark. In *CVPR*, pages 12596–12606, 2021. [6](#), [8](#)
- [42] 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. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [43] Yufan Liu, Jiajiong Cao, Bing Li, Chunfeng Yuan, Weiming Hu, Yangxi Li, and Yunqiang Duan. Knowledge distillation via instance relationship graph. In *CVPR*, pages 7096–7104, 2019. [2](#)
- [44] Fuchen Long, Ting Yao, Zhaofan Qiu, Xinmei Tian, Jiebo Luo, and Tao Mei. Gaussian temporal awareness networks for action localization. In *CVPR*, pages 344–353, 2019. [3](#), [8](#)
- [45] Zelun Luo, Jun-Ting Hsieh, Lu Jiang, Juan Carlos Niebles, and Li Fei-Fei. Graph distillation for action detection with privileged modalities. In *ECCV*, pages 166–183, 2018. [2](#)
- [46] Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *Fitnets: Hints for thin deep nets*. In *ICLR*, 2017. [2](#)
- [47] Sauradip Nag, Xiatian Zhu, Yi-Zhe Song, and Tao Xiang. Proposal-free temporal action detection via global segmentation mask learning. In *ECCV*, 2022. [1](#), [8](#)
- [48] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3967–3976, 2019. [2](#)
- [49] Zhiwu Qing, Haisheng Su, Weihao Gan, Dongliang Wang, Wei Wu, Xiang Wang, Yu Qiao, Junjie Yan, Changxin Gao, and Nong Sang. Temporal context aggregation network for temporal action proposal refinement. In *CVPR*, pages 485–494, 2021. [2](#)
- [50] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016. [2](#)
- [51] Sucheng Ren, Yong Du, Jianming Lv, Guoqiang Han, and Shengfeng He. Learning from the master: Distilling cross-modal advanced knowledge for lip reading. In *CVPR*, pages 13325–13333, 2021. [2](#)
- [52] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. [2](#), [4](#)
- [53] Dingfeng Shi, Yujie Zhong, Qiong Cao, Jing Zhang, Lin Ma, Jia Li, and Dacheng Tao. React: Temporal action detection with relational queries. In *ECCV*, 2022. [2](#), [8](#)
- [54] Han Shi, Jiahui Gao, Hang Xu, Xiaodan Liang, Zhenguo Li, Lingpeng Kong, Stephen Lee, and James T Kwok. Revisiting over-smoothing in bert from the perspective of graph. In *ICLR*, 2019. [6](#)
- [55] Zheng Shou, Jonathan Chan, Alireza Zareian, Kazuyuki Miyazawa, and Shih-Fu Chang. Cdc: Convolutional-deconvolutional networks for precise temporal action localization in untrimmed videos. In *CVPR*, pages 5734–5743, 2017. [2](#), [3](#), [8](#)
- [56] Zheng Shou, Dongang Wang, and Shih-Fu Chang. Temporal action localization in untrimmed videos via multi-stage cnns. In *CVPR*, pages 1049–1058, 2016. [1](#)
- [57] Changyong Shu, Yifan Liu, Jianfei Gao, Zheng Yan, and Chunhua Shen. Channel-wise knowledge distillation for dense prediction. In *ICCV*, pages 5311–5320, 2021. [2](#)
- [58] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Neurips*, pages 568–576, 2014. [1](#), [2](#)
- [59] Nako Sung, Minkyu Kim, Hyunwoo Jo, Youngil Yang, Jingwoong Kim, Leonard Lausen, Youngkwan Kim, Gayoung Lee, Donghyun Kwak, Jung-Woo Ha, and Sunghun Kim. Nsm1: A machine learning platform that enables you to focus on your models. *arXiv preprint arXiv:1712.05902*, 2017. [8](#)
- [60] Jing Tan, Jiaqi Tang, Limin Wang, and Gangshan Wu. Relaxed transformer decoders for direct action proposal generation. In *ICCV*, pages 13526–13535, 2021. [2](#), [3](#), [5](#), [6](#), [8](#)
- [61] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020. [2](#)
- [62] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Neurips*, 30, 2017. [6](#)
- [63] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, pages 20–36, 2016. [1](#)
- [64] Qiang Wang, Yanhao Zhang, Yun Zheng, and Pan Pan. Rcl: Recurrent continuous localization for temporal action detection. In *CVPR*, pages 13566–13575, 2022. [8](#)
- [65] Tao Wang, Li Yuan, Xiaopeng Zhang, and Jiashi Feng. Distilling object detectors with fine-grained feature imitation. In *CVPR*, pages 4933–4942, 2019. [2](#)

- [66] Andreas Wedel, Thomas Pock, Christopher Zach, Horst Bischof, and Daniel Cremers. An improved algorithm for tv-l1 optical flow. In *Statistical and geometrical approaches to visual motion analysis*, pages 23–45. Springer, 2009. 1
- [67] Xi Wei, Tianzhu Zhang, Yan Li, Yongdong Zhang, and Feng Wu. Multi-modality cross attention network for image and sentence matching. In *CVPR*, pages 10941–10950, 2020. 5
- [68] Kun Xia, Le Wang, Sanping Zhou, Nanning Zheng, and Wei Tang. Learning to refactor action and co-occurrence features for temporal action localization. In *CVPR*, pages 13884–13893, 2022. 8
- [69] Fanyi Xiao, Yong Jae Lee, Kristen Grauman, Jitendra Malik, and Christoph Feichtenhofer. Audiovisual slowfast networks for video recognition. *arXiv preprint arXiv:2001.08740*, 2020. 1
- [70] Junfei Xiao, Longlong Jing, Lin Zhang, Ju He, Qi She, Zongwei Zhou, Alan Yuille, and Yingwei Li. Learning from temporal gradient for semi-supervised action recognition. In *CVPR*, pages 3252–3262, 2022. 1, 3
- [71] Bo Xiong, Haoqi Fan, Kristen Grauman, and Christoph Feichtenhofer. Multiview pseudo-labeling for semi-supervised learning from video. In *ICCV*, pages 7209–7219, 2021. 3
- [72] Yuanjun Xiong, Yue Zhao, Limin Wang, Dahua Lin, and Xiaou Tang. A pursuit of temporal accuracy in general activity detection. *arXiv preprint arXiv:1703.02716*, 2017. 2
- [73] Mengmeng Xu, Juan Manuel Perez Rua, Xiatian Zhu, Bernard Ghanem, and Brais Martinez. Low-fidelity video encoder optimization for temporal action localization. *Neurips*, 34:9923–9935, 2021. 2
- [74] Mengmeng Xu, Chen Zhao, David S. Rojas, Ali Thabet, and Bernard Ghanem. G-tad: Sub-graph localization for temporal action detection. In *CVPR*, pages 10156–10165, 2020. 1, 3, 6, 7, 8
- [75] Haosen Yang, Wenhao Wu, Lining Wang, Sheng Jin, Boyang Xia, Hongxun Yao, and Hujie Huang. Temporal action proposal generation with background constraint. In *AAAI*, 2022. 6
- [76] Ke Yang, Peng Qiao, Dongsheng Li, Shaohe Lv, and Yong Dou. Exploring temporal preservation networks for precise temporal action localization. In *AAAI*, 2018. 2
- [77] Jun Yuan, Bingbing Ni, Xiaokang Yang, and Ashraf A Kasim. Temporal action localization with pyramid of score distribution features. In *CVPR*, pages 3093–3102, 2016. 2
- [78] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l1 optical flow. In Fred A. Hamprecht, Christoph Schnörr, and Bernd Jähne, editors, *Pattern Recognition*, pages 214–223. Berlin, Heidelberg, 2007. Springer Berlin Heidelberg. 1
- [79] Runhao Zeng, Wenbing Huang, Minghui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *ICCV*, pages 7094–7103, 2019. 1, 2, 8
- [80] Chenlin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *ECCV*, 2022. 1, 2, 3, 6, 7, 8
- [81] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, pages 11953–11962, 2022. 2
- [82] Chen Zhao, Ali K Thabet, and Bernard Ghanem. Video self-stitching graph network for temporal action localization. In *ICCV*, pages 13658–13667, 2021. 8
- [83] Long Zhao, Xi Peng, Yuxiao Chen, Mubbasir Kapadia, and Dimitris N Metaxas. Knowledge as priors: Cross-modal knowledge generalization for datasets without superior knowledge. In *CVPR*, pages 6528–6537, 2020. 2
- [84] Peisen Zhao, Lingxi Xie, Chen Ju, Ya Zhang, Yanfeng Wang, and Qi Tian. Bottom-up temporal action localization with mutual regularization. In *ECCV*, pages 539–555, 2020. 1, 2, 8
- [85] Yue Zhao, Yuanjun Xiong, and Dahua Lin. Recognize actions by disentangling components of dynamics. In *CVPR*, pages 6566–6575, 2018. 1
- [86] Yue Zhao, Yuanjun Xiong, Limin Wang, Zhirong Wu, Xiaou Tang, and Dahua Lin. Temporal action detection with structured segment networks. In *ICCV*, pages 2914–2923, 2017. 2
- [87] Zhaohui Zheng, Rongguang Ye, Ping Wang, Dongwei Ren, Wangmeng Zuo, Qibin Hou, and Ming-Ming Cheng. Localization distillation for dense object detection. In *CVPR*, pages 9407–9416, 2022. 2
- [88] Zixin Zhu, Wei Tang, Le Wang, Nanning Zheng, and Gang Hua. Enriching local and global contexts for temporal action localization. In *ICCV*, pages 13516–13525, 2021. 2